ANL/EVS-14/5

Water Use Optimization Toolset Project: Development and Demonstration Phase Draft Report

Environmental Science Division

About Argonne National Laboratory
Argonne is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC under contract DE-AC02-06CH11357. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne and its pioneering science and technology programs, see www.anl.gov.

DOCUMENT AVAILABILITY

Online Access: U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via DOE's SciTech Connect (http://www.osti.gov/scitech/)

Reports not in digital format may be purchased by the public from the National Technical Information Service (NTIS):
U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandra, VA 22312
www.ntis.gov
Phone: (800) 553-NTIS (6847) or (703) 605-6000
Fax: (703) 605-6900
Email: **orders@ntis.gov**

Reports not in digital format are available to DOE and DOE contractors from the Office of Scientific and Technical Information (OSTI):
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
www.osti.gov
Phone: (865) 576-8401
Fax: (865) 576-5728
Email: **reports@osti.gov**

Disclaimer
This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor UChicago Argonne, LLC, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, Argonne National Laboratory, or UChicago Argonne, LLC.

ANL/EVS-14/5

Water Use Optimization Toolset Project: Development and Demonstration Phase Draft Report

prepared by
John R. Gasper (lead author), Thomas D. Veselka, Matthew R. Mahalik,
John W. Hayse, and Samrat Saha
Argonne National Laboratory

Mark S. Wigmosta, Nathalie Voisin, Cynthia Rakowski, and Andre Coleman
Pacific Northwest National Laboratory

Thomas S. Lowry
Sandia National Laboratory

December 6, 2013

ACKNOWLEDGEMENTS

The authors would like to thank the U.S. Department of Energy (DOE) office of Energy Efficiency and Renewable Energy's Water Power Program for providing funding support for this project under the Advanced Water Power Program Announcement, DE-FOA-0000070.

CONTENTS

ACKNOWLEDGEMENTS iii

1 EXECUTIVE SUMMARY 1-1

1.1 Project Objective and Goal 1-1
1.2 Project Approach 1-1
1.3 Project Findings 1-2
1.4 Future Project Direction 1-5

2 PROJECT DESCRIPTION 2-1

2.1 Background and Overview 2-1
2.2 Model Development 2-1
2.2.1 Water Use Optimization Toolset Project Goals and Objective 2-1
2.2.2 WUOT Model Components 2-2
2.2.2.1 User Interface 2-2
2.2.2.2 Hydrologic Forecasting 2-3
2.2.2.3 Seasonal Hydrosystems Analysis 2-5
2.2.2.4 Environmental Performance 2-5
2.2.2.5 Day-Ahead Scheduling and Real-Time Scheduling 2-5
2.3 WUOT Site Demonstrations 2-5
2.3.1 WUOT Compatibility with Other Analytical Tools 2-7
2.3.2 Technical Review Team 2-7

3 WUOT COMPONENT CONCEPT AND DESIGN 3-1

3.1 User Interface 3-1
3.1.1 WUOT Concept 3-1
3.1.2 Hardware Requirements 3-1
3.1.3 Software Architecture 3-1
3.1.4 Managing WUOT Data 3-2
3.1.5 Building a Project 3-2
3.1.6 Database Design 3-6
3.1.7 Providing Tool Inputs 3-7
3.1.8 Invoking the Tools 3-8
3.1.9 Viewing Tool Results 3-8
3.2 Enhanced Hydrologic Forecasting System 3-13
3.2.1 Concept 3-13
3.2.2 Design 3-14
3.2.2.1 The Variable Infiltration Capacity Model 3-15
3.2.2.2 Forecast Products and WUOT Integration 3-22
3.3 Seasonal Hydrosystems Analysis 3-26

CONTENTS (Cont.)

3.3.1 Concept 3-26
3.3.2 Design 3-26
3.3.3 Mathematical Model 3-27
3.3.3.1 Heat Balance 3-28
3.3.3.2 Water Density 3-31
3.3.3.3 Withdrawal Dynamics 3-31
3.3.3.4 Inflow Dynamics 3-33
3.3.3.5 Density Mixing 3-34
3.3.3.6 Diffusion Mixing 3-35
3.3.3.7 Volume/Area/Elevation Curves 3-36
3.3.3.8 River Routing 3-36
3.3.3.9 Inflow Temperatures 3-38
3.3.3.10 Optimization 3-38
3.3.4 Data Requirements 3-44
3.3.5 Model Output 3-45
3.4 Environmental Performance 3-49
3.4.1 Concept 3-49
3.4.2 Design 3-49
3.4.2.1 IRF Modes 3-52
3.4.2.2 Algorithm Components of IRF Tool 3-53
3.4.2.3 Calculating IRF Score Using IRF Planning Mode 3-59
3.4.2.4 Integration of the IRF Tool in the WUOT 3-60
3.4.2.5 Calculating IRF Score Using IRF Historic Mode 3-64
3.4.3 IRF Tool Development 3-64
3.5 Day-Ahead and Real-Time Scheduling 3-65
3.5.1 Overview of Physical Processes and Modeling Challenges 3-65
3.5.2 Concept 3-70
3.5.3 Components and Modeling Processes 3-70
3.5.4 Design 3-71
3.5.4.1 Objective Function for Demonstration Sites 3-74
3.5.5 System Networks and Building Blocks 3-75
3.5.5.1 Boundary Node 3-75
3.5.5.2 Storage Node 3-76
3.5.5.3 Conversion Node 3-78
3.5.5.4 Links 3-81
3.5.5.5 Junction Node 3-83
3.5.6 Simulation Timeline 3-84
3.5.6.1 Time Constructs 3-85
3.5.6.2 Network Balance and Flow 3-86
3.5.6.3 Constraints 3-86
3.5.6.4 Sub-BU Time Considerations 3-88
3.5.6.5 Inflow Forecast Uncertainty 3-92

CONTENTS (Cont.)

3.5.6.6 Solving the Optimization Problem 3-95
3.5.6.7 Application of Nonlinear Functions and Relationships 3-95
3.5.6.8 Models of the Nonlinear Relationships and the Validation Process 3-96
3.5.6.9 Reservoir Surface Elevation 3-96
3.5.6.10 Tailwater Elevation 3-97
3.5.6.11 Turbine Power Conversion Efficiency 3-97
3.5.6.12 Iteration Process to Estimate Head Value 3-98
3.5.6.13 Nonlinear Validation Process for Ancillary Services 3-99
3.6 References 3-100

4 WUOT DEMONSTRATIONS 4-1

4.1 Concept 4-1
4.2 Design 4-1
4.2.1 Protocol for Conduct of Water Use Optimization Tool Set Site Demonstrations 4-1
4.2.1.1 Phase 1: Pre-Demonstration Actions 4-1
4.2.1.2 Phase 2: Initial Demonstration 4-2
4.2.1.3 Phase 3: Site Demonstration Reporting 4-3
4.3 Demonstration Site Descriptions 4-3
4.3.1 Oroville-Thermalito Complex Demonstration Site 4-3
4.3.1.1 Hyatt PGP 4-6
4.3.1.2 TDD and Powerplant 4-6
4.3.1.3 Fish Barrier Dam and the Feather River Fish Hatchery 4-7
4.3.1.4 Thermalito Power Canal 4-8
4.3.1.5 Thermalito Forebay Dam, Forebay, and PGP 4-8
4.3.1.6 Thermalito Afterbay Dam and Afterbay 4-9
4.3.1.7 Feather River Flow and Reservoir Operating Criteria 4-10
4.3.1.8 Oroville-Thermalito Complex Resource Overview 4-13
4.3.2 Colorado River Storage Project Demonstration Site 4-16
4.3.2.1 Operating Constraints and Operational Guidelines 4-18
4.3.2.2 Glen Canyon Dam 4-18
4.3.2.3 Flaming Gorge and Fontenelle Dams 4-20
4.3.2.4 Gunnison Basin and the Aspinall Cascade 4-21
4.3.2.5 Collbran Project 4-25
4.3.2.6 Rio Grande Project 4-25
4.3.2.7 Small Projects 4-25
4.3.2.8 Operational Business Process 4-26
4.3.2.9 Long-Term Planning and Purchases 4-26
4.3.2.10 Daily Scheduling 4-27
4.3.2.11 Real-Time Scheduling 4-27

CONTENTS (Cont.)

4.3.2.12 Management and Marketing Office Objectives 4-28
4.3.2.13 Ancillary Services 4-30
4.3.2.14 Reclamation Unit Dispatch Guidelines 4-31
4.4 Toolset Demonstration Results and Discussion 4-32
4.4.1 Aspinal Cascade 4-32
4.4.1.1 Hydroforecasting 4-32
4.4.1.2 HydroSCOPE 4-47
4.4.1.3 Environmental Performance 4-69
4.4.1.4 Stand-Alone CHEERS Retrospective Analysis of CRSP Day-Ahead and Real-Time Scheduling 4-91
4.4.1.5 CHEERS Retrospective Analysis in an Integrated WUOT Framework 4-139
4.4.1.6 CRSP Retrospective Summary 4-143
4.4.2 Oroville Complex 4-145
4.4.2.1 Hydroforecasting 4-145
4.4.2.2 Seasonal Model 4-153
4.4.2.3 Environmental Performance 4-153
4.4.2.4 Application of CHEERS for Oroville-Thermalito Complex Day-Ahead and Real-Time Scheduling 4-202
4.5 References 4-212

APPENDIX A: WUOT NATIONAL LABORATORY TEAM STAFF A-1

APPENDIX B: TECHNICAL REVIEW TEAM MEMBERS B-1

APPENDIX C: ACRONYMS AND ABBREVIATIONS C-1

APPENDIX D: DATA REQUIREMENTS D-1

APPENDIX E: EQUATION NOMENCLATURE E-1

APPENDIX F: ADAPTIVELY REFINED DYNAMIC PROGRAM FOR LINEAR SPLINE F-1

APPENDIX G: PREDICTED WATER-TO-POWER CONVERSION EFFICIENCY CURVES FOR HYDROPOWER GENERATING UNITS IN THE ASPINALL CASCADE G-1

APPENDIX H: HEAD LOSS APPROXIMATIONS AT BLUE MESA AND MORROW POINT H-1

CONTENTS (Cont.)

APPENDIX I: HOURLY CHEERS PRICES I-1

APPENDIX J: CUSTOMIZED INFLOW FORECAST ERROR RESERVOIR BUFFERS J-1

APPENDIX K HISTORICAL CRYSTAL RESERVOIR OPERATING VIOLATIONS K-1

APPENDIX L COMPARISONS OF CHEERS MILP AND NONLINEAR POST-PROCESSOR RESULTS L-1

APPENDIX M STREAMFLOW ASSIMILATION IN EHFS M-1

FIGURES

2-1 Water Use Optimization Toolset Conceptual Design 2-2

2-2 Tool Set Integration 2-3

2-3 Oroville Complex, Feather River; and Colorado River Storage and Seedskadee Projects 2-4

2-4 WUOT Project Demonstration Sites 2-6

3-1 Project Manager Screen 3-3

3-2 Node Creation Screen 3-4

3-3 Sample Network Using All Node Types 3-6

3-4 Form to Invoke and Monitor the Day-Ahead and Real-Time Tool 3-9

3-5 Node Annotations Screen 3-10

3-6 Result Annotations Screen 3-10

3-7 Result Data Viewed on Network Window with Timeline Slider 3-11

3-8 Time Series Graph Screen 3-12

FIGURES (Cont.)

3-9 Time Series Table Screen 3-12

3-10 Locations of Seasonal Ensemble Forecasts from the University of Washington/Princeton University Forecast Model 3-15

3-11 Component Diagram of the Forecasting Sequence 3-16

3-12 Representation of Water and Energy Balance in a VIC Model Grid Cell 3-17

3-13 Daily Medium-Range Forecast for the Feather River at Oroville on 4/1/2013 3-21

3-14 Day-Ahead Ensemble Flow Forecast for Water Year 1998 over the Feather River Basin in Oroville, California 3-22

3-15 Seasonal Flow Forecast for Inflow into the Oroville Reservoir, California 3-23

3-16 The HydroSCOPE Simulation and Optimization Tool Conceptual Model 3-27

3-17 The Cone of Withdrawal for Three Outlet Structures at Different Elevations in a Stratified Water Column 3-33

3-18 Geometric Configuration for Each Test Case and Reach 3-40

3-19 Test Case 1 Comparison Plot of the Inflow Hydrograph and the Discharge from Reach 1, Calculated by the EMDE Method and a HEC-RAS Model 3-41

3-20 Test Case 2 Comparison Plot of the Inflow Hydrograph and the Discharge from Reach 3, Calculated by the EMDE Method and a HEC-RAS Model 3-41

3-21 Test Case 3 Plot of EMDE Calculated Discharge from Reach 2 Compared to the USGS Gauge 09144250 Hydrograph 3-42

3-22 Test Case 3 Plot of EMDE Calculated Discharge from Reach 3 Compared to the HEC-RAS Model and the USGS Gauge 09152500 Hydrograph 3-42

3-23 Communication Structure of the HydroSCOPE-DAKOTA Framework 3-43

3-24 Illustration of the Daily, Weekly, and Monthly Multipliers Used for Optimizing the Reservoir Release Schedules 3-44

3-25 Calibrated Vertical Temperature Profiles for the Three Reservoirs of the Aspinall Cascade on the Gunnison River 3-47

FIGURES (Cont.)

3-26 Plot of Maximum and Minimum Daily Temperatures for HydroSCOPE and Data from USGS Gauge 09152500 for the 6-month Period from April through September 2009; Plot Showing the Daily Variation Produced from HydroSCOPE as Compared to the Daily Maximum and Minimum Temperatures from the Gauge 3-48

3-27 Screenshot of the IRF Tool Depicting the "Objectives" Tab, Which Allows Users to Define Environmental Objectives 3-50

3-28 Hypothetical Example to Illustrate the Relationship between Environmental Performance and Power Revenue 3-61

3-29 Hypothetical Example Showing the Results of Optimization after 1,000 Simulations 3-62

3-30 Layout Depicting Data Needs and Sources for IRF Tool and Their Exchange with the Other WUOT Components during Optimization 3-63

3-31 The IRF Tool Is Coded Using the C# Programming Language and Interacts with an SQL Database 3-65

3-32 Reservoir Elevation Level as a Function of Water Storage 3-67

3-33 Tailrace Elevation as a Function of Outflow Rate 3-67

3-34 Maximum Turbine Flow Rate as a Function of Head 3-68

3-35 Turbine Efficiency as a Function of Both Turbine Output and Head 3-69

3-36 Relationship between Reservoir Volume and Elevation 3-77

3-37 Hypothetical SOS2 Curve Relating Generating Unit Power Output to Turbine Flow Rate 3-80

3-38 Hypothetical Water Travel Time Distribution Function 3-82

3-39 Illustration of a Wave Function Modeled as TTD 3-83

3-40 Major CHEERS Time Construct Classifications and Sub-period Classifications 3-85

3-41 CHEERS Constraint Classifications over Time and Space 3-87

FIGURES (Cont.)

3-42 Operating Range with and without Ancillary Services 3-90

3-43 Regulation Service Impacts on BU Generation Levels 3-91

3-44 Ramping Impacts on BU Generation Levels 3-92

3-45 Illustration of CHEERS Elevation Buffers 3-93

3-46 Forecast Error Probability Distribution Function 3-94

3-47 Power Output Curve with Ramping and Regulation Service 3-100

4-1 Location of the Oroville-Thermalito Complex in California..... 4-4

4-2 Map of the Oroville-Thermalito Complex 4-5

4-3 Historical Monthly Average Water Storage Levels for Lake Oroville..... 4-13

4-4 Historical Monthly Average Water Storage Levels for Thermalito Reservoirs 4-14

4-5 Elevation Characteristics of the Oroville-Thermalito Complex Reservoirs 4-14

4-6 Oroville-Thermalito Complex Powerplant Capacities and Historical Generation Levels 4-15

4-7 SLCA/IP Power Plants..... 4-17

4-8 Map of the Aspinall Cascade and Surrounding Area..... 4-21

4-9 Aspinall Cascade Reservoir Storage Capacities and Power Plant Generating Capacities 4-22

4-10 Routing Network for the Gunnison River Basin..... 4-34

4-11 Top: 1980–2000 Calibrated Monthly Flow Climatology at Blue Mesa, with Observed Naturalized Flow, Observed Regulated Flow, and EHFS Simulation; Middle: Mean Monthly Flow for 2001–2010 Validation Period; Bottom: Mean Monthly Flow for the 2009 Demonstration Period..... 4-36

4-12 The EHFS Nowcast and Observed Flow at Blue Mesa for 2009 4-38

FIGURES (Cont.)

4-13 Observed Inflow at Blue Mesa for 2009 and EHFS 1-day Lead Ensemble Forecast without Streamflow Assimilation or Post-processing 4-38

4-14 Relationship between Simulated Streamflow and Spatially Averaged Model Soil Moisture in Lowest Layer for the Area Upstream of Blue Mesa in the Gunnison River Basin..... 4-39

4-15 Effect of Streamflow Assimilation on the Model's Soil Moisture State for Three Subbasins in the Gunnison River Basin..... 4-40

4-16 Observed Flow at Blue Mesa for 2009 and EHFS 1-day Lead Ensemble Forecast with Streamflow Assimilation, without Post-processing..... 4-41

4-17 Observed Flow at Blue Mesa for 2009 and EHFS 1-day Lead Ensemble Forecast with Data Assimilation and Post-processing 4-42

4-18 Observed Flow at Blue Mesa for 2009 and EHFS 2-day Lead Ensemble Forecast with Data Assimilation and Post-processing 4-43

4-19 Observed Flow at Blue Mesa for 2009 and EHFS 3-day Lead Ensemble Forecast with Data Assimilation, with Post-processing..... 4-44

4-20 12-month Seasonal Ensemble Flow Forecast for Blue Mesa Issued April 1, 2009, with Forecast Flows to March 30, 2010.....4-446

4-21 12-month Seasonal Ensemble Flow Forecast for Grand Junction Issued April 1, 2009, with Forecast Flows to March 30, 2010 4-45

4-22 Evaluation of the April 1, 2009, EHFS Volumetric Probabilistic Forecast with Respect to the Equivalent Forecast Issued by the NWS Colorado Basin River Forecast Center 4-46

4-23 Conceptualization of the HydroSCOPE Model for the Aspinall Cascade Demonstration Site..... 4-48

4-24 Plot of the Blue Mesa Target Elevations and the Actual Elevations for 2009 4-50

4-25 Plot of the Power Generation Curve for Crystal Reservoir 4-54

4-26 Historical Monthly Electricity Prices for April 2009 through March 2010 and the Average Monthly Prices Used in the Retrospective Analysis 4-55

FIGURES (Cont.)

4-27 Schematic for Implementing the Penalty Function 4-56

4-28 Plots of the Tradeoff between the Total Revenue and the IRF Score Produced by DAKOTA during an Optimization Run 4-57

4-29 Plots Showing the Range of Variability of Flows Sampled by DAKOTA during the Optimization at the Gunnison Gauge 4-59

4-30 Plots Showing the Range of Variability of Flows Sampled by DAKOTA during the Optimization at the Whitewater Gauge 4-60

4-31 Optimized and Historical Releases for Blue Mesa Reservoir 4-61

4-32 Optimized and Historical Releases for Morrow Point Reservoir 4-62

4-33 Optimized and Historical Releases for Crystal Reservoir 4-63

4-34 Optimized and Historical Releases at the Gunnison Gauge 4-64

4-35 Optimized and Historical Releases at the Whitewater Gauge 4-65

4-36 Differences in the Daily Volumes between the Optimized and Historical Operations for Blue Mesa Reservoir 4-66

4-37 Differences in the Daily Volumes between the Optimized and Historical Operations for Morrow Point Reservoir 4-67

4-38 Differences in the Daily Volumes between the Optimized and Historical Operations for Crystal Reservoir 4-68

4-39 Relationship between Minimum Daily Discharge at the Whitewater Gauge and Suitability for Fish Larval Entrainment into Floodplain Areas in the Lower Gunnison River 4-71

4-40 Relationship between Within-Day Change in Stage at Whitewater and Suitability for Juvenile Fish in the Lower Gunnison River 4-71

4-41 Relationship between Mean Hourly Discharge and Availability of Backwater Habitat for Endangered Fish in the Lower Gunnison River 4-72

FIGURES (Cont.)

4-42 Relationship between Instantaneous Discharge and Suitability of Conditions as Indicated by Objective Points Associated with Minimum Peak Flow Requirements in the Lower Gunnison River 4-75

4-43 Relationship between Mean Hourly Discharge and Availability of Trout Spawning Habitat for Endangered Fish in the Black Canyon and Gunnison Gorge Reaches of the Gunnison River 4-77

4-44 Objective Environmental Performance Score and Site Scores for a 16-year Historic Period for the Gunnison River Using the IRF Historic Tool 4-80

4-45 Historic Hydrographs for the Two Sites in the Gunnison River during a 16-year Period 4-81

4-46 Objective Performance Score for Objectives 1.1–1.13 for a 16-year Period for ESWA Site in the Gunnison River Using the IRF Historic Tool 4-82

4-47 Objective Performance Score for Objectives 2.1–2.9 for a 16-year Period for BCGG Site in the Gunnison River Using the IRF Historic Tool 4-83

4-48 Daily IRF and Site Environmental Performance Scores Based on Observed Facility Operation and IRF Integrated WUOT Optimized Simulation for the Gunnison River Using the IRF Planning Tool for April 1, 2009, to September 30, 2009 4-85

4-49 Daily Objective Performance Score for Objectives 1.1–1.13 based on Optimized WUOT Simulation for the Gunnison River Using the IRF Planning Tool for April 1, 2009, to September 30, 2009, at the ESWA Site 4-86

4-50 Daily Objective Performance Score for Objectives 2.1–2.9 based on Optimized WUOT Simulation for the Gunnison River Using the IRF Planning Tool for April 1, 2009, to September 30, 2009, at the BCGG Site 4-87

4-51 Actual Observed and WUOT Optimized Hydrographs for the Two Sites in the Gunnison River from April 1, 2009, to September 30, 2009 4-88

4-52 Comparison of Daily Objective Performance Score for Objective 1.12 between Observed Facility Operation and Optimized WUOT Simulation for the Gunnison River Using the IRF Planning Tool for April 1, 2009, to September 30, 2009, at the ESWA Site 4-89

FIGURES (Cont.)

4-53 CHEERS Model Run Sequence and Data Flows for Stand-Alone Retrospective Analyses 4-93

4-54 CRSP Supply and Demand Energy Balance 4-96

4-55 Screen Capture of the CRSP Daily Scheduling Sheet 4-97

4-56 Morrow Point Power Plant Day-Ahead and Actual Unit Generation Level for April 20, 2009, and April 21, 2009 4-97

4-57 CHEERS Network Developed for the CRSP Demonstration Site 4-98

4-58 Aspinall Reservoir Graphics Added to CRSP Daily Scheduling Sheets 4-101

4-59 Blue Mesa Reservoir Elevation as a Function of Water Storage Volume 4-102

4-60 Morrow Point Reservoir Elevation as a Function of Water Storage Volume 4-103

4-61 Crystal Reservoir Elevation as a Function of Water Storage Volume 4-103

4-62 Blue Mesa Tail Water Elevation as a Function of Water Release Rate 4-105

4-63 Morrow Point Tail Water Elevation as a Function of Water Release Rate 4-105

4-64 Crystal Tail Water Elevation as a Function of Water Release Rate 4-106

4-65 Power Output as a Function of Water Release Rate, April 2009 4-108

4-66 Power Conversion Factor as a Function of Water Release Rate, April 2009 4-109

4-67 Power Output as a Function of Water Release Rate, April 2009 4-110

4-68 Power Conversion Factor as a Function of Water Release Rate, April 2009 4-110

4-69 Piecewise Linear Curves for Power Output and Power Conversion Factor, Blue Mesa 4-113

4-70 Piecewise Linear Curves for Power Output and Power Conversion Factor, Morrow Point 4-113

FIGURES (Cont.)

4-71 Crystal Unit 1 Power Outputs as a Function of Turbine Release Rate and Head.....4-115

4-72 Market Prices Used for the CHEERS Retrospective Analysis 4-116

4-73 Actual and CHEERS Real-Time CRSP System Energy Balances 4-118

4-74 Actual CRSP and CHEERS Real-Time Power Market Transactions.....4-118

4-75 Comparison of Actual Aspinall Plant-Level Generation with CHEERS Real-Time Optimization Results.....4-120

4-76 Comparison of Actual Apsinall Unit-Level Generation with CHEERS Real-Time Results.....4-122

4-77 CHEERS Aspinall Water Releases and Reservoir Elevations.....4-124

4-78 Actual and CHEERS Real-Time CRSP System Energy Balances 4-125

4-79 Actual CRSP and CHEERS Real-Time Power Market Transactions.....4-126

4-80 Comparison of Day-Ahead Scheduled and Actual Apsinall Plant-Level Generation with CHEERS Real-Time Results 4-127

4-81 Comparison of Actual Apsinall Unit-Level Generation with CHEERS Real-Time Results.....4-128

4-82 May 4, 2009, Actual and CHEERS CRSP System Energy Balances under Prescheuled Hydropower Projections and Actual Real-Time Conditions.....4-130

4-83 Comparison of Actual Apsinall Unit-Level Generation with CHEERS Results.....4-132

4-84 CHEERS Model Run Sequence and Data Flows for WUOT Integrated Retrospective Analyses.....4-140

4-85 Comparison of Actual Daily Water Release with Day-Ahead Prescheduled Level and Optimized HydroSCOPE Amounts 4-142

4-86 Comparison of Actual Inflow/Side Flows with Persistence and EHFS Forecasts.....4-144

4-87 Feather River Basin, VIC Routing, and the Forecast Point at Oroville.....4-146

FIGURES (Cont.)

4-88 Monthly Flow at Oroville, California, during the Calibration and Validation Periods for Observed Naturalized Flows and Modeled Flows 4-148

4-89 Daily Naturalized Observed Flow and EHFS Simulated Flow at Oroville, California 4-149

4-90 Observed Flow at Oroville for 2011 and EHFS 1-, 2-, and 3-day Lead Ensembles.....4-151

4-91 Seasonal Ensemble Forecast at Oroville, Issued April 1, 2011, with Forecast Flows to March 30, 2012 4-152

4-92 Evaluation of the April 1, 2011, EHFS April–July Volumetric Probabilistic Forecast for the Feather River Basin 4-152

4-93 Conceptual Model for the Seasonal Tool.....4-154

4-94 Relationship between Flows at the THA Gauge and the GRL Gauge.....4-155

4-95 The Multiplier Function for the Releases from Oroville 4-156

4-96 The Perfect Forecast Supplied by the EHFS Includes Naturalized Flows That Must Be Adjusted to Match the Water Balance in the Reservoir; a Linear Transformation Is Used to Make the Adjustment.....4-158

4-97 Difference in the Forecasted Inflow Volume for a 6-month Period from Each Date before and after the Adjustment Was Made.....4-159

4-98 The Penalty Multiplier is a Function of the Cumulative Time That All the Objectives Are out of Compliance.....4-161

4-99 Elevation Objectives for the Afterbay during the Retrospective Analysis Simulation Period.....4-162

4-100 Historical Elevation for Oroville as Compared to the High-Flow and Low-Flow Storage Requirements 4-163

4-101 Schematic of an Eight-Shutter Intake System with Four Shutters Still in Place 4-164

4-102 Outflow Rates are Calculated for the Mid-elevation of Each Open Shutter Using an Empirical Representation of the BOR Experimental Results 4-165

4-103 Plot of the Power Generation Curve for the Hyatt Power Plant 4-167

FIGURES (Cont.)

4-104 Plot of the Power Generation Curve for the Thermalito Forebay....4-168

4-105 Plots Showing the Range of Variability of the Total Releases from Oroville Sampled by DAKOTA during the Optimization Process....4-170

4-106 Plots of the Optimization Results for Selected Dates of the Analysis....4-172

4-107 The Difference between the Optimized Temperature and the Historical Temperature at Robinson's Riffle....4-173

4-108 Plot of the Daily Result Releases, the Water Balance Releases, and the Corresponding Water Surface Elevations for Oroville....4-175

4-109 Plot of the Daily Volume Difference for Releases from Oroville of the Daily Results and the WB Results....4-176

4-110 Relationship between Instantaneous Depth of Water Column and Suitability Points for Chinook Salmon Spawning in the High Flow Channel Section of the Feather River....4-179

4-111 Relationship between Instantaneous Depth of Water Column and Suitability Points for Steelhead Trout Spawning in the High Flow Channel Section of the Feather River....4-180

4-112 Discharge Stage Relationship Developed for the Low Flow Channel of Feather River Using Three-Segment Piecewise Linear Regression....4-190

4-113 Discharge Stage Relationship Developed for the High Flow Channel of Feather River Using Three-Segment Piecewise Linear Regression....4-191

4-114 Actual Observed Hydrographs for the HFC Site in the Feather River, California, from 1986 to 2012....4-192

4-115 Annual Site IRF Scores and Objective Performance Scores for Objectives 1.1–1.8 for a 17-year Period at the Feather River HFC Site Using the IRF Historic Tool....4-193

4-116 Daily IRF and Site Environmental Performance Scores Based on Observed Facility Operation and IRF Integrated WUOT Optimized Simulation for the Feather River Using the IRF Planning Tool for April 1, 2011, to September 30, 2011....4-195

FIGURES (Cont.)

4-117 Daily Objective Performance Score for Objectives 1.1–1.8 Based on Optimized WUOT Simulation for the Feather River Using the IRF Planning Tool for April 1, 2011, to September 30, 2011, at the HFC Site 4-196

4-118 Comparison of Daily Objective Performance Score for Objective 1.4 between Observed Facility Operation and Optimized WUOT Simulation for the Feather River Using the IRF Planning Tool for April 1, 2011, to September 30, 2011, at the HFC Site 4-198

4-119 Actual Observed and WUOT Optimized Hydrographs for the Two Sites in the Feather River from April 1, 2011, to September 30, 2011 4-199

4-120 Daily Objective Performance Score for Objectives 2.1–2.3 based on Optimized WUOT Simulation for the Feather River Using the IRF Planning Tool for April 1, 2011, to September 30, 2011, at the LFC Site 4-200

4-121 Relative Improvement in Generation and Revenue of CHEERS Model Scenarios over the Historic Record for CDWR Operations May–September 2011 4-211

H-1 Blue Mesa Unit 1 Head and Generation Relationship for Different Water Release Rates H-1

H-2 Blue Mesa Unit 1 Head Loss Approximation, April 2009, Min Head = 320.77 ft, Max Head = 334.27 ft H-2

H-3 Morrow Point Unit 1 Head and Generation Relationship for Different Water Release Rates H-3

H-4 Morrow Point Unit 1 Head Loss Approximation, April 2009, Min Head = 388.99 ft, Max Head = 404.83 ft H-4

I-1 Process for Estimating Hourly WECC Market Hub Prices I-1

I-2 Process for Estimating Total Hourly Loads in the U.S. Portion of WECC I-2

I-3 Illustration of Load Duration Curve Shaping to Match a Target Load Factor I-3

I-4 Original and Shaped Chronological Load Curve I-3

J-1 Forecast Error versus Exceedance Probability J-1

FIGURES (Cont.)

K-1 Stretches of Time during the Retrospective Analysis in which Historic Information Indicates Reservoir Constraints Were Violated at Crystal, Together with Inflow/ Sideflow Amounts and Reservoir Elevation Levels K-2

L-1 Comparison of Hourly Aspinall Reservoir Elevations Computed by the CHEERS MILP and the Nonlinear Routine L-2

L-2 Comparison of Hourly Unit-Level Generation Levels Computed by the CHEERS MILP and the Nonlinear Routine L-3

M-1 Relationship between Simulated Streamflow and Spatially Averaged Model Soil Moisture in Layer 3 for the Area Upstream of Blue Mesa in the Gunnison Basin..... M-2

M-2 Effect of Streamflow Assimilation on the Model Soil Moisture State for Three Subbasins in the Gunnison River Basin M-4

TABLES

3-1 Forecast Point Lookup Is Manually Established to Provide a Relationship between the On-the-Ground Latitude/Longitude and the EHFS Model Cell Latitude/Longitude 3-24

3-2 Listing of Reach Geometries and Hydraulic Properties for Testing the EMDE River Routing Method 3-39

3-3 Listing of Time Variant Input Variables to HydroSCOPE 3-45

3-4 Listing of Static Input Variables to HydroSCOPE 3-46

4-1 Feather River HFC Minimum Flow Requirements 4-11

4-2 Oroville-Thermalito Complex Reservoir Operating Criteria 4-12

4-3 Operating Constraints Prior to 1991 and under the ROD 4-19

4-4 Characteristics of Aspinall Cascade Resources 4-23

TABLES (Cont.)

4-5 Forecast Point Lookup Manually Established to Provide a Relationship between the Specified Nodes of the Toolset System, the EHFS Routing Network at a One-Eighth-Degree Spatial Resolution, and On-the-Ground Latitude/Longitude of Existing Gauge Stations Whenever Possible 4-33

4-6 Monthly Streamflow Simulation Performance in the Gunnison River Basin at Blue Mesa 4-35

4-7 Maximum and Minimum Operating Elevations for the Aspinall Cascade 4-53

4-8 Baseline and Optimized Values and Percent Change for the IRF Score and Total Revenue 4-58

4-9 Polynomial Coefficients for the Reservoir Elevation Functions 4-104

4-10 Polynomial Coefficients for the Tail Water Elevation Function 4-106

4-11 Polynomial Coefficients for the Conversion Factor Function, Blue Mesa, April 2009 4-108

4-12 Polynomial Coefficients for the Blue Mesa Power Output Function, April 2009 4-109

4-13 Polynomial Coefficients for the Power Output Function, Morrow Point, April 2009 4-111

4-14 Polynomial Coefficients for the Conversion Factor Function, Blue Mesa, April 2009 4-111

4-15 Piecewise Linear Curve Breakpoints, Blue Mesa 4-114

4-16 Piecewise Linear Curve Breakpoints, Morrow Point 4-114

4-17 Intercepts and Slopes for Crystal Head-Generation Linear Approximations 4-115

4-18 Comparison of Actual Purchase and Sales Quantities with CHEERS Real-Time Results 4-119

4-19 Comparison of Actual Purchase Costs and Sales Revenues with CHEERS Real-Time Results 4-120

4-20 Daily Summary of Actual and CHEERS Real-Time Generation 4-123

TABLES (Cont.)

4-21 Comparison of Actual Purchase and Sales Quantities with CHEERS Real-Time Results 4-126

4-22 Comparison of Actual Purchase Costs and Sales Revenues with CHEERS Real-Time Results 4-127

4-23 Comparison of Actual Generation with CHEERS Real-Time Results 4-128

4-24 Comparison of Actual Purchase and Sales Quantities with CHEERS Real-Time Results 4-131

4-25 Comparison of Actual Purchase Costs and Sales Revenues with CHEERS Real-Time Results 4-131

4-26 Comparison of Actual Generation with CHEERS Results 4-133

4-27 Comparison of Monthly Net Costs 4-135

4-28 Monthly Comparison of Actual Generation with CHEERS Results 4-136

4-29 Monthly Comparison of Actual Start/Stop Cycles CHEERS Results 4-136

4-30 Monthly Comparison of Actual Ramping with CHEERS Results 4-138

4-31 CHEERS Stand-Alone Mode Run Sensitivity Analysis 4-138

4-32 Comparision of Actual and HydroSCOPE Monthly Water Releases 4-143

4-33 Comparision of Power System Metric for Actual and CHEERS WUOT Integrated Model Results Assuming a One Unit Start/Stop Limit Per Day 4-145

4-34 The Forecast Point Lookup is Manually Established to Provide a Relationship between the One-Eighth-Degree Routing Network, the Node in the Integrated Toolset, and Corresponding On-the-Ground Latitude/Longitude When Available 4-147

4-35 Daily Streamflow Simulation Performance in the Feather River Basin at Oroville, California 4-148

4-36 Power Calculation Coefficients for the Hyatt Power Plant 4-166

4-37 Power Calculation Coefficients for the Thermalito Forebay 4-168

TABLES (Cont.)

4-38 Power and IRF Scores and the Percent Improvement for Selected Dates of the Analysis Period 4-173

4-39 Power and IRF Scores of the Daily Results and Water Balance Release Schedules 4-176

4-40 Constraints Related to Water for the Fish Hatchery Site 4-187

4-41 Discharge, Elevation, and Stage Values for Two Transects in the Low Flow Channel of Feather River 4-189

4-42 Discharge, Elevation, and Stage for Two Transects in the High Flow Channel of Feather River 4-189

4-43 Computed Bed Elevations at Four Cross-Sections in the High and Low Flow Sections of the Feather River 4-190

4-44 Absolute Average Error between Forecast and Actual Hourly Price for Energy and Ancillary Services 4-208

H-1 Intercepts and Slopes of Head-Generation Linear Approximations, Blue Mesa Unit 1 ..... H-2

H-2 Intercepts and Slopes of Head-Generation Linear Approximations, Morrow Point Unit 1 ..... H-3

L-1 Comparison of Total Daily Unit-Level Generation Levels Computed by the CHEERS MILP and the Nonlinear Routine L-4

**WATER USE OPTIMIZATION TOOLSET PROJECT:
DEVELOPMENT AND DEMONSTRATION PHASE DRAFT REPORT
DECEMBER 6, 2013**

1 EXECUTIVE SUMMARY

This report summarizes the results of the development and demonstration phase of the Water Use Optimization Toolset (WUOT) project. It identifies the objective and goals that guided the project, as well as demonstrating potential benefits that could be obtained by applying the WUOT in different geo-hydrologic systems across the United States.

A major challenge facing conventional hydropower plants is to operate more efficiently while dealing with an increasingly uncertain water-constrained environment and complex electricity markets. The goal of this 3-year WUOT project, which is funded by the U.S. Department of Energy (DOE), is to improve water management, resulting in more energy, revenues, and grid services from available water, and to enhance environmental benefits from improved hydropower operations and planning while maintaining institutional water delivery requirements. The long-term goal is for the WUOT to be used by environmental analysts and deployed by hydropower schedulers and operators to assist in market, dispatch, and operational decisions.

1.1 PROJECT OBJECTIVE AND GOAL

The objectives of the project are to develop and demonstrate an integrated set of advanced analytical tools that would allow hydropower operators and planners to optimize power generation, economics, water management, and environmental performance and to demonstrate the effectiveness of the integrated toolset at hydropower sites that are representative of different operational, market structures, and environmental conditions across the United States.

1.2 PROJECT APPROACH

The WUOT, which was developed and demonstrated by a team of DOE National Laboratory staff (Appendix A) and advised by a technical review committee of hydropower operators (Appendix B), is a suite of advanced, analytical tools that integrate water, power, and environmental performance. The WUOT was designed to simulate key factors affecting hydropower operations, including water availability, short- and long-term water and power demands, and environmental performance. The WUOT allows users to evaluate impacts from natural and anthropogenic changes in water flows within river/reservoir systems as they pertain to hydropower generation, revenues, and environmental performance. The WUOT optimization simultaneously considers, over time, the interactions among all water and power resources, hydropower economics, ancillary services, customer loads, and environmental constraints. The

WUOT is comprised of four models linked via a graphical user interface (GUI) and common database. These models include the Enhanced Hydrologic Forecasting System (EHFS), which projects water availability, HydroSCOPE, which simulates long-term reservoir and river conditions, power production, and revenue; the Index of River Functionality (IRF), which calculates the environmental performance associated with a time series of hydropower operations for various habitats across a basin; and the Conventional Hydropower Energy and Environmental Resource System (CHEERS), which develops near-term schedules for power generation, ancillary services, and water release.

The WUOT is being demonstrated via retrospective analyses comparing historical operational performance at multiple U.S. sites with toolset-projected performance. Power marketers, day-ahead and real-time schedulers, and plant operators provided system configuration and detailed input data, along with feedback on model design and performance for these assessments. The WUOT is also being applied by operators to assist in evaluating operations at several of the demonstration sites.

1.3 PROJECT FINDINGS

Retrospective analyses have been completed for two of the demonstration sites, the Aspinall Cascade in the Colorado River Storage Project and the Oroville-Thermalito Complex on the Feather River in California. For the analyses, staff representing each of the demonstration sites chose a 6-month historical period of interest. Toolset components were run in both integrated and stand-alone modes for those periods and the results were compared to actual operating performance. Model results suggest that alternative operational regimes potentially would have improved the value of demonstration site resources to the grid while enhancing or not degrading environmental performance and complying with water-related obligations for non-power uses.

The EHFS provided medium-range and seasonal ensemble forecasts for the retrospective analyses. For the Aspinall Cascade demonstration, in the Gunnison River Basin, streamflow forecasts were provided at Blue Mesa Dam, Morrow Point Dam, tributary inflow between Blue Mesa and Morrow Point, Crystal Dam, tributary inflow between Morrow Point and Crystal Dam, the North Fork of the Gunnison just above the confluence with the Gunnison, the Uncompahgre above the confluence with the Gunnison, Delta, and Grand Junction. For the Oroville-Thermalito Complex demonstration in the Feather River Basin, streamflow forecasts were provided at Oroville.

The EHFS provided four major products for use by the other WUOT models, including the following:

- *Seasonal Climate Forecast Forcings*. The seasonal forcing group includes meteorological forcing data at one-eighth-degree spatial grid resolution for a horizon of 366 days. This hourly meteorological forecast dataset provides the

sub-daily variability of the weather required by the seasonal optimization tool, HydroSCOPE.

- *Seasonal Streamflow Forecast*. The seasonal forecast group includes an ensemble of daily data (49 traces + average trace) that is issued once every 7 days out to 1 year and includes not only the daily flow at specific forecast points, but also hydrologic fluxes.

- *Medium-Range Meteorological Forecasts*. The medium-range ensemble meteorological forecasts are derived from the pre-processed Global Forecast System (GFS) medium-range weather forecasts. It includes one-eighth-degree gridded, 13-day, 15-member, hourly meteorological data, issued daily for the period of 1990–2011. In this configuration, the hourly disaggregation is based on the solar zenith angle; however, the daily dataset has prediction skill.

- *Medium-Range Hydrologic Forecasts*. The medium-range forecast group is structured similarly to the seasonal forecast group; however, the medium-range forecast group only includes forecasts out to 13 days for 15 ensemble traces and an average trace for the period of evaluation.

Using output from EHFS, HydroSCOPE was employed for the retrospective analysis to project successive, daily optimizations over the 6-month analysis periods of each site (183 optimizations for each site). The daily release volumes for the first day of each optimization were provided to the day-ahead planning model (CHEERS), which then optimized those releases over the next 24 hours. This “telescopic” approach to optimization means that seasonal objectives are implicitly included in the daily optimization process and is the primary reason behind the WUOT’s integrated design.

For both study sites, the seasonal model results show that over their respective 6-month analysis periods there is little room for improvement with regard to their environmental performance, but there is a 3–4% potential improvement for their power performance. This indicates that the current set of downstream flow and elevation targets at each site limits the ability to influence environmental performance by only changing releases from the reservoirs. For the Aspinall Cascade, utilizing sets of ensemble forecasts in a manner that reduces the operational downside may be the most effective method for increasing the system’s long-term performance; this helps operators avoid conditions that degrade environmental performance while maintaining or improving power performance. For the Oroville Complex, there may be a potential for improvement in both power generation and environmental performance if the optimization were to include optimizing the Oroville intake shutter positions and/or optimizing the system during the winter months to maximize system flexibility in the summer months.

In the retrospective analyses, the IRF was engaged to evaluate the environmental performance of operational schedules developed by other WUOT tools. At the Colorado River Storage Project’s (CRSP’s) Aspinall Cascade on the Gunnison River, integrated results indicate an overall modest improvement in environmental performance under the hypothetical optimized

operation relative to the actual observed facility operation during the same period. The improvement stems from modification of the reservoir releases, whereby the optimized flow regime was able to enhance environmental performance by reshaping the hydrograph so that overall the objectives are accomplished to a greater degree. This allowed the accomplishment of an objective related to the maintenance of spawning habitat for endangered fish to its full potential in the optimized schedule, whereas in the observed operation this objective was accomplished at a level that was approximately 30% less than the optimized operation. It is important to note that overall environmental performance did not deteriorate during the optimization process, and that at the same time a significant economic benefit occurred.

At the Oroville-Thermalito Complex on the Feather River, overall environmental performances (as indicated by the IRF scores) were better for almost all days during the modeling period, compared to the observed operation during the same period. From an individual objective perspective, the greatest improvement in environmental performance was observed in an objective pertaining to accomplishing Chinook salmon outmigration by releasing a spring pulse-flow. The level of accomplishment for this objective was approximately 20% higher in the optimized schedule than in the observed operation. Even though improvement in overall environmental performance in the optimized operation was somewhat moderate, it is important to note that the enhancement in environmental condition was not accomplished by compromising on the power-generation, and that this took place while meeting the power- and water-delivery commitments of the facility and respecting infrastructural and downstream flood-control constraints.

In evaluating short-term performance over the 6-month period of the retrospective analysis, CHEERS performed daily runs to optimize day-ahead planning and real-time operations on an hourly scale. The model runs were configured to reflect the current business practices and dispatch guidelines used by the actual schedulers. Using WUOT-generated daily release volumes (HydroSCOPE) and measures of environmental performance (IRF), CHEERS developed optimized hydropower water release and dispatch operation schedules; CHEERS also developed methods for economic balancing of other loads and resources including long- and short-term power contracts, interchanges, and transactions on the spot market. Results were compared to historical operational performance.

The current planners and dispatchers have decades of experience in efficient operations of the systems within many economic, environmental and regulatory constraints. The retrospective analysis indicates that in general their performance has been excellent. However, analyses of CHEERS model results also indicate that there are some areas that could be fine-tuned to further improve operations. For example, at the Aspinall Cascade demonstration site in the CRSP, hourly operations of three particular days can provide an understanding of the potential for improvement. One day deals with an early-spring, energy-short hydropower condition, and the other day with a late-spring, energy-long hydropower condition. CHEERS results indicate that the Aspinall units could be operated at higher levels of efficiency, with less ramping, and with fewer starts and stops. Furthermore, other system operations can be adjusted to reduce costly purchases of balancing energy from the variable spot market, especially, as detailed in the third day's examination, when it becomes necessary to cope with an inaccurate

inflow forecast in real time. Model results from using CHEERS in both stand-alone and integrated modes and using several variations of operational rules and system assumptions suggest that over the 6-month retrospective analysis period, WUOT-optimized schedules could save an estimated $500,000 in system costs while producing an additional 6 GWh of energy, all while releasing the same amount of water that was released historically.

The Oroville-Thermalito Complex demonstration has an operational goal to maximize revenue from the sale of energy and ancillary services. Due to the business-sensitive nature of CDWR operating information in the CAISO market, the model results and historical operation information cannot be presented in this report at the same level of detail as those for the CRSP. Results do indicate that while CDWR schedulers are already operating units at high efficiency points and making excellent use of limited water resources, there are significant potential gains in revenue to be made. Several other variations of WUOT-optimized operational scenarios demonstrate opportunities for improved values as a result of higher unit availability, more accurate market price forecasts, different unit operating guidelines, and the production and sale of ancillary services.

Results of the demonstrations indicate that the WUOT can be a useful means to help manage risks associated with hydrological uncertainty and adverse environmental impacts, reduce costs, increase unit efficiency and plant capacity factors, and enhance the economic value of hydropower resources.

1.4 FUTURE PROJECT DIRECTION

Future direction of the project includes refining individual tools and the user interface based on feedback from demonstration site users and the review committee of hydropower experts. Future plans also include the development of additional WUOT documentation and technology transfer materials, along with the training and transfer of the model to demonstration site staff.

2 PROJECT DESCRIPTION

2.1 BACKGROUND AND OVERVIEW

A major challenge facing conventional hydropower plants is operating more efficiently while dealing with increasingly uncertain water-constrained environments and complex electricity markets. This challenge was recognized by the DOE's Energy Efficiency and Renewable Energy Water Power Program, which held a series of discussions and a workshop in 2003 with hydropower stakeholders to define related research and development needs. One of these was the need for an optimization tool that went beyond the capability of existing tools to integrate power generation, water management, and environmental performance to facilitate more effective and efficient hydropower operations.

Under a subsequent competitive procurement (Advanced Water Power Program Announcement, DE-FOA-0000070, topic area 3, "Supporting Research and Testing for Hydropower"), a team of DOE National Laboratories (Argonne National Laboratory [Argonne], Oak Ridge National Laboratory [ORNL], Pacific Northwest National Laboratory [PNNL], and Sandia National Laboratories [SNL]) was charged with developing and demonstrating such a tool. The resulting WUOT project is described in this report. The WUOT allows evaluation of operational efficiency by projecting energy production, ancillary services, revenues generated, and environmental performance scores resulting from user-defined operational scenarios. The WUOT differs from other currently available tools since it integrates water forecasting, reservoir and stream flow routing, power system, and environmental simulation models (Figure 2-1) into a single integrated package. Although the WUOT was designed to be operated as a toolset, each of its components can be run individually or in conjunction with non-WUOT tools. The WUOT was demonstrated at two sites with widely ranging hydropower operations and geo-hydrologic conditions: the Oroville-Thermalito Complex on Feather River in California, and the Colorado River Storage Project (CRSP) in the Colorado River Basin. Results for a third demonstration site, the Conowingo Dam Complex on the Susquehanna River, will be the subject of a separate report. The development and demonstration of the WUOT was guided by a technical review team of hydropower operators and planners.

2.2 MODEL DEVELOPMENT

2.2.1 Water Use Optimization Toolset Project Goals and Objective

The goal of this 3-year WUOT project is to improve water management, resulting in more energy, revenues, and grid services from available water, and to enhance environmental benefits from improved hydropower operations and planning while maintaining institutional water delivery requirements. The long-term goal is for the WUOT to be used by environmental analysts and deployed by hydropower schedulers and operators to assist in market, dispatch, and operational decisions.

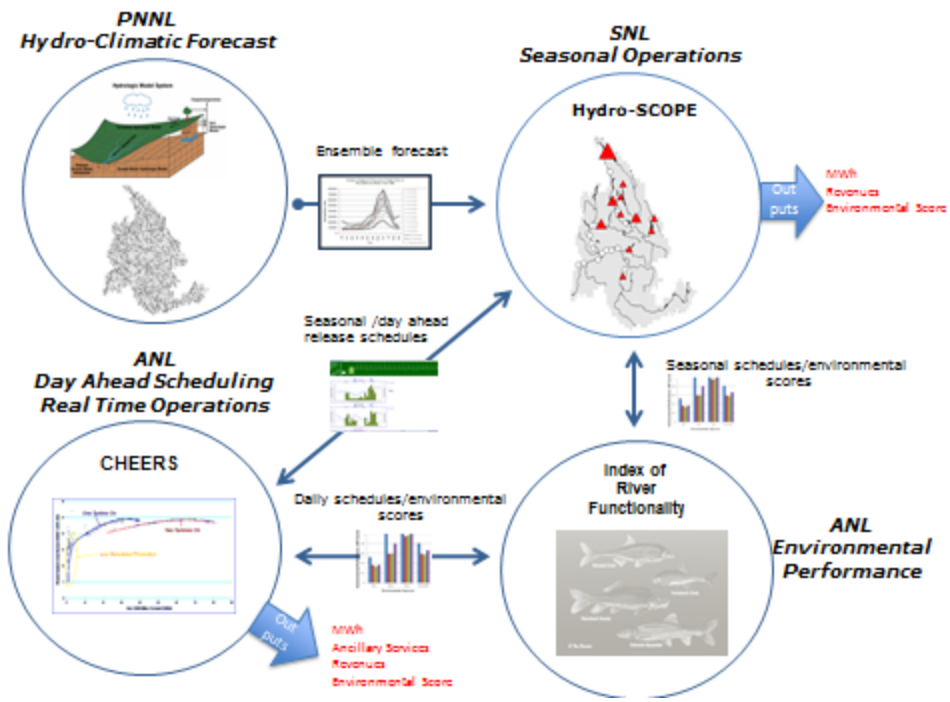

FIGURE 2-1 Water Use Optimization Toolset Conceptual Design

The objectives of the project are to develop and demonstrate an integrated set of advanced analytical tools that would allow hydropower operators and planners to optimize power generation, economics, water management, and environmental performance, as well as to demonstrate the effectiveness of the integrated toolset at hydropower sites that are representative of different operational, market structure, and environmental conditions across the United States.

2.2.2 WUOT Model Components

The toolset consists of four analytical tools plus a graphical user interface and shared database.

2.2.2.1 User Interface

A graphical user interface (GUI) provides a way for the user to define and characterize the hydro-system under evaluation. It communicates with a common database that is shared by all tools and contains both data evaluation and display software. Figure 2-2 shows toolkit component interactions with the GUI and common database. Using the GUI, the user creates a topology or network for the hydropower system being analyzed and enters related characteristics and performance data (Figure 2-3). The user-defined network includes objectives, system relationships, and constraints, as well as several options for managing time and spatial boundary conditions that allow analysis at any level of granularity consistent with the application, project

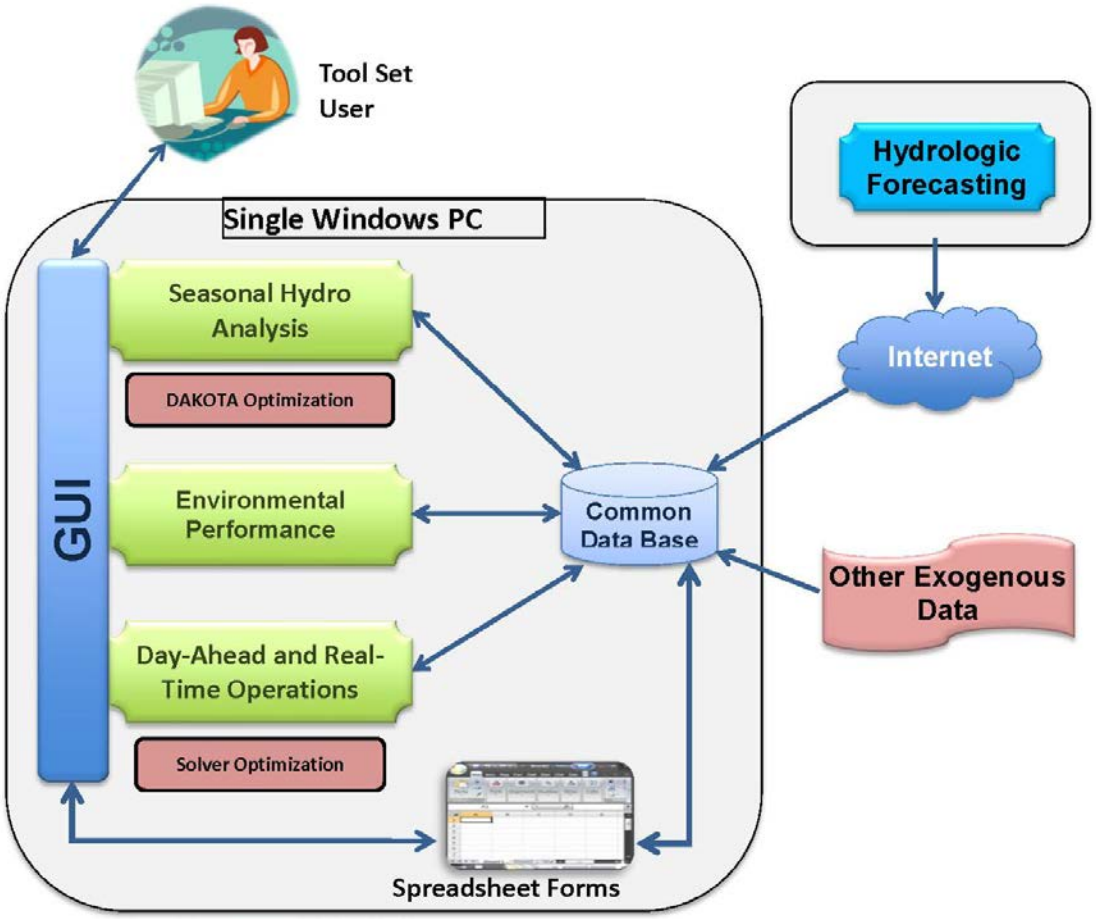

FIGURE 2-2 Tool Set Integration

resources, data availability, and required simulation accuracy. A common database is used to store and access input and output information shared among tools. This insures internally consistent inputs and outputs among toolset components and minimizes user resources required to operate the toolset. A set of data display and evaluation tools facilitate analysis of results.

2.2.2.2 Hydrologic Forecasting

The Enhanced Hydrologic Forecasting System (EHFS) component is a spatially distributed modeling system that provides daily to seasonal ensemble inflow forecasts for use by the Seasonal Hydrosystems Analysis, Day-ahead Scheduling and Real-time Operation and Environmental Performance Analysis tools. The tool reduces forecast uncertainty through the use of remotely sensed spatial data, high-resolution meteorological forecasts, and real-time stream flow updating; this allows improved system operations by reducing the need for overly conservative projections of water availability.

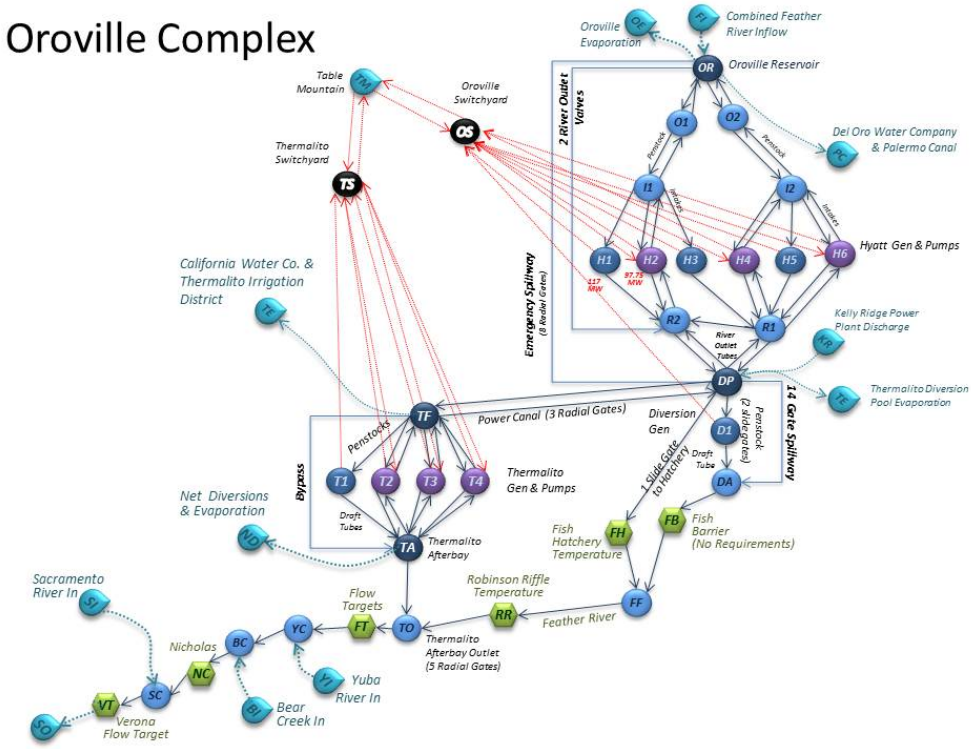

(a)

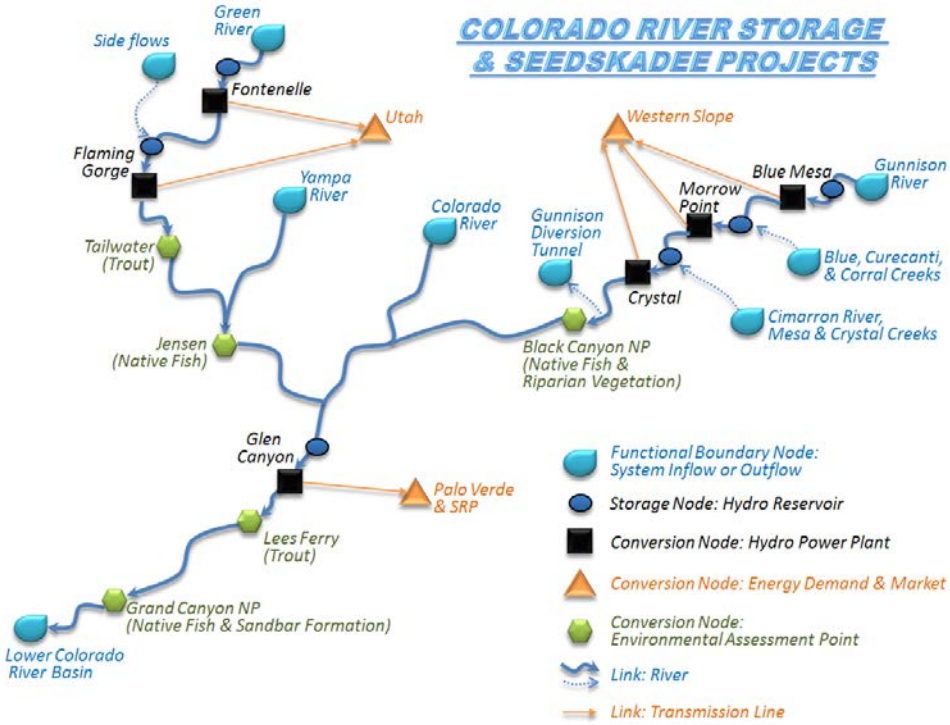

(b)

FIGURE 2-3 (a) Oroville Complex, Feather River; and (b) Colorado River Storage and Seedskadee Projects

2.2.2.3 Seasonal Hydrosystems Analysis

Hydro-SCOPE is a coupled one-dimensional reservoir and river routing model that simulates reservoir and river temperatures, power production, and revenue, as well as downstream river conditions, as a function of inflows, meteorological conditions, and power and water demand. Hydro-SCOPE includes multiple objective optimization for evaluating tradeoffs between operational and environmental factors. The tool allow users to balance seasonal and multi-seasonal forecasts of energy demand and water availability/water demand against power generation capacities, operational constraints, competing water users, and environmental performance.

2.2.2.4 Environmental Performance

The Index of River Functionality (IRF) component incorporates environmental objectives into the toolset by computing the environmental performance measures associated with a time series of hydropower operations for various habitats at specific locations across a basin. IRF scores typically are a function of flow, habitat, and population dynamics. Environmental performance scores are developed based on how well river conditions accomplish user-defined environmental objectives (e.g., fish spawning, fish growth) in terms of timing, magnitude, duration, and frequency of occurrence. The IRF allows users to evaluate differences in the environmental performance of various operating scenarios.

2.2.2.5 Day-Ahead Scheduling and Real-Time Scheduling

The Conventional Hydropower Energy and Environmental Resource System (CHEERS) modeling system creates schedules for power generation, ancillary services (regulation up and down, spin reserves, and non-spinning reserves), and water releases. These schedules are driven by multiple objectives, simultaneously solving for energy and environmental goals.

2.3 WUOT SITE DEMONSTRATIONS

One design parameter for the project was that the WUOT must be sufficiently flexible to be utilized in different geo-hydrologic basins and power marketing structures across the United States. The project team canvassed hydropower operators across the United States to identify candidate sites, and four sites were chosen to demonstrate WUOT flexibility: the Oroville-Thermalito Complex on Feather River in California, the Aspinall Cascade in the upper portion of the Colorado River Basin, the Conowingo Dam Complex on the Susquehanna River on the Maryland/Delaware border, and the Skagit River Hydroelectric Project on the Skagit River in Washington. These sites represent a wide range in environmental and hydropower operational conditions and are located in different power market conditions (Figure 2-4). The organizations responsible for operating these systems agreed to provide the resources necessary to conduct the

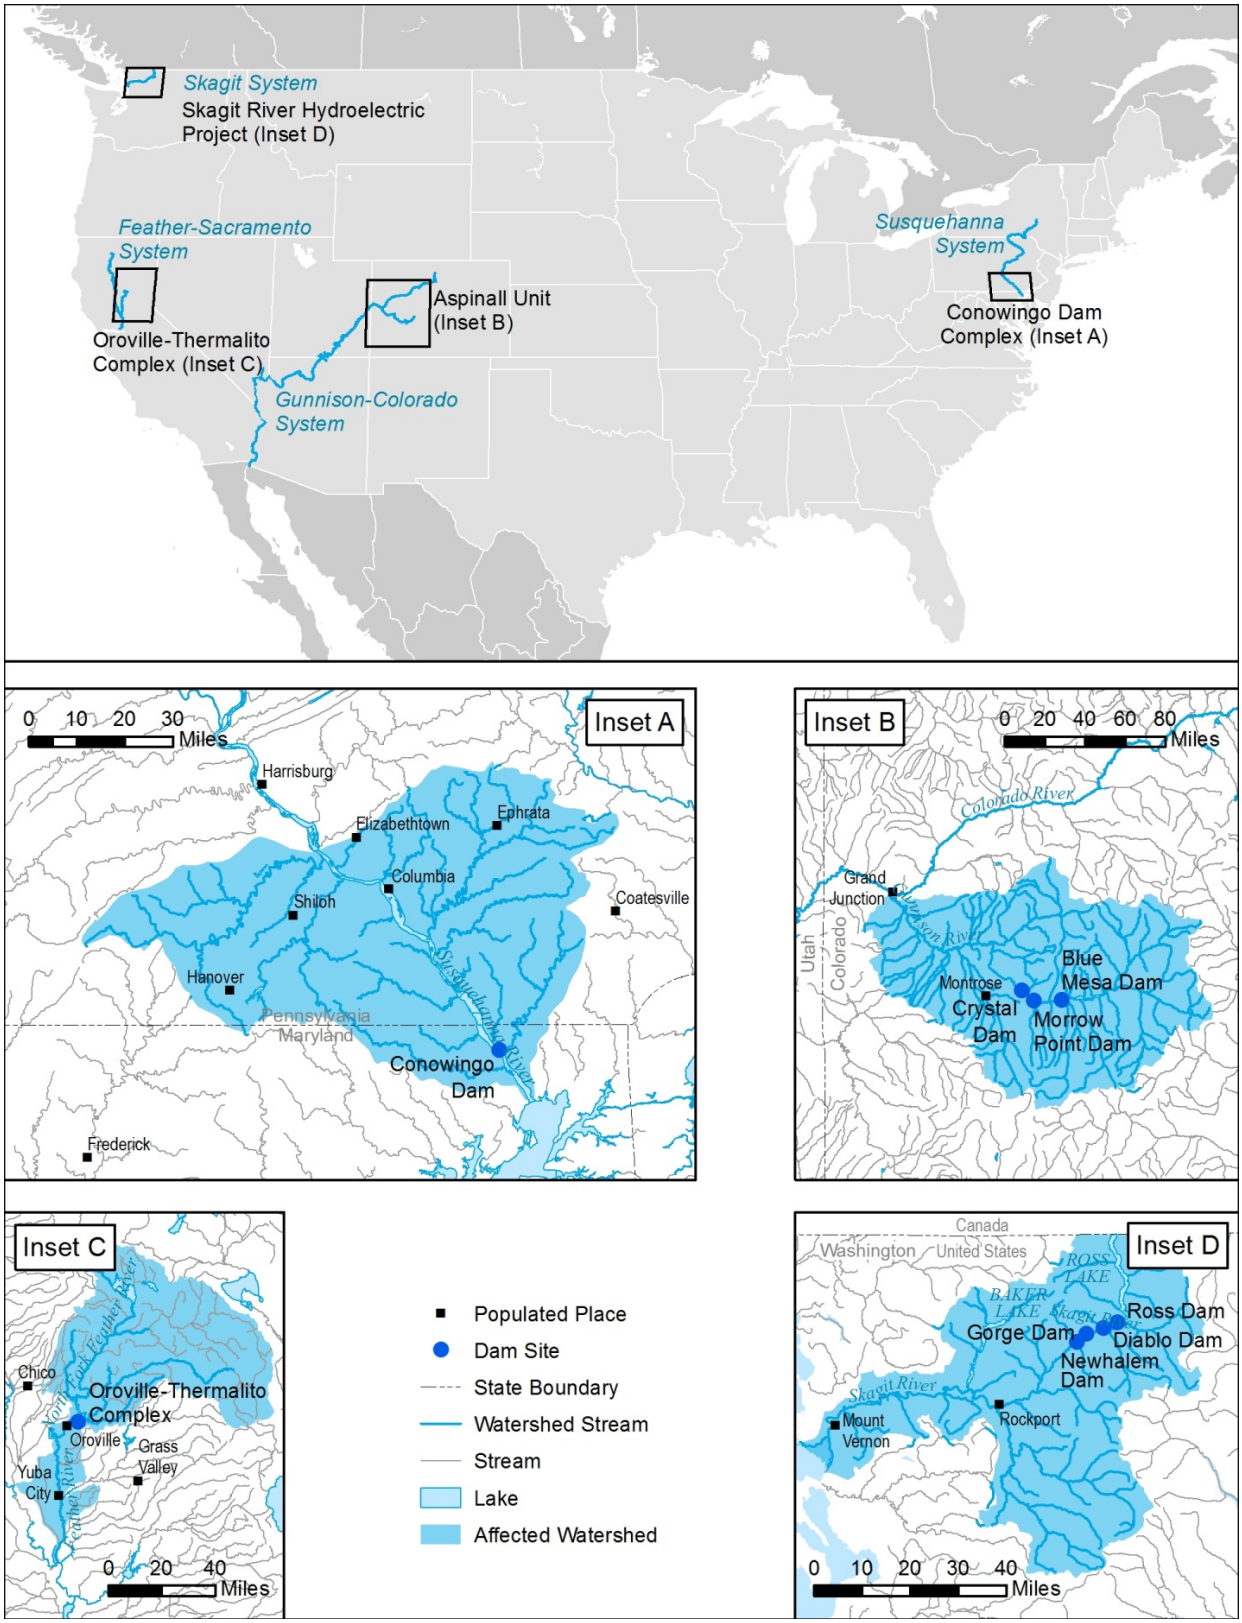

FIGURE 2-4 WUOT Project Demonstration Sites

demonstrations. Demonstration site operators became members of the project's technical review team, providing system design information, input data and review of results; they are integrally involved in all phases of the demonstration process, including the following:

1. Designing models,
2. Establishing model performance objectives/measures,
3. Creating site topologies,
4. Defining baseline and test conditions,
5. Identifying data needs and availability, and
6. Reviewing results.

Setup and operation of the WUOT requires staff resources to develop site topologies, identify and obtain input and characterization information, populate the common database, make model runs, and evaluate model outputs. The extent and nature of these requirements is dependent on the complexity of the hydropower system being studied and on the level of analysis to be conducted.

2.3.1 WUOT Compatibility with Other Analytical Tools

One design parameter of the WUOT was that it be readily compatible with other hydropower planning and optimization tools. Although no commercial tools were identified that meet the WUOT design parameter of integrating power, water management, and environmental performance, there are commercial and proprietary tools (Appendix A) that address topics and issues that are similar to WUOT toolset components. Many hydropower organizations have invested resources in the development and application of these tools. The WUOT is designed to build on these investments by accessing data from or providing data to other tools via data bridges to the WUOT common database. For example, if an operator already employs an alternate river routing tool, reservoir elevation and river flow data could easily be imported to CHEERS and the environmental tool via the common database. Alternatively, inflow forecasting information could be imported to a detailed dispatch model.

2.3.2 Technical Review Team

Development and demonstration of the WUOT was guided by a technical review team (Appendix B). The technical review team is comprised of representatives from organizations participating in the site demonstrations and from other key hydropower stakeholders. The study team canvassed industry, governmental, and other hydropower stakeholder organizations to identify staff with the interest, availability, and capability to participate in the technical review

team. Extensive familiarity with hydropower systems and analytical models, as well as the ability to participate in demonstration of the WUOT, were key considerations in the selection of members. Throughout the project, technical review team members provided reviews of toolset component concepts, designs, and results, as well as access to demonstration sites and related data.

3 WUOT COMPONENT CONCEPT AND DESIGN

3.1 USER INTERFACE

3.1.1 WUOT Concept

The objective of the WUOT user interface is to provide the hydropower operator or planner with a means to define and characterize the components of hydropower system under evaluation, select and run the WUOT tools, and view and evaluate WUOT results. The framework developed to meet this objective consists of a graphical user interface (GUI) through which the user designs the system network; accesses each of the analytical tools; stores information in a common database that provides information shared by the toolset components; and interacts with a set of data analysis and visualization tools for evaluating model results.

3.1.2 Hardware Requirements

All of the WUOT components can be installed on and executed from a single Windows personal computer (PC), with the exception of the hydrologic forecasting tool. Due to the intense level of computational power required by that tool it operates on a server at PNNL, automatically performing a forecasting operation once per day, the results of which can be downloaded into the WUOT via the Internet. There are no specific hardware requirements for the PC on which WUOT is installed, and any modern PC will be capable of running the software. However, as with any other computationally demanding software, stronger hardware will result in faster performance. Parts of both the WUOT interface and the individual tools take advantage of multi-core central processing units (CPUs) when available. The software has been tested on 32-bit Windows XP and 64-bit Windows 7 operating systems. Although they have not been specifically tested, other versions of Windows are likely compatible.

3.1.3 Software Architecture

The WUOT interface has been developed using the Microsoft .NET platform and is written in C#. WUOT can be installed on a PC simply by placing a folder in the PC's file system that contains the program's executable files and other important files and folders. The PC on which WUOT is installed is required to have the Microsoft .NET Framework 4 installed. A typical PC likely already has the framework installed, but if not it can be downloaded for free from Microsoft. One mechanism by which the WUOT user can enter data into the system is through dynamic Excel workbooks, which are invoked from the interface. Therefore, the PC must also have installed a copy of Microsoft Excel 2010. A MySQL database is used to hold all of the WUOT data. MySQL is free software that can be downloaded from the Internet and installed on any PC. Typically, the MySQL service is installed and run on the same PC as WUOT, but this does not have to be the case. In the configuration screen the user sees when

opening the interface, the user provides connection parameters for the MySQL server to be used for the current session. The server specified can be located on the local PC or a on a remote PC. This ability allows multiple WUOT users in different locations to share the same database on a non-concurrent basis; the current version of the software is not designed for simultaneous use by multiple users. Via this interface, the user is also able to specify connection parameters to a second MySQL database, the database on the PNNL server that holds the hydrologic forecasting tool results. Access to this database is controlled via a username and password supplied by PNNL, as well as by an IP address filter in place on the PNNL server.

The day-ahead and real-time tool requires the use of a mixed-integer linear program solver software. Many such solvers exist, both commercial and non-commercial. While the day-ahead and real-time tool has been designed such that it may be integrated with other solvers for a minimal amount of effort by the WUOT developers, it is currently compatible only with the commercial LINGO 12 solver, and therefore a copy of that solver must also be installed on the PC. For its optimization routine, the seasonal tool requires the use of the DAKOTA (Design Analysis Kit for Optimization and Terascale Applications) software, which can be downloaded and installed for free from SNL.

3.1.4 Managing WUOT Data

The WUOT tools are designed such that they can be operated completely independently of one another, or in an integrated mode in which any combination of tools works together and shares inputs and results. Regardless of the mode of operation, all of the tools use a common database schema, meaning that all of the data for all of the tools is contained together in a single entity. This entity or schema is referred to as a “project.” A project contains all of the data for a specific application of WUOT. For example, if a WUOT user is managing one hydropower system in California, and a separate hydropower system in Colorado, then it is likely that the user will have represented each system in a separate project. The first screen the user sees when opening the interface is the Project Manager screen (shown in Figure 3-1), from which a new project can be created or an existing project can be opened, copied, renamed, or deleted. When creating a new project, the user has the option to create a completely blank project, or to start with a user-defined template that already has basic information in place about the system to be modeled. The Project Manager also provides the option to export the entire contents of a project into a single plain text file (referred to as a “dump”) that can be imported into any WUOT instance at a later time. This feature facilitates the easy backup, archival, or transfer of the project.

3.1.5 Building a Project

Each of the four WUOT tools has a specific purpose, and to achieve its purpose each may only need to deal with a particular type of information, for only selected hydropower system components, and over a certain time horizon using a certain time step. For example, the

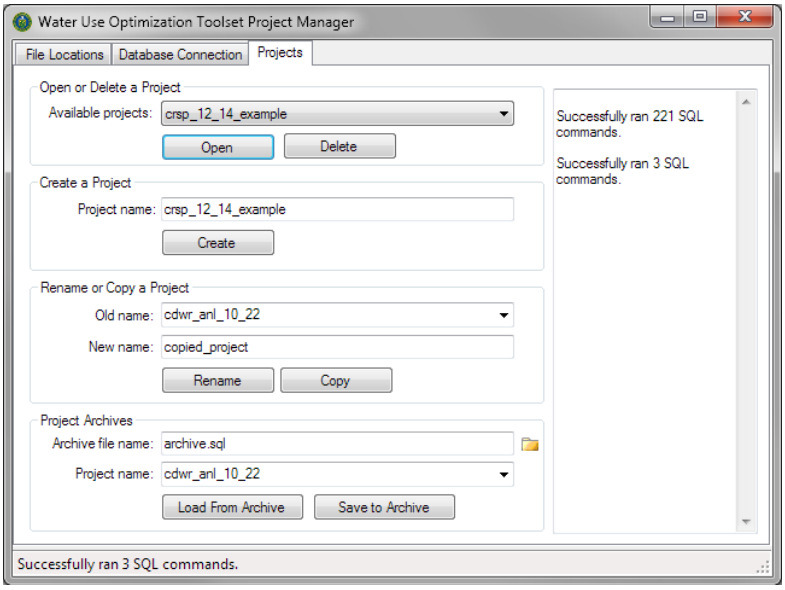

FIGURE 3-1 Project Manager Screen

hydrologic forecasting tool is concerned only with water inflows at certain points in the system, is not concerned about anything to do with electricity production, and produces a daily forecast for a length of 1 year. By contrast, the day-ahead and real-time tool is concerned with both water and electricity at every point in the system, and typically produces an hourly schedule for only a few days. Even though each tool has a markedly different scope and focus, each uses the same network representation and supporting data. This design feature facilitates data consistency and sharing of both inputs and results among the tools.

Using the GUI, the user creates a network representing the hydropower system under study by placing nodes and links on a "blank canvas." Nodes represent individual network components (e.g., hydro-reservoirs, power plants, energy demand centers, river gauges). Links represent the connections among those components (e.g., rivers, bypass channels, transmission lines). Figure 3-2 shows the creation of a new node in the GUI. The user right-clicks on a blank area to bring up a context menu from which the "new node" option can be selected. After choosing that option, a new screen appears in which the user enters basic information about the node. Methods for entering more detailed information are accessed separately and discussed later. Links are created in a similar way. The visual appearance of the node or link can be modified as desired. Nodes or groups of nodes can be selected and dragged to other positions, and the view of the network can be zoomed in or out or panned in any direction.

A physical resource conveyed through the network to and from nodes and along links is referred to as a commodity. Commodities that are currently simulated at demonstration sites include water, electricity, and money. Each commodity has one or more associated attributes. For example, the water commodity might have attributes of flow rate, elevation, temperature, and dissolved oxygen content. These commodities with associated attributes flow into, through, and out of network nodes and links. Along the various network pathways a commodity may be

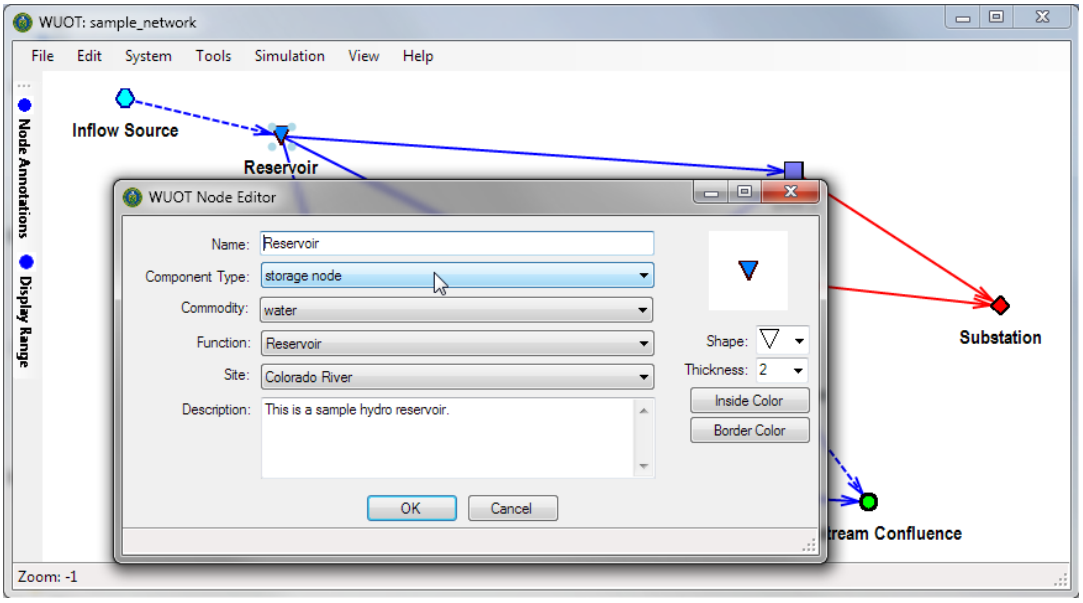

FIGURE 3-2 Node Creation Screen

created, converted to a different commodity, or combined with another commodity to form a new one. Attribute values may change over simulated time as commodities flow through the network. All commodities and attributes are user-defined, although certain commodities and attributes are included by default in a new project. For example, the water commodity and associated flow rate attribute are included by default in every project because that is the only commodity-attribute combination with which it makes sense for the hydrologic forecasting tool to operate. The day-ahead and real-time tool in particular can make use of new user-defined commodities and attributes, which will be described in detail later.

Node and link types are generic, enabling a component to represent a wide range of physical objects and processes. The key advantage of this approach is its flexibility: because the approach uses generic—rather than specific—components, the user can model any physical object without depending on the developers of the model to have anticipated the need for that object. For example, the interface does not provide one node type that has been specifically designed to represent a water reservoir and another node type that has been specifically designed to represent a battery. Instead, the interface provides a generic "storage node," and the user supplies the details about what is being stored (e.g., water, electricity, or something else) and storage characteristics (e.g., minimum and maximum storage capacity, maximum rate of release, rate of loss over time, etc.). In addition to the storage node, the interface provides the following components:

1. A functional boundary node represents the entrance or exit of a commodity into or out of the modeled network, such as water flows into a hydro-reservoir;

2. A conversion node represents a location for the transformation of one commodity into another, such as a hydropower plant that converts water into electricity;\

3. A junction node represents the joining together of several homogeneous commodities, such as the confluence of two rivers; and

4. A link carries the modeled commodities between nodes.

Commodity attribute properties are computed at various points in the system. These computations can be used simply for informational purposes, or the user can place limitations on attributes in order to constrain system operation. Figure 3-3 shows a sample network that utilizes each of the different node types.

Water enters the system via the functional boundary node labeled "Inflow Source" in the upper left corner of the screen. In this sample network, inflow values would likely be provided as output from the hydrologic forecasting tool. However, as with each of the tools, if the user elects not to use a particular tool then the values that would have otherwise been produced may be provided exogenously by the user. Water flows along a link from the Inflow Source to the storage node labeled "Reservoir," and a value for the *flow rate* attribute of the *water* commodity is calculated and stored to represent this movement. At the Reservoir node the water may be held over time until it is released, and values for the attributes of *volume* and *elevation* are calculated for the water stored here. Either the seasonal or the day-ahead and real-time tool may decide how much water to release from the Reservoir node and when to release it. Water released from the Reservoir node may take any of three pathways. One pathway leads down a link representing a bypass tube into a junction node representing the plant's tailrace. The other two pathways each lead to a conversion node representing a turbine and generating unit. At these nodes the water commodity is used to produce the *electricity* commodity, and the associated *energy* attribute flows to the functional boundary node representing a substation. The water that is used in the conversion process moves down to the Tailrace node. From the Tailrace node the water continues flowing downstream, until it is joined with a sideflow and arrives at a functional boundary node representing a confluence. The environmental tool may have objectives present at this confluence, for which it will calculate a score using the incoming flow rates determined by one or more of the other tools.

Each link must carry one and only one commodity (although each commodity can potentially have many attributes), and must also flow in only one direction. When a pair of nodes requires a particular commodity to flow in two directions between them (for example, a pumped storage plant requires water to flow in both directions between the upper reservoir and the turbine/pump), then the user may add a second link flowing in the opposite direction. There is no limit to the number of links that may enter or exit a given node.

It should again be emphasized that while all tools use the same network representation, each tool uses the network differently, or uses only the subset of the network that is pertinent to its objective. For example, the hydrologic forecasting tool is concerned only with the water

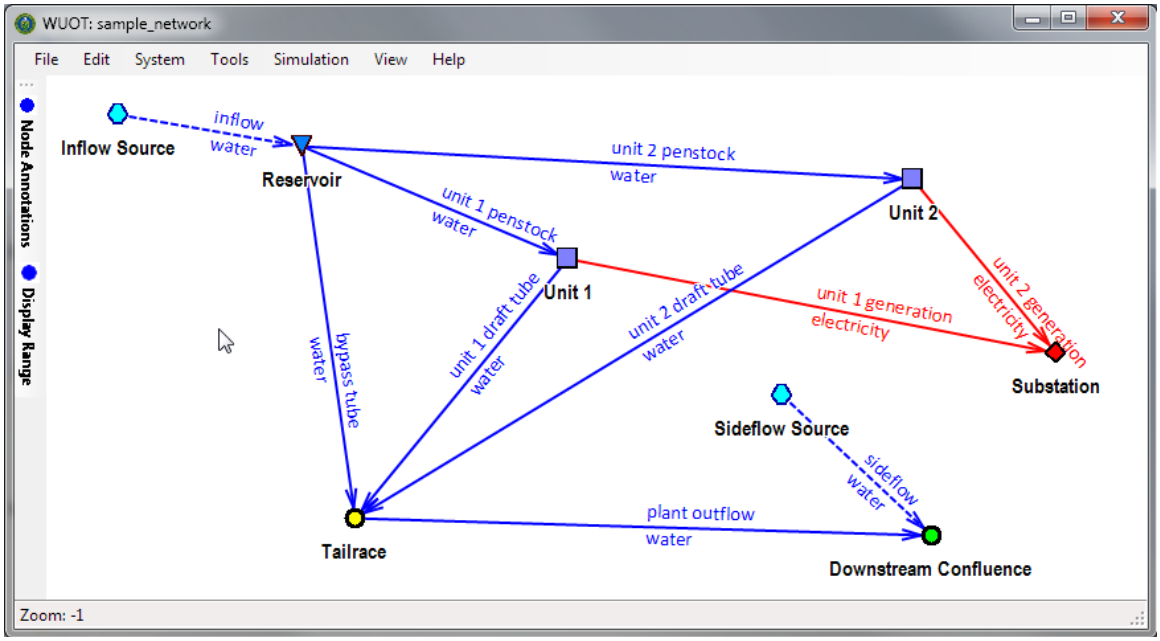

FIGURE 3-3 Sample Network Using All Node Types

inflows at functional boundary nodes, and ignores other nodes. Likewise, the environmental tool ignores any nodes that do not have environmental objectives defined. The seasonal tool determines reservoir releases using a plant-level rather than a unit-level representation, so it will ignore the individual generator nodes. The day-ahead and real-time tool operates at the lowest level of granularity and uses every node and link in the network. Even though the individual tools have these inherent differences and operate at different spatial and temporal resolutions, the whole toolset is designed in such a way that the tools understand each other's outputs and operate in an integrated fashion.

3.1.6 Database Design

Just as commodities and attributes are customizable, so are the units of measure used to represent them. Certain common units of measure are included by default in a new project, but the user can add and define new units of measure and provide the calculation that should be used to convert one unit another. For consistency, all values associated with any given attribute are saved in the database using the same unit, which is designated by the user. When entering input data or viewing result data, the user can elect to use a different unit than the one used for database storage, and a conversion between the interface unit and the database unit will be performed automatically as necessary.

Just as all attribute values are stored using a common unit, all time values in the database are stored in terms of Universal Coordinated Time (UTC). This ensures that there will not be any confusion when dealing with network components spread over different time zones, or when transitioning between standard and daylight saving time. The user may specify a local time zone

to use when entering inputs or viewing results, and the time values will be converted between UTC and local time as necessary.

All numeric time series data (for example, the water elevation level at a reservoir, or the energy output from a hydro plant) stored in the database is associated with a particular "data source." The definition of a data source includes an identifier for the WUOT tool or other external source (such as a SCADA server) that produced the values, a timestamp for when the values were produced, and an indicator of the time interval at which the data is stored (hourly, daily, 15 minutes, etc.). Every time one of the tools is executed, a new data source is created to hold the results of that particular execution. In this way, results from every past execution of a tool can be reviewed by an operator of the model or used as input into another execution of one of the tools. The standardization of the data source format and the identification and selection of data sources through the GUI is the mechanism by which the tools are able to easily share data. For example, the user may wish to use the environmental tool to calculate a score for a series of future river gauge flows that was calculated 2 days ago by the day-ahead and real-time tool. In this case, the user would open a menu to invoke the environmental tool; in that menu, the user would choose to use as input the data source that was produced 2 days ago by the day-ahead and real-time tool.

Storing tool results as separate series, rather than as a single series that is continuously overwritten with the latest information, facilitates auditing of past model outcomes and allows the toolset to be used in a "what-if" analysis capacity. For purposes of data management, once a particular data source is no longer needed by the user it can be deleted at will.

3.1.7 Providing Tool Inputs

Once the network has been created, the user must enter information that characterizes the properties, limitations, and objectives of the system. Depending on the nature of the particular data to be entered and the tool for which the data is intended, the GUI provides different mechanisms for data entry. For example, the user may right-click on a particular node and select from the context menu that appears the option to edit the node's environmental objectives, at which point a form will appear that is tailored to characterize such objectives. As another example, the user can click on a top-level system menu to access dynamic Excel workbooks that allow viewing and editing of overall system specifications, such as the commodity and attribute definitions. More detailed information about the input data required by each tool is provided in each tool's section of this report.

When the WUOT is applied at a site, the network design and model application will certainly involve data that must be updated on a frequent basis in order for the toolset to produce meaningful and useful results. Such data might include the current observed reservoir elevation levels, energy market price forecasts, the availability of generating units, and so on. Although such frequently updated data could be entered into the interface manually using the methods described above, doing so would be time-consuming and error-prone. In order for the toolset to be useful for actual hydropower operators, such data needs to be entered into the toolset quickly

and automatically. However, this need is complicated by the fact that every site could potentially have a different data storage format, a different mechanism for accessing the data, and varying security protocols for direct electronic systems access by other tools. To address all of these problems, the toolset provides a standardized method for accepting data in a flat text file format. The WUOT modeler first determines what site data will be needed for the particular application and how it should be applied in the toolset, and then enters "site transfer rules" representing this information into the toolset. A paraphrased sample site transfer rule is, "Take the values from flat file series #18 and apply those values on an hourly basis as the maximum possible water release through turbine number 2." The site's information technology staff writes a procedure that, when invoked, accesses the site's data systems and produces a flat file containing the necessary data. The flat file is posted to a location accessible to WUOT, which can then import and use the data in subsequent model runs.

3.1.8 Invoking the Tools

A top-level system menu in the GUI provides access to invoke a run of each of the tools individually. When clicking on each option, a new form will appear from which the user can set key run parameters, start the execution of the tool, and monitor the tool's progress. Figure 3-4 shows an example of such a form being used to monitor the progress of the day-ahead and real-time tool.

Details regarding the structure and parameters of each tool's run are provided in the tool's section of this report. As discussed earlier, the hydrologic forecasting tool performs its calculations on a server located at PNNL, so the "run" invoked for that tool is actually a procedure that connects to the PNNL server's database via the Internet and downloads the results of the forecasting run as directed by the user. While the environmental tool can be directly invoked by the user for the purpose of evaluating a set of hydrologic conditions, typically that tool is invoked automatically without user intervention by the seasonal tool as it performs its optimization.

3.1.9 Viewing Tool Results

The GUI provides various mechanisms to view and analyze model results depending on the nature of the data to be analyzed and the context of the tool that produced it. Numeric and text-based data, both static and time series, can be viewed in the network alongside the nodes and links. Each node has eight slots around its perimeter in which the user may elect to display different types of information. Figure 3-5 shows the screen used to specify what data should be shown in each slot. A similar form containing three slots is available for links.

The categories of available information are listed on the right side of the screen, and the user can drag each item into one of the empty slots. The chosen information will then be shown in that position for every node or link in the network. The font, size, and style of the text shown in the network can be customized by the user, and the color of the text can be set to automatically

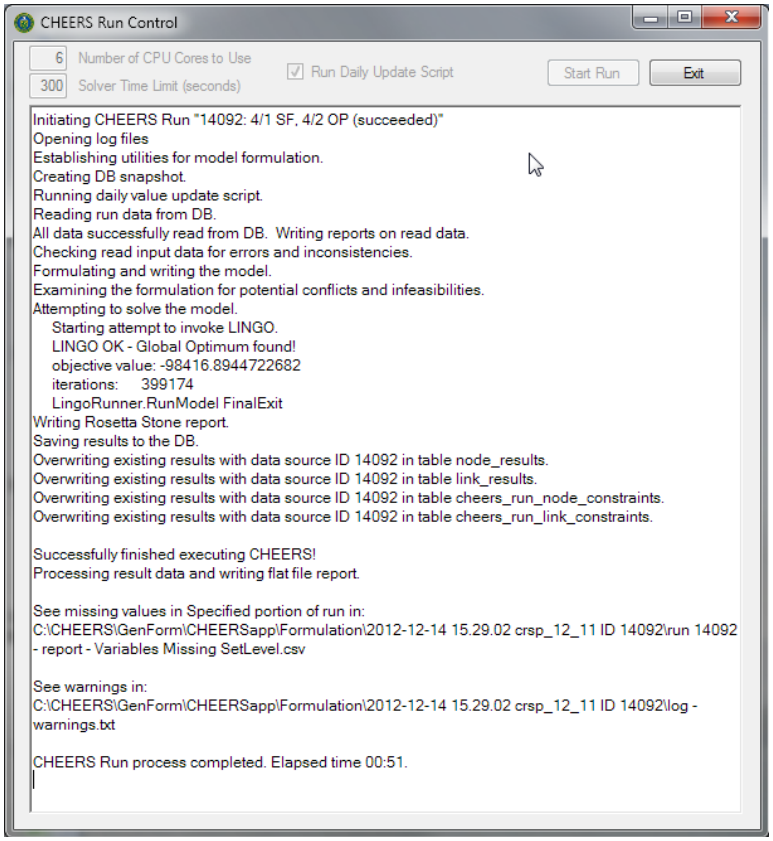

FIGURE 3-4 Form to Invoke and Monitor the Day-Ahead and Real-Time Tool

match the color of the node or link to which it belongs, making data identification clearer in a crowded network. Some of this information is static, such as the node name or the node's associated commodity, and other information is time-based. For time-based information the user can show up to three series, and can specify a different attribute for each commodity in each series. These options limit the total amount of information clutter on the screen by allowing the user to target exactly the information of interest. For example, say the user would like to see the flow rate attribute of the water commodity and the energy attribute of the electricity commodity. Figure 3-6 shows the menu in which the user makes this specification by selecting the desired attribute from a drop-down list associated with each commodity.

If the user then drags this series, "Display Type A," to the slot below the node, then every node in the network that deals with the water commodity will show the incoming and outgoing flow rate, and every node that deals with the electricity commodity will show the incoming and outgoing energy rate. A node that deals with both water and electricity (for example, a conversion node that takes in water and sends out water and electricity) will show both types of values, with an abbreviation indicating which is which. If a node deals with neither water nor electricity, then nothing will be shown in that node's bottom slot.

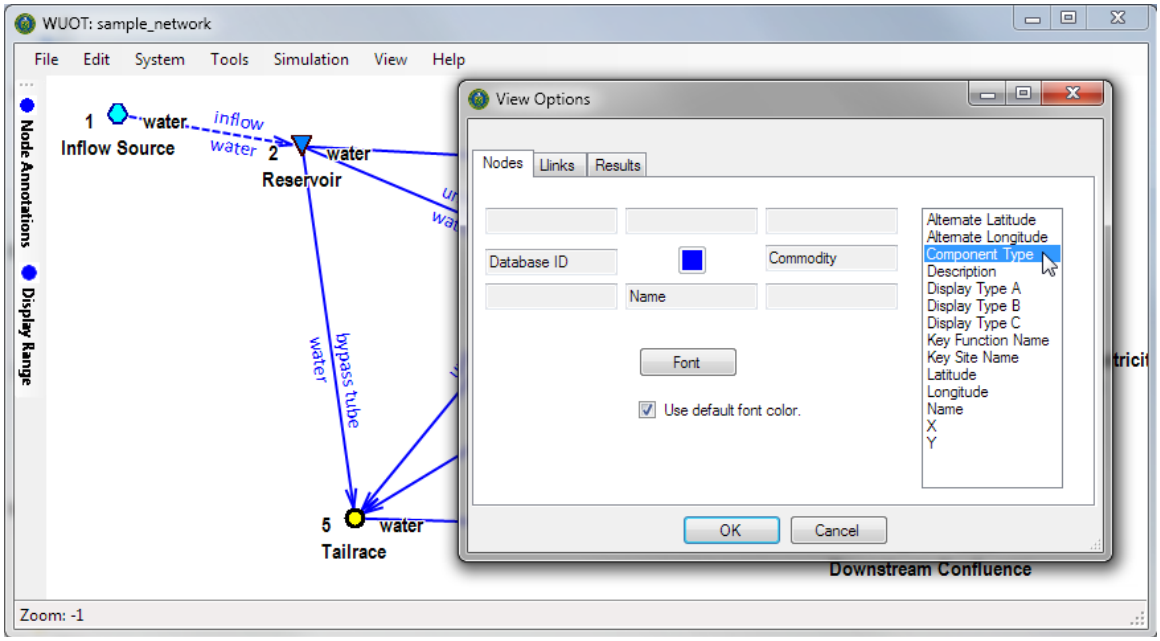

FIGURE 3-5 Node Annotations Screen

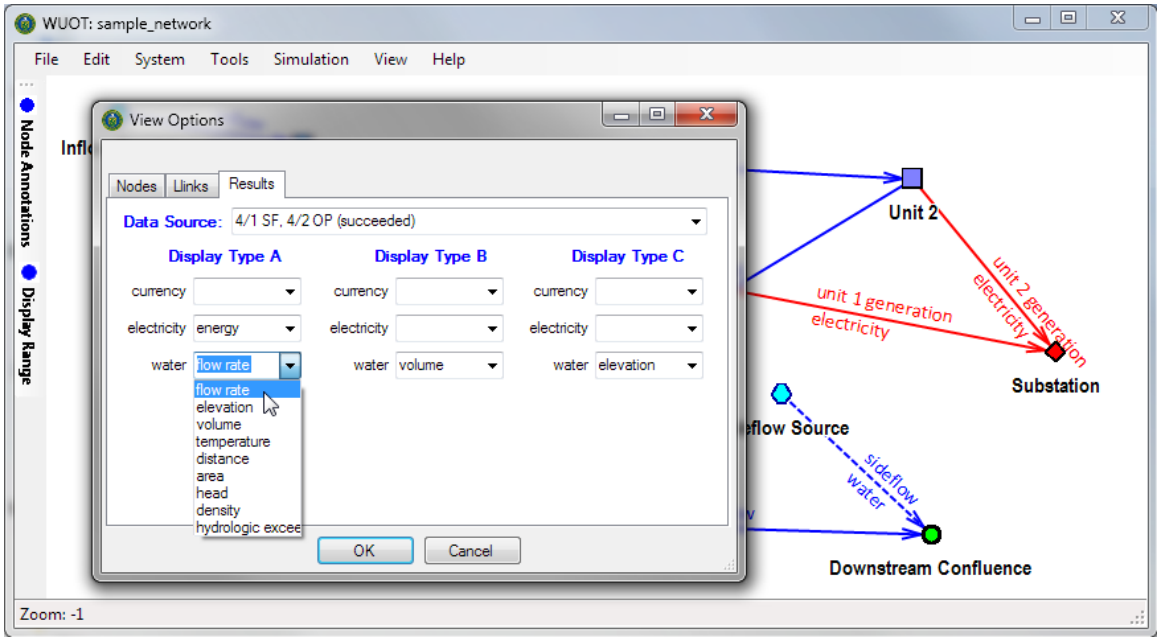

FIGURE 3-6 Result Annotations Screen

In the network view, the user can select from a list the particular data source associated with the data to be displayed, and can access a timeline slider control to step through all the time intervals available in the selected data source. Instead of viewing only one time interval, the user can also view a range of time intervals and apply different mathematical functions to the range of values selected. For example, if the selected data source uses an hourly interval, the user might elect to show the average reservoir elevation over the course of a 24-hour period, or the total energy produced by a generating unit over an 8-hour period.

Figure 3-7 is an example of day-ahead and real-time result data being shown with the timeline slider. For this snapshot in time the user can see the flow of water or electricity along each link, the amount of water and/or electricity entering and/or exiting each node, and the volume of water in the reservoir.

Time series graphs can also be produced using the GUI. Figure 3-8 is an example of a time series graph showing generation from each of the units along with the reservoir elevation. Using the controls on the form, the user can add, edit, hide, or remove series; change the visual appearance of a series; zoom in or out of a particular area on the graph; or copy a picture of the graph. If the user configures a graph and would like to recall the graph configuration later when examining a different data set, the configuration can be quickly saved and recalled. By clicking the "View Table" button, a table can be accessed that contains all of the numeric data shown in the graph. As shown in Figure 3-9, data from this table may be selected and copied to quickly move it to another application.

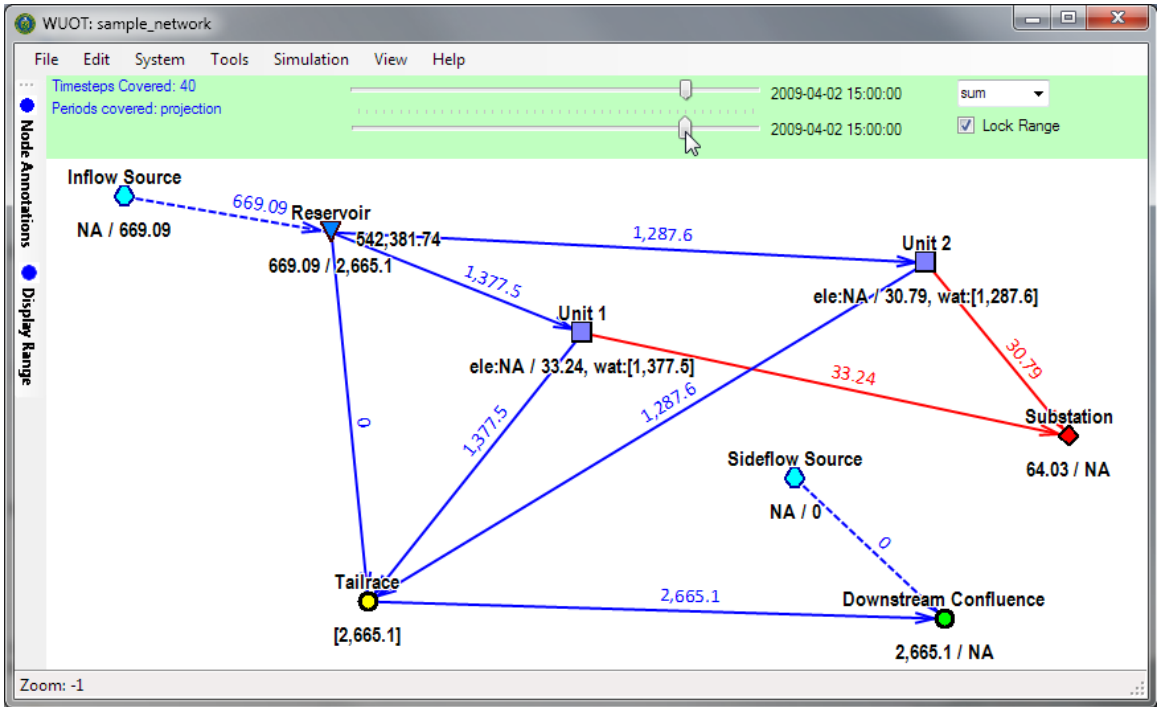

FIGURE 3-7 Result Data Viewed on Network Window with Timeline Slider

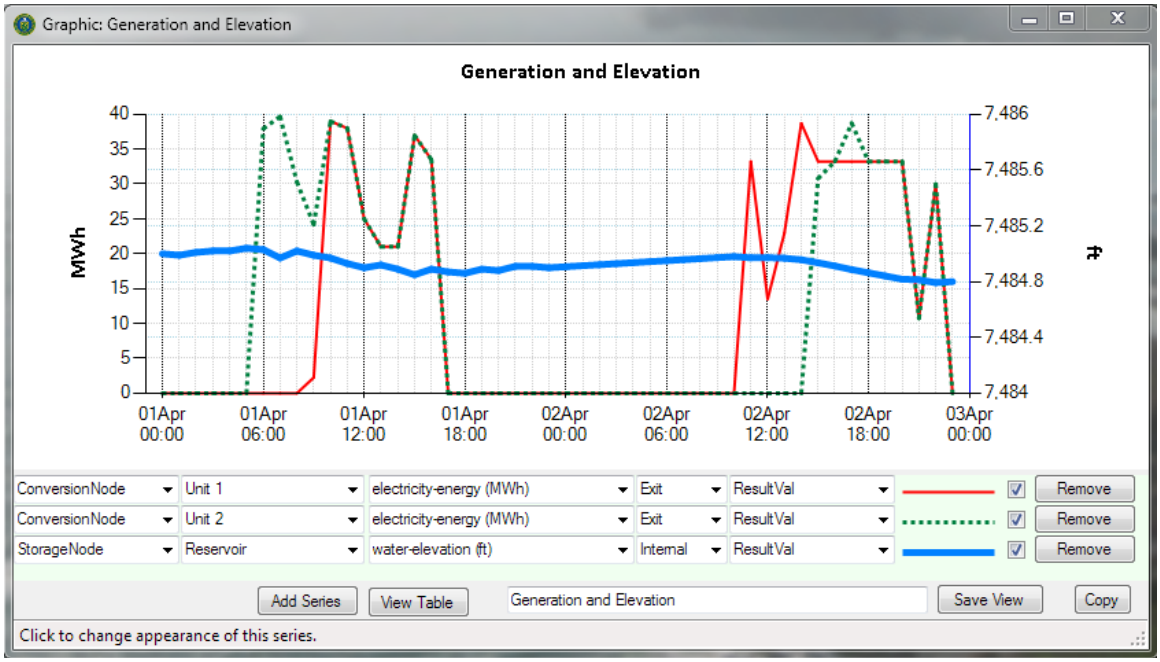

FIGURE 3-8 Time Series Graph Screen

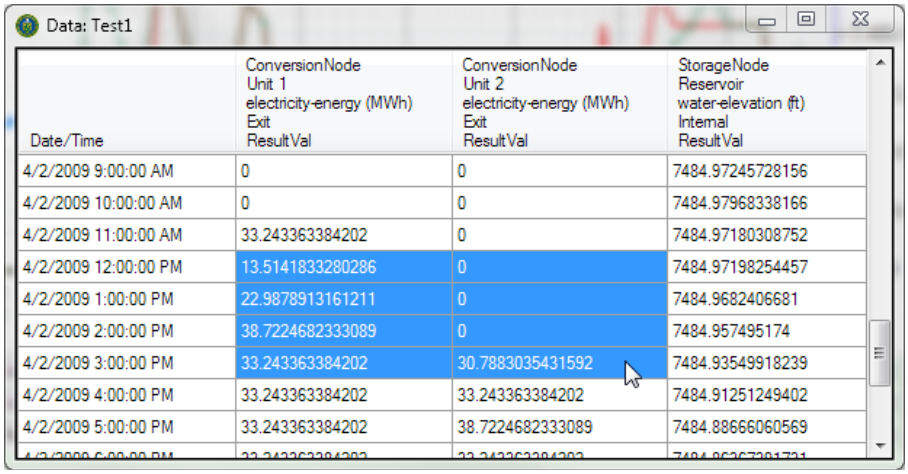

Data: Test1

Date/Time	ConversionNode Unit 1 electricity-energy (MWh) Exit ResultVal	ConversionNode Unit 2 electricity-energy (MWh) Exit ResultVal	StorageNode Reservoir water-elevation (ft) Internal ResultVal
4/2/2009 9:00:00 AM	0	0	7484.97245728156
4/2/2009 10:00:00 AM	0	0	7484.97968338166
4/2/2009 11:00:00 AM	33.243363384202	0	7484.97180308752
4/2/2009 12:00:00 PM	13.5141833280286	0	7484.97198254457
4/2/2009 1:00:00 PM	22.9878913161211	0	7484.9682406681
4/2/2009 2:00:00 PM	38.7224682333089	0	7484.957495174
4/2/2009 3:00:00 PM	33.243363384202	30.7883035431592	7484.93549918239
4/2/2009 4:00:00 PM	33.243363384202	33.243363384202	7484.91251249402
4/2/2009 5:00:00 PM	33.243363384202	38.7224682333089	7484.88666060569

FIGURE 3-9 Time Series Table Screen

In addition to the methods described above, model data can exported through the interface to plain text files or to dynamic Excel workbooks from which it can be investigated.

When the WUOT is applied at a site, the site's staff may wish to have WUOT results imported into their business system. This can be accomplished using the same flat file procedure described in Section 3.7, although of course in this case it is WUOT that produces the flat file, and the site that imports it.

3.2 ENHANCED HYDROLOGIC FORECASTING SYSTEM (EHFS)

3.2.1 Concept

Hydroelectric power production is one of many competing demands for available water; these demands include other priority uses such as irrigation, recreation, thermoelectric cooling, and environmental performance. Each demand has specific temporal and spatial requirements that must be balanced against other priority hydrosystem functions such as flood control. An accurate forecast of the magnitude and timing of water supply is essential for optimal management. Such forecasts set the parameters that define water resource supply and demand at various timescales, as well as the associated uncertainty. Forecasts assist managers in integrated planning and operations to balance competing water uses against current and future supply while protecting against the possibility of water or energy shortages and excesses with real-time actions. In the absence of appropriate forecast products, the hydrosystem may be operated using overly conservative constraints to meet multiple water use objectives and mitigate the impacts of hydrologic extremes (flood, drought). However, customized forecast products with increased forecast accuracy and lead times may allow improved system optimization by relaxing overly conservative constraints without increasing risk. For example, water that may have been held in storage for summer irrigation or to meet environmental flow requirements could be released to maximize power generation if reliable forecasts indicate adequate future water supply.

Some agencies use the National Weather Service (NWS) flow forecast directly from the River Forecast Center (RFS), which generates flow forecasts for predetermined locations and include regulated flow and probabilistic flow forecasts, which are based on the chance of exceeding different flood thresholds. RFS flow forecasts are constrained by the current setup of the hydrologic model (lumped model) and the nature of the flow forecast because the flow is regulated with potentially significant uncertainties in the regulation, since it is often based on communication with the multiple control agencies. Therefore, many utilities use their own flow forecasts (e.g., Bonneville Power Administration), contract third parties to provide appropriate flow forecasts, and/or simply use precipitation forecasts or past climatology. The objective of this EHFS's task is to develop and demonstrate a national medium-range (1–13 days ahead of time) to seasonal forecasting system that incorporates enhanced forecasting approaches and algorithms, spatiotemporal datasets, and automated data acquisition and processing at user-specified locations of interest. These forecasts provide reservoir inflows for our demonstration partners with an emphasis on the lead times required by their operations and at the exact locations required by the toolset in order to ensure the multi-objective optimizations within the WUOT can be properly executed (i.e., in particular, the locations required by the environmental performance tool). The EHFS provides a consistent, national approach for multiscale ensemble streamflow forecasting that incorporates (1) enhanced forecasting approaches; (2) a distributed, physics-based hydrologic model; (3) current spatiotemporal datasets; and (4) automated data acquisition and processing. This EHFS is integrated into the WUOT to inform reservoir and power operations and to guide environmental performance decision making.

3.2.2 Design

This project integrates state-of-the-science ensemble streamflow forecasting systems developed by PNNL, the University of Washington (UW) (Wood and Lettenmaier 2006), and Princeton University to provide a next-generation tool, collectively referred to as the EHFS. The PNNL portion of the EHFS was supported by DOE and the National Aeronautics and Space Administration (NASA); the UW/Princeton piece of the EHFS was largely developed with support from the National Oceanic and Atmospheric Administration (NOAA) and NASA. Uncertainty in forecasts is addressed explicitly through the generation of ensembles of streamflow forecasts using hydrologic models that are driven by ensembles of possible future meteorological realizations. The UW/Princeton forecasting system provides ensemble seasonal volumetric forecasts throughout the nation (Figure 3-10).

The PNNL EHFS is built on the UW/Princeton seasonal forecast system, which has been enhanced under this project through the use of remotely sensed data, automated spatial and temporal data assimilation, and the added capability for both medium-range and seasonal forecasts for watersheds over a range of scales throughout the United States. More specifically, the EHFS has been improved over its predecessor by the following:

- Daily medium-range ensemble flow forecasts;
- River flow assimilation and improved snow data assimilation;
- Customized flow forecast locations, including natural flow and contributing flow;
- Decreased computational time by allowing the use of multiple processors;
- A flexible and generic environment for extended or new applications, including change in spatial resolutions, multiple hydrology models, multiple forecast products, and new basins;
- Compatibility for writing output to a Web-accessible database;
- Improved error detection and reporting, enabling a faster and more efficient resolution of problems; and
- A template for forecast verification of the ensemble mean forecast and uncertainties.

The following are some of the feature/capability highlights of the EHFS:

- A complete and readily updateable multi-temporal scale national database (medium range to seasonal);

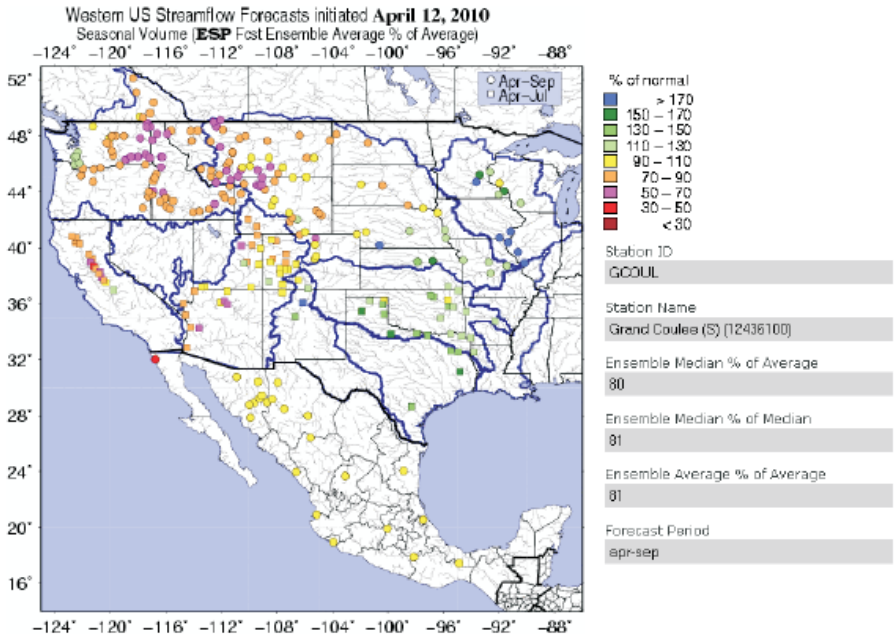

FIGURE 3-10 Locations of Seasonal Ensemble Forecasts from the University of Washington/Princeton University Forecast Model (Source: http://www.hydro.washington.edu/forecast/westwide)

- Automated data acquisition and processing of dynamic and real-time data;
- Advanced, automated data assimilation;
- Integrated GIS-based data and physics-based models;
- Utilization of an ensemble-based methodology for both medium-range and seasonal forecasts to support uncertainty analyses of generated forecasts; and
- Integration into the Water Use Optimization Toolbox.

3.2.2.1 The Variable Infiltration Capacity (VIC) Model

The forecast consists of two major components: (1) a "nowcast" that is used to estimate the current distribution of water stored within the basin; and (2) a "forecast" that estimates the change in water storage and resulting streamflow under the forecasted meteorology (see Figure 3-11). A time series of observed meteorology is used to drive (spin-up) the hydrologic model during the nowcast. The spin-up period is long enough that the influence of assumed initial conditions (ICs) at the start of the simulation is removed and the model state reflects a best estimate of current conditions prior to the forecast. At this time, observed snow water equivalent (SWE) data (e.g., the USDA Natural Resources Conservation Service Snowpack Telemetry

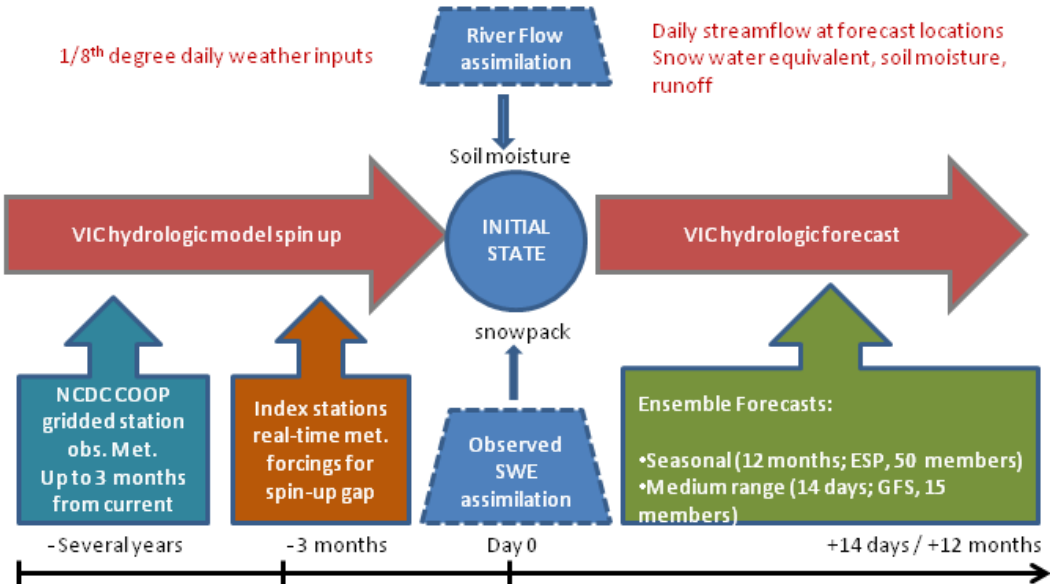

FIGURE 3-11 Component Diagram of the Forecasting Sequence (adapted from Wood and Lettenmaier [2006]) (During the "nowcast," the hydrologic model is spun up using historical meteorological data, then uses recently observed meteorological data to generate the model representation of the basin water state prior to the hydrologic forecast. Model snowpack and soil moisture states are then updated through automated data assimilation of observed snow water equivalent and streamflow. The model is then driven by medium-range and seasonal meteorological ensemble forecasts to generate a spatially explicit ensemble of streamflow forecasts.)

[SNOTEL]) may be used to refine model estimates of water storage in the snowpack. Observed streamflow may also be used to update the soil water content (see Figure 3-11). The model is then driven by an ensemble of meteorological forecasts to generate an ensemble of streamflow forecasts (Figure 3-11) over different temporal lead times and resolutions.

The EHFS approach uses the Variable Infiltration Capacity (VIC) model (Liang et al. 1994) to track the basin water balance and predict how the basin will respond to the forecasted meteorology. The model is "driven" by daily meteorology, including precipitation, wind, and air temperature, which varies spatially throughout the basin. In its various forms, VIC has been applied to watersheds including the Columbia River, the Ohio River, the Arkansas-Red Rivers, and the Upper Mississippi River; the model has also been applied globally. VIC is a grid-based, semidistributed model that tracks the movement and storage of precipitation in time and space, accounting for rain and snow interception by vegetation, transpiration, evaporation, snow accumulation and melt, and water movement and storage in the soil. Each computational grid cell may contain a number of vegetation and soil types by presenting its fractional area within the cell (Figure 3-12). The current moisture state of the basin and the location and amount of water stored in the snowpack and basin soils dictate how the hydrology of the basin will respond to future weather conditions at medium-range and seasonal time scales. Values have been defined

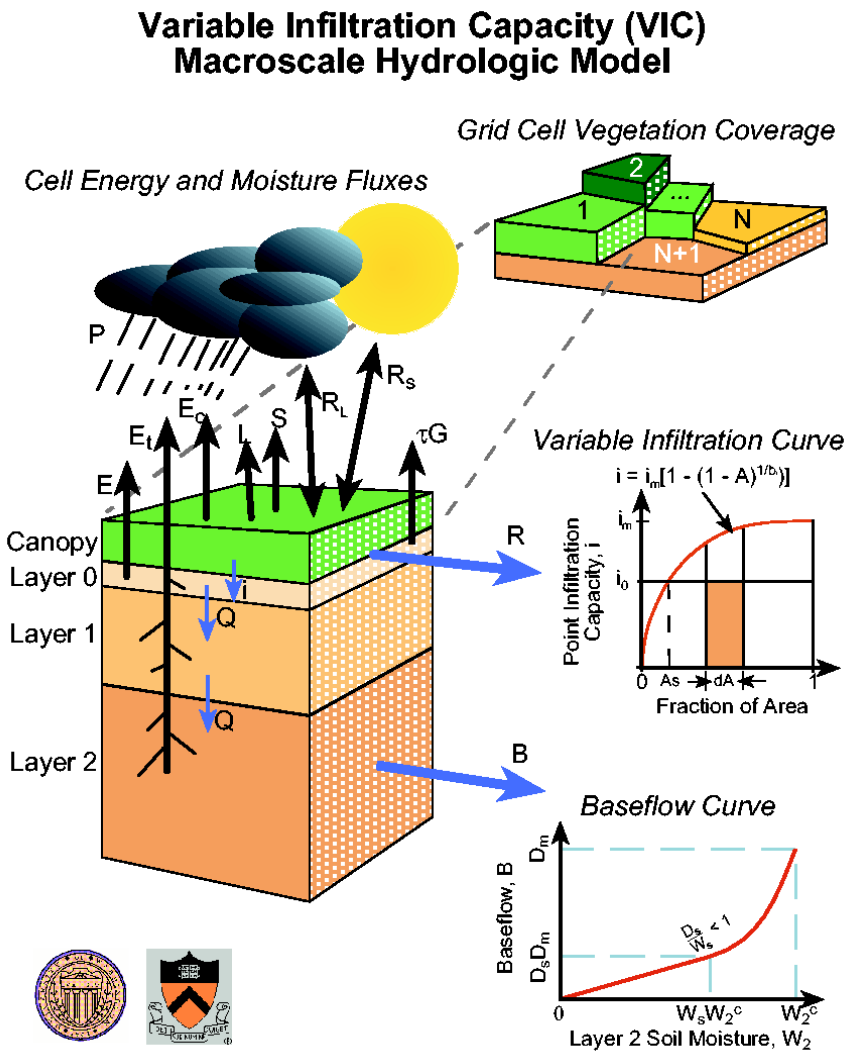

FIGURE 3-12 Representation of Water and Energy Balance in a VIC Model Grid Cell (each grid cell may contain multiple vegetation and soil types)

for all VIC input parameters for each one-eighth-degree grid cell across the continental United States (with grid cell size increasing from approximately 126 to 172 km^2 from the northern to southern latitudes of the United States).

The EHFS runs VIC daily to track the current basin water balance (Figure 3-12) and predict how the basin will respond to forecasted meteorology. In addition, the EHFS also utilizes a routing model (Lohmann et al. 1996, 1998) to transport runoff and baseflow from each grid cell to and through the channel network. A simple linear transfer function model (unit hydrograph) is used to route flow from within a cell to the cell outlet. The flow exiting the grid cell is then routed to the forecast location based on a second independent specified unit hydrograph. The runoff transport is time invariant and linear.

The VIC model is forced by a daily one-eighth-degree gridded observed meteorological dataset (Maurer et al. 2002; Wood and Lettenmaier 2006) during the spin-up period. In the medium-range forecast mode, VIC is forced with a daily 13-day-long, 15-member ensemble weather forecast derived from downscaled Global Forecast System (GFS) meteorological forecasts (Voisin et al. 2010). In the seasonal forecast mode, VIC is forced by a daily, 12-month-long, 49-member ensemble forecast (resampling from 1960–2010 climatology). These medium-range and seasonal forcings are part of the EHFS products. The following meteorological inputs and units provide the model forcings:

- Daily total precipitation (mm);
- Daily maximum air temperature (°C);
- Daily minimum temperature (°C); and
- Average daily wind speed (m/s).

Embedded in VIC are algorithms (Thornton and Running 1999; Kimball et al. 1997; Bras 1990) that allow the derivation of the full subdaily energy forcing required to properly run the hydrology model (as denoted by an asterisk in the list below). A range of outputs are then available for each model time step for all grid cells and can be mapped to provide meteorological and energy-balance input for other WUOT components. Output parameters of direct relevance for the WUOT include the following:

- Year;
- Month;
- Day;
- *Hour;
- Precipitation (mm);
- *Net shortwave radiation at the surface (W/m^2);
- *Net short and longwave radiation at the surface (W/m^2);
- *Relative humidity (fraction);
- *Specific humidity (kg/kg);
- *Incoming longwave radiation below the canopy (W/m^2);
- *Air temperature (°C);

- *Atmospheric density (kg/m^3);
- *Atmospheric pressure (kPa);
- *Near surface vapor pressure (kPa);
- Near surface wind speed (m/s);
- Runoff (mm); and
- Base flow (mm).

VIC Model Calibration, Streamflow Assimilation, and Post Processing. A flow forecast system consists of a series of forecast models and processing steps. First, the hydrological model needs to be calibrated with respect to observed natural flow and forced with an observed meteorological dataset. Next, the weather forecast needs to be adjusted for eventual systematic errors with respect to the observed meteorological dataset; this is handled by the EHFS pre-processor. Flow forecasts are also driven by soil moisture and snowpack conditions. Streamflow data assimilation is performed in order to improve the initial conditions on the day of the forecasts. Finally, a gap might exist between the near-real-time simulated flow and the observed flow, creating a *de facto* bias for the short-lead-time flow forecasts. A post-processor ensures the connection of the short-lead-time flow forecasts to the near-real-time observed flow. Each will be discussed in more detail in the sections on the demonstration basins; this section is just a brief overview.

Model calibration adjusts the soil parameterization and routing parameters of the VIC grid cells. In an iterative process, the suite of soil parameters are adjusted, the flux for each grid cell calculated, and flow routed; the routed flow at specific forecast points are compared to observed streamflow data. Nash-Sutcliff efficiencies (NSE), bias, root-mean-square errors (RMSE), and other parameters are calculated for the run, then an improved parameter set is generated and run until the set of parameters converge to stable values for optimal performance (Nash and Sutcliffe 1970). The soils input parameter file produced in the final run is then used in the forecast framework.

A pre-processor is applied to the medium-range weather forecast in order to spatially downscale to the grid size of the hydrological model, remove any systematic bias in the mean errors, and improve the probabilistic information of the ensemble weather forecast for input into the hydrological model.

Streamflow assimilation is used to adjust the initial condition moisture state of the soils file (“state file”) based on near-real-time observed streamflows. In the Gunnison River Basin, soil moisture values in the deepest model soil layer, the layer responsible for the generation of base flow, are replaced by direct substitution with values estimated from the observed streamflow. In the Feather River Basin, there is a faster hydrologic response, so the moisture

values in the middle soil layer were adjusted. Soil moisture values are updated for all areas upstream of a given streamflow gauge in such a way that the spatial patterns of wetter and drier cells are preserved.

Post-processing of the streamflow forecast results at the individual forecast points can be used to further improve model performance in some basins. In particular, this post-processing allows the simulated real-time flow to be reconnected with the observed flow on the day of the forecast, and therefore reduces the discrepancy between short-lead-time flow forecast and observed flow. A simple post-processing has been used at individual forecast points for the Gunnison. Near-real-time streamflow data were used to offset the flow nowcast by the difference between the observed and simulated flow on the day of the forecast.

Nowcast. Daily historical meteorology is used as input to retrospective daily runs of VIC to generate the climatology that establishes "normal," thus providing and distinguishing between the nowcast and forecast runs for anomalies and percent of normal. In addition, these data are used for model calibration and validation. The observational data are used to generate a spatially distributed forcing dataset in the method of Maurer et al. (2002). A database of daily meteorological data has been established from NOAA's Cooperative Observer (COOP) meteorological stations via the National Climatic Data Center's (NCDC's) Summary of the Day dataset and all available Environment Canada stations in the Canadian portion of the Columbia River Basin. In order to derive near-real-time hydrologic initial conditions (snowpack and soil moisture on the day of the forecast), the hydrology model is forced by near-real-time observed meteorology. The near-real-time (yesterday's) meteorological observations are automatically collected daily from the available meteorological stations through the NOAA High Plains Regional Climate Center using the Applied Climate Information System (ACIS). The real-time database extends from 2010 to the present time and takes over where the historical meteorology leaves off. This brings the system from the retroactive states to the current day, and ultimately produces the current hydrologic state in preparation for making the hydrologic forecast. Currently, the hydrologic nowcast portion of the EHFS is running daily at PNNL over the states of California and Colorado.

Medium-Range Streamflow Forecasts. Run-of-the-river dams rely largely upon upstream reservoir management at the seasonal scale; therefore, these dams are typically responding to short- and/or medium-range flow forecasts. Similarly, smaller reservoirs in regions with high interannual variability benefit from accurate short- and medium-range flow forecasts to provide flood control and release a conservative but appropriate amount of water to conserve storage for other operational objectives, such as summer irrigation, recreation, and hydropower.

The GFS retrospective forecast dataset is used as the source of ensemble weather forecasts in the current implementation of the EHFS. The publicly available GFS forecasts provide a long-term dataset appropriate for training the downscaling approach and for the evaluation of the flow forecasts over a long period. This more-than 30-year retrospective forecast dataset (1979–2011) includes 14-day forecasts at 2.5° spatial resolution with a daily time step,

which are derived from the 1998 version of the GFS (Hamill et al. 2006). An updated version of the dataset with a 1° spatial resolution using the current operational version of the GFS model has just been made available and is being integrated into the EHFS. Forecast variables from the GFS model include precipitation, daily average air temperature, and zonal and meridional wind components.

An analog approach (Voisin et al. 2010) is used to (1) calibrate the information in the forecast ensemble (bias, probabilistic information); (2) downscale the forecast variables to the scale of the hydrology model; and (3) derive minimum and maximum temperature from the daily average temperature. The extended 1950–2010 daily gridded meteorological dataset (Maurer et al. 2002) at one-eighth-degree spatial resolution is the source for analogs, therefore ensuring consistency with the nowcast period. The medium-range forecast portion of the EHFS is run and updated daily. Each day, an ensemble forecast is run and data such as that shown in Figure 3-13 is produced. As seen in the case study sections, we develop retrospective flow forecast datasets in order to verify our medium-range flow forecasts with respect to persistence. For example, the day-ahead forecast streamflow was extracted from medium-range forecasts for each day of 1998. Figure 3-14 shows the day-ahead ensemble streamflow forecasts, the average forecast streamflow, and observed flows for water year 1998 over the Feather River Basin at Oroville, California.

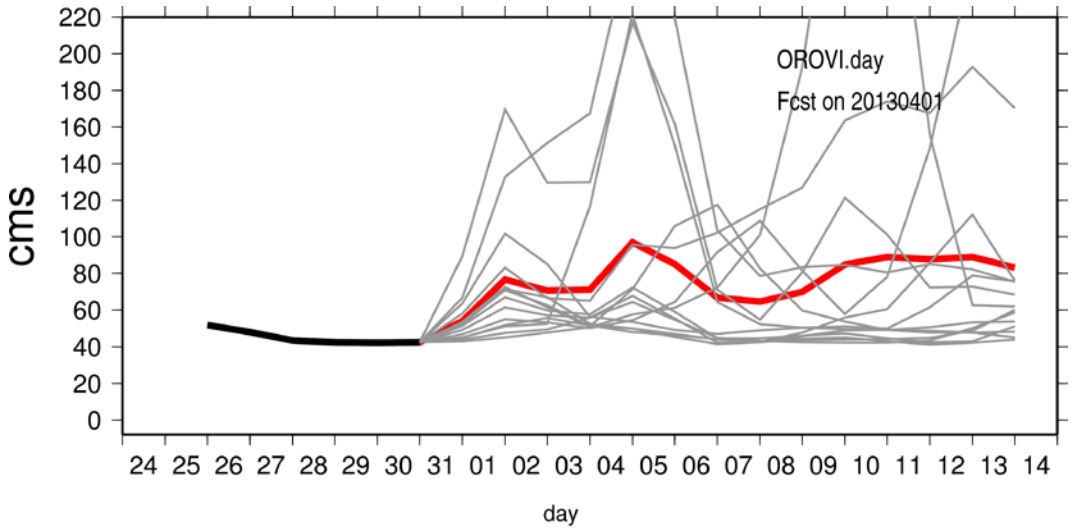

FIGURE 3-13 Daily Medium-Range Forecast for the Feather River at Oroville on 4/1/2013 (The red line is the forecast ensemble average; the gray lines show the individual ensemble members.)

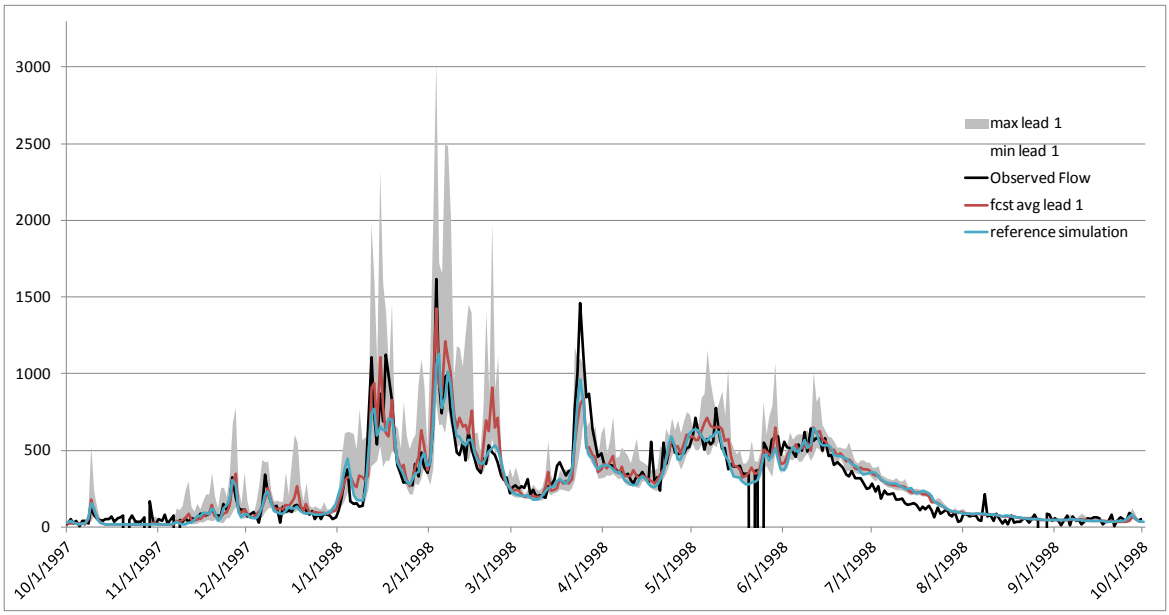

FIGURE 3-14 Day-Ahead Ensemble Flow Forecast for Water Year 1998 over the Feather River Basin in Oroville, California

Seasonal Streamflow Forecasts. Reservoirs with significant storage capacities rely on seasonal volumetric flow forecasts for their management. Shukla and Lettenmaier (2011) have shown that improvement in seasonal climate forecast alone will lead to better seasonal hydrologic forecast skill throughout the year in most parts of the northeastern and southeastern United States; however, in the western United States, the forecast skill is improved mainly during the fall and winter months. For the northwestern United States, where a significant portion of the United States' hydropower is generated, initial conditions tend to drive seasonal flow forecasts in the spring and summer months. The EHFS employs the Extended Streamflow Prediction (ESP) approach (Day 1985) used by the River Forecast Centers since the mid-1970s. The ESP relies on initial conditions and a resample of seasonal weather forecasts (traces) from previous years (1960–2010; 49 traces); this approach brings consistency between the nowcast and forecast systems. The seasonal forecast portion of the EHFS is typically run on a weekly basis. Where applicable, the ensemble forecast is filtered for traces targeting suitable climate indices for potential fine tuning of the ESP forecasts. An ensemble average forecast is also evaluated with respect to the climatology of different climate index phases (Figure 3-15).

3.2.2.2 Forecast Products and WUOT Integration

For the WUOT system to function efficiently, a relational database is key to the storage and retrieval of parameters, model run scenarios, required data, and, fundamentally, operating as the vehicle for intermodel communication and data flow. Argonne is leading the implementation of a MySQL relational database, referred to as the common database. A database schema, or data

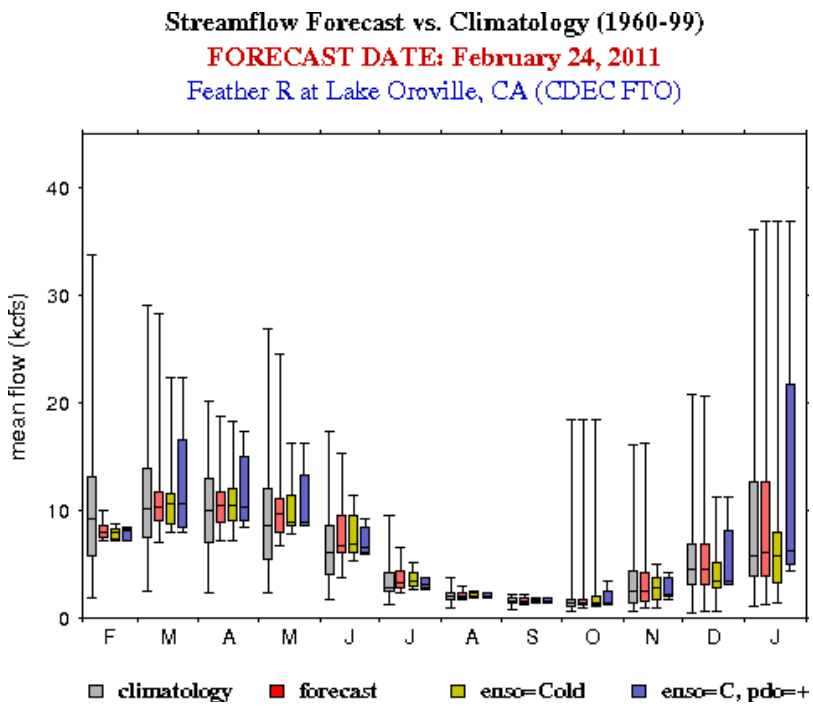

FIGURE 3-15 Seasonal Flow Forecast (2011) for Inflow into the Oroville Reservoir, California

structure, has been developed by members of the multi-laboratory team. The database schema provides a mechanism for designing data structure and relationships; it can be thought of as a blueprint for data that is structured inside of a database. Due to the volume of data produced for the medium- and seasonal-range forecasts, the retrospective forecasts, and required meteorological forcing data, it was deemed most efficient to house a Web-accessible database at PNNL that could be remotely queried by the individual components of the WUOT. This largely reduces the latency of data being moved from one system to another and makes information available near the completion of the model runs. From the perspective of the other tool components, once the connection to the PNNL Web database is established, there is no difference in the functionality to retrieve EHFS data; it is as if the database resides locally.

Natively, the EHFS writes its outputs to ASCII files in a complex date-based directory structure for medium- and seasonal-range forecasts, in addition to the meteorological forcing data and other supplemental data. In order to integrate with and pass data to the other WUOT modeling components, model output data needed to be structured into a relational database. The common database, the database used by the remainder of the WUOT models, uses MySQL and thus was also chosen for use on the EHFS.

Using the existing structure of the model output data, a working series of database schemas were developed to meet the needs of the model itself and the data integration process, keeping in mind the need to integrate additional demonstration basins and potentially other sources of hydrologic forecast data such as those from NOAA-NWS or industry. Currently, the

database only houses VIC-based forecasts produced by the EHFS. The database schema that was adopted takes into consideration current and future needs and is adaptable and scaleable. Because the EHFS operates on a one-eighth-degree modeling grid cell, a latitude and longitude pair are used as a top-tier element to access the remainder of the data in the EHFS database schema. To help with the integration process for the established forecast point locations (i.e., a functional boundary node in the CHEERS system), a lookup table is manually generated that relates the actual latitude/longitude location on the ground to the one-eighth-degree model grid cell latitude/longitude; this information is paired with a named code of the forecast point location (see Table 3-1). The lookup table is stored in the WUOT common database and allows the retrieval of data by simply using a named code, rather than inserting latitude/longitude pairs.

TABLE 3-1 Forecast Point Lookup Is Manually Established to Provide a Relationship between the On-the-Ground Latitude/Longitude and the EHFS Model Cell Latitude/Longitude

Forecast Point	Latitude	Longitude	Feature Description	EHFS Latitude	EHFS Longitude
BLMSA	38.453370	–107.334052	Blue Mesa Dam	38.4375	–107.3125
MRWPT	38.451700	–107.537900	Morrow Point Dam	38.4375	–107.4375
MRWPS			(side flow between BLMSA and MRWPT)		
CRYST	38.510402	–107.623774	Crystal Dam	38.4375	–107.5625
CRYSS			(side flow between MRWPT and CRYST)		
NFORK	38.782700	–107.837000	Just above the confluence of North Fork Gunnison and Gunnison	38.8125	–107.6875
UNCOM	38.756800	–108.090200	Uncompahgre River	38.6875	–108.1875
DELTA	38.753	–108.078	Delta (below confluence with the Uncompahgre River)	38.8125	–108.313
GRNJC	38.983	–108.45	Whitewater (happens to be in preexisting Grand Junction GRNJC grid cell)	38.9375	–108.438
OROVI	39.536266	–121.482576	Oroville Dam	39.5625	–121.4375

The database schema is structured into four major groupings beyond latitude and longitude:

1. *Seasonal Climate Forecast Forcings* The seasonal climate forecast products consists of a one-time resampling of the climatology and, therefore, is a static dataset in the database. The seasonal forcing group includes meteorological forcing data at a one-eighth-degree spatial grid resolution for a horizon of 366 days. The dataset includes hourly meteorological data from 1960 to 2010 and includes data parameters such as precipitation, air temperature, vapor pressure, atmospheric density, wind speed, specific humidity, incoming shortwave radiation, and incoming longwave radiation. The hourly dataset allows the representation of 49 plausible climates for the next year, with the understanding that the actual seasonal flow forecast is driven largely by the initial conditions. This hourly meteorological forecast dataset has no forecast skill per say; however, it gives the sub-daily variability of the weather required by the Hydroscope seasonal optimization tool.

2. *Seasonal Streamflow Forecast.* The seasonal forecast group includes an ensemble of daily data (49 traces + average trace) that is issued once every 7 days out to 1 year and includes not only the daily flow at specific forecast points, but also hydrologic fluxes including evaporation, precipitation, runoff, baseflow, soil moisture (top, middle, bottom layers), snow water equivalent, and hydrologic exceedance (based on 1960–2010 water year climatology).

3. *Medium-Range Meteorological Forecasts.* The medium-range ensemble meteorological forecasts are derived from the pre-processed GFS medium-range weather forecasts. It includes one-eighth-degree gridded, 13-day, 15-member, hourly meteorological data, issued daily for the period from 1990 to 2010. It includes data parameters such as precipitation, air temperature, vapor pressure, atmospheric density, wind speed, specific humidity, incoming shortwave radiation, and incoming longwave radiation. In this configuration, the weather forecasts are processed and have prediction skill at a daily time scale. The daily weather forecast data is disaggregated to an hourly time-step and has no specific subdaily prediction skill. It should be noted that with the hourly data, the precipitation is set as uniform across the day and other parameters are driven by a diurnal cycle.

4. *Medium-Range Hydrologic Forecasts.* The medium-range forecast group is structured similarly to the seasonal forecast group; however, the medium-range forecast group only includes forecasts out to 13 days for 15 ensemble traces and an average trace for the period from 1990 to 2010.

The EHFS database currently contains data for the Gunnison River in Colorado, and the Feather River in California.

3.3 SEASONAL HYDROSYSTEMS ANALYSIS

3.3.1 Concept

The objecive of HydroSCOPE, the Seasonal Hydrosystems Analysis portion of the WUOT, is to address seasonal-scale (1 to 6 months) operations and planning to help operators, planners, and policy makers ensure that current-day operations are optimal with regards to projected future conditions. HydroSCOPE has also been designed for use as a strategic planning tool and for optimizing under uncertainty.

3.3.2 Design

HydroSCOPE is a systems-level hydropower simulation and optimization tool that couples basin-scale simulations of water quality and quantity with optimization software. The simulation engine simulates changes in temperature, energy production, revenue, available water, and environmental performance due to alterations in project operations and/or projected future conditions. HydroSCOPE is seamlessly coupled with the DAKOTA optimization software (available at dakota.sandia.gov); the tool returns suggested operational schedules that will maximize the probability of achieving a set of user-defined objectives specific to that basin.

HydroSCOPE is run in one of two modes: simulation mode and optimization mode. Simulation mode executes the model deterministally to allow users to simulate one parameter set at a time in order to explore specific system behavior or to screen scenarios that may require further investigation. In optimization mode, the user defines sets of objectives (i.e., metrics), constraints, and changeable input variables for performing multi-objective optimization, parameter estimation, uncertainty quantification, and/or sensitivity analysis. Objectives are defined to maximize or minimize particular metrics (e.g., maximize energy production or minimize environmental impacts) or to achieve a target performance (e.g., maintain a defined reservoir elevation). Constraints are defined as physical (e.g. flow limits through an outlet structure) or statutory (e.g., meeting a downstream temperature or flow requirement).

HydroSCOPE is seamlessly integrated with the environmental performance model to allow environmental objectives to be used directly in the optimization process. HydroSCOPE is able to utilize both single-trace and ensemble forecasts from the hydrological forecast model. Output from HydroSCOPE provides a range of operational strategies to the day-ahead planning model in the form of daily release volumes to ensure that the day-ahead operations are optimally aligned with seasonal scale objectives. Flow rates, depths, velocities, and temperature are used by the environmental performance model for evaluating the environmental metrics. Future versions of HydroSCOPE will provide two-way communication with the day-ahead planning model where the day-ahead model supplies detailed 1–7 day operations to HydroSCOPE.

The simulation engine of HydroSCOPE is developed on a system-dynamics framework that focuses on the temporal dynamics at particular points in the system while managing the spatial detail in a way that minimizes the computational overhead. This approach requires more data processing up front, but it results in a simulation model that is very quick to run (even for multi-project, basin-scale simulations) and that can provide real-time feedback to the user. The model is written in the MATLAB programming language using an object-oriented approach and is converted to a C-code executable for integration into the larger toolset.

3.3.3 Mathematical Model

The reservoir model uses a one-dimensional (1-D), finite-difference approximation to simulate the vertical temperature profile and the outflow temperature. The model accounts for heat exchange across the reservoir surface, mixing at the inflow points and outlet structures, mixing due to wind shear and density gradients, and ice formation and melting. The stream-reach model uses a 1-D, finite-difference approximation in the horizontal direction, and accounts for heat exchange across the water surface as well as ice formation and melting. River routing is handled through an exponential material delay method that allows for quicker execution and coarser spatial resolutions. HydroSCOPE is typically run using a 6-hour timestep, although this can be changed as appropriate. A schematic of the HydroSCOPE conceptual model is shown in Figure 3-16.

The model consists of two components, reservoirs and rivers. Each reservoir is treated as a 1-D vertically discretized system while each river is split into a number of reaches that can range in length from 1 to 30 miles.

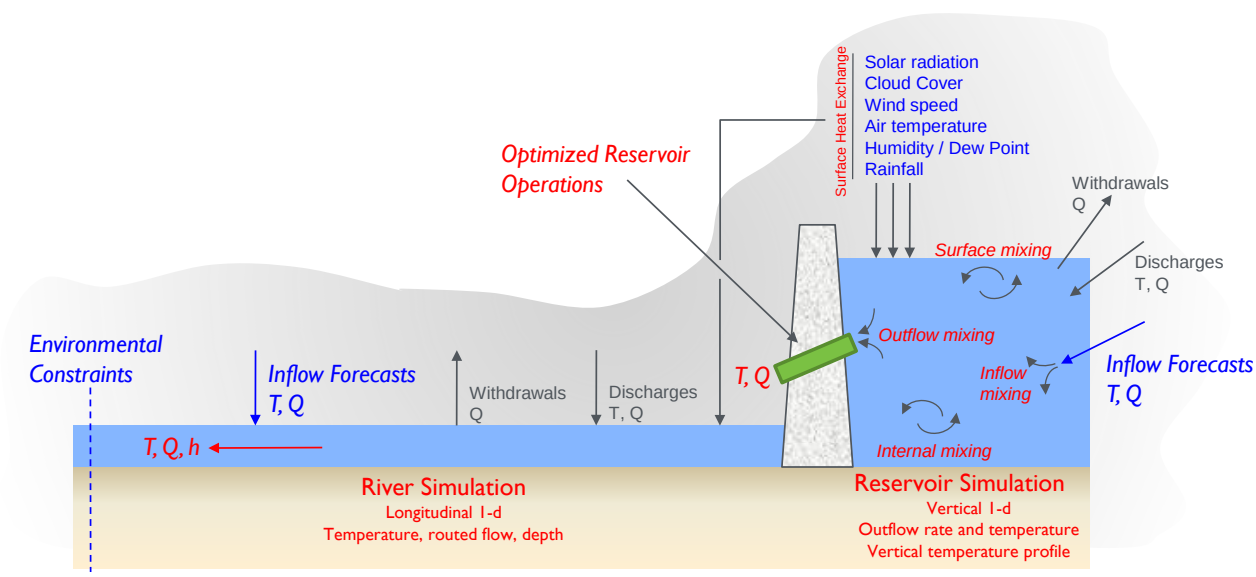

FIGURE 3-16 The HydroSCOPE Simulation and Optimization Tool Conceptual Model (Items in red are calculated by HydroSCOPE while items in blue may be dynamically supplied by other models within the larger toolset.)

The model is designed to simulate the heat and water balance throughout the system dynamically through time. Each component of the heat balance is comprised of many sub-processes that are defined by their governing equation(s), the system of which is solved with a finite difference scheme in both time and space.

3.3.3.1 Heat Balance

The form of the governing equation for the temperature within each layer of a 1-D reservoir system as used in the model is as follows:

$$\forall \frac{\partial T}{\partial t} = \Delta z Q_z \frac{\partial T}{\partial z} + \Delta z A_z D_c \frac{\partial^2 T}{\partial z^2} + Q_{in} T_{in} - Q_{out} T + \frac{A H_s}{\rho C_p} + T \frac{\partial \forall}{\partial t} \tag{3-1}$$

where T (°C) is the temperature, Q_z (m^3/s) is the vertical flow rate of water, z (m) is the space coordinate in the vertical direction, $\forall$ (m^3) is the volume, A_z (m^2) is the area as a function of depth, D_c (m^2/s) is the effective diffusion coefficient, ρ (kg/m^3) is the density, H_s [J/(m^2s)] is the heat exchange across the water surface, A_s (m^2) is the surface area of the air/water interface, Q_{in} [m^3/s] is the inflow to the reservoir, T_{in} (°C) is the temperature of the inflow, Q_{out} (m^3/s) is the reservoir outflow, and C_p [Ws/(kg °C)] is the specific heat of water. The model is capable of handling multiple inflows and outflows.

The first term on the right-hand side represents the vertical heat transfer due to advection. The second term is the transfer of heat due to dispersion. The third and fourth terms represent the heat flux from any inflows and outflows. The fifth and sixth terms represent the heat exchange across the air/water interface and change in heat due to changes in layer volume.

River reaches are assumed to be vertically mixed; thus, as applied to a river reach, the governing equation for temperature becomes the following:

$$\forall \frac{\partial T}{\partial t} = \Delta x Q_x \frac{\partial T}{\partial x} + Q_{in} T_{in} - Q_{out} T + \frac{A H_s}{\rho C_p} \tag{3-2}$$

where x (m) is the space coordinate in the longitudinal direction (i.e., in the direction of flow) and all other terms are defined above. The inflow terms include tributary inflow, inflows from industrial and/or municipal discharges, and inflows from groundwater. Outflows include withdrawals for municipal, agricultural, and/or industrial use, and losses to groundwater. The dispersion term in the longitudinal direction is negligible as compared to the other terms and is ignored for the exponential materials delay river routing approach explained below. The advective term represents the net heat flux as a result of the flow in the river. Note that Q_x on the upstream side can be different than Q_x on the downstream side if any inflows or outflows are included.

The surface heat exchange for both rivers and reservoirs is calculated as follows (Thomann and Mueller 1987):

$$H_s = H_{LW} + H_{SW} - H_{BR} - H_C - H_E + H_{RF} \tag{3-3}$$

where H_{LW} is the longwave radiation, H_{SW} is the shortwave radiation, H_{BR} is the back radiation, H_C is the conductive heat flux, and H_E is the evaporative heat flux. The units for each term are given in W/m^2.

The heat flux from longwave radiation is calculated using Brunt's equation, as implemented in the CE-QUAL-W2 model (Cole and Wells 2006):

$$H_{LW} = \begin{cases} 5.31x10^{-13} T_{aK}^{6}\left(1+0.0017C_c^2\right)\varepsilon & \text{for} \quad T_a \geq 5\ °C \\ 5.62x10^{-8} T_{aK}^{4}\left[1-0.26\exp\left(-7.77x10^{-4}T_a^2\right)\right]\left(1+0.0017C_c^2\right)\varepsilon & \text{for} \quad T_a < 5\ °C \end{cases} \tag{3-4}$$

where T_{aK} is the air temperature in degrees Kelvin, T_a is the air temperature in degrees Celsius, C_c is the cloud cover (on a scale of 0–10, zero being clear and ten being completely cloudy), and ε is the emissivity of water (0.97).

Shortwave radiation at the water surface is supplied by the hydrologic forecasting tool and is calculated as a surface component and a penetrating component. The surface component accounts for the rapid absorption of heat within the first foot or so of the surface while the penetrating component represents the solar radiation that penetrates beyond the initial absorption zone. Both components are described as exponentially decreasing with depth, with the surface component using an extinction coefficient that is much larger than the penetrating extinction coefficient. The equation for shortwave radiation is given as follows:

$$H_{SW} = \left[\left(1-S_a\right)\exp\left(-\lambda z\right) + S_a \exp\left(-\lambda_s z\right)\right]S_r \tag{3-5}$$

where S_a is the percentage of incoming solar radiation that is absorbed in the upper layer, λ (1/m) is the penetrating light extinction coefficient, λ_{ss} is the surface light extinction coefficient, and S_r (W/m^2) is the incoming solar radiation. A value of 10/m for λ_{ss} results in approximately 95% of the surface component solar radiation being absorbed in the first foot of water. The incoming solar radiation, S_r, is a function of time and is adjusted for the rivers to account for shading effects from the adjacent topography and streamside vegetation. Shading effects for the reservoirs are assumed to be negligible. The shading percentages are calculated as a pre-processing procedure prior to running the model. Since the streams are assumed to be vertically mixed, 100% of the shortwave radiation is added to the streams with no accounting for extinction in the vertical direction (i.e., $H_{sw} = S_r$). For shallow streams where light penetrates to the bottom, an assumption is made that an amount of heat equivalent to the solar radiation reaching the bottom is radiated back to the water in the form of longwave radiation.

Back radiation is calculated as follows:

$$H_{BR} = \varepsilon \sigma T_s^4 \tag{3-6}$$

where T_s is the temperature of the water at the surface, σ is the Stefan-Boltzmann constant (5.6704×10^{-8} [W/(m^{2}K^4)]), and ε is defined above.
The conductive heat flux is calculated as follows:

$$H_C = \beta W_f (T_s - T_a) \tag{3-7}$$

where β is the is the Bowen ratio set as a unit-less constant equal to 0.47, and W_f is the wind function given by the following:

$$W_f = a + bW_s^c \tag{3-8}$$

where W_s [m/s] is the wind speed at 2 m above the water surface, and a, b, and c are coefficients usually given by a = 4.6, b = 9.2, and c = 2.

The evaporative heat flux is given by the following:

$$H_E = W_f (e_s - e_a) \tag{3-9}$$

where e_S and e_a (mmHg) are the saturation and atmospheric vapor pressures, respectively. The saturation vapor pressure is calculated using the following:

$$e_s = \begin{cases} \exp\left[2.3026\left(7.5T_s/\left(T_s + 237.3\right) + 0.6609\right)\right] & \text{for} \quad T_s > 0\ °\text{C} \\ \exp\left[2.3026\left(9.5T_s/\left(T_s + 265.5\right) + 0.6609\right)\right] & \text{for} \quad T_s \le 0\ °\text{C} \end{cases} \tag{3-10}$$

where T_S (ºC) is the surface water temperature. The atmospheric vapor pressure is supplied to the model from the hydrologic forecasting tool or can be calculated as a pre-processing exercise using the following:

$$e_a = \begin{cases} \exp\left[2.3026\left(7.5T_d/\left(T_d + 237.3\right) + 0.6609\right)\right] & \text{for} \quad T_d > 0\ °\text{C} \\ \exp\left[2.3026\left(9.5T_d/\left(T_d + 265.5\right) + 0.6609\right)\right] & \text{for} \quad T_d \le 0\ °\text{C} \end{cases} \tag{3-11}$$

where T_d [°C] is the dewpoint temperature.

3.3.3.2 Water Density

Water density is calculated using the following regression equation:

$$\rho = ((((a_1 T - a_2)T + a_3)T - a_4)T + a_5)T + a_6 \tag{3-12}$$

where ρ (kg/m^3) is the water density, T (ºC)is the water temperature, and a_1 through a_6 are regression coefficients given by the following:

$a_1 = 6.536332 \times 10^{-9}$
$a_2 = 1.120083 \times 10^{-6}$
$a_3 = 1.001685 \times 10^{-4}$
$a_4 = 9.09529 \times 10^{-3}$
$a_5 = 6.793952 \times 10^{-2}$
$a_6 = 999.842594$

3.3.3.3 Withdrawal Dynamics

The range of depths that contribute to the outflow from the reservoir (the "cone of withdrawal") is calculated using the method in Schneider et al. (2004) as a function of the outflow rate and the density gradient. This effect is calculated for a point sink in the model using the following:

$$d = \left(\frac{C_{bi} Q}{N}\right)^{0.3333} \tag{3-13}$$

and for a line source using the following:

$$d = \left(\frac{C_{bi} 2q}{N}\right)^{0.5} \tag{3-14}$$

where d (m) is the withdrawal zone half-height, Q (m^3/s) is the outflow rate, N (Hz) is the internal buoyancy frequency, q (m^2/s) is the outflow per unit width, and C_{bi} is the boundary interference coefficient. The outflow per unit width, q, is calculated using the width of the line sink.

The withdrawal zone half-height, d, is calculated twice, once in the positive (upward) direction and again in the negative (downward) direction from the withdrawal elevation structure. When calculated in a vertically layered scheme, d is determined using equations (3-13) or (3-14) for each layer moving away from the withdrawal layer and then comparing that value to the distance from the mid-elevation of the layer to the withdrawal structure elevation. The

final value of *d* in both directions is based on the first layer where the distance from the mid-elevation of that layer to the withdrawal structure elevation is less than *d*. If this condition never occurs, then *d* is given as the mid-elevation of the top and/or bottom layer. The boundary interference coefficient is equal to 2 if the outflow structure is within 10% of the total depth to the surface or bottom, and 1 elsewhere.

The internal buoyancy frequency, *N*, is defined as the frequency at which a parcel of water would oscillate if it was displaced from its original state. It is approximated using the following:

$$N^2 \approx \frac{g}{\rho}\frac{d\rho}{dz} \Rightarrow N = \sqrt{\left|\rho_z - \rho_0\right| g \Big/ \left[\left(E_z - E_0\right)\rho_0\right]} \tag{3-15}$$

where ρ_z (kg/m^3) is the water density at depth E_z (m), and ρ_0 (kg/m^3) is the water density at the mid-elevation of the outlet structure, E_0 (m).

The outflow velocities are determined by the following:

$$V_z = 1 - \left[\frac{\rho_z - \rho_0}{\rho_{L\max}}\right] \tag{3-16}$$

where V_z [m/s] is the normalized velocity at depth *z* (m), and ρ_{Lmax} is calculated as follows:

$$\rho_{L\max} = \max(\left|\rho_{d+} - \rho_0\right|, \left|\rho_{d-} - \rho_0\right|, N_z) \tag{3-17}$$

where ρ_{d+} (kg/m^3) is the water density at the positive half-height distance, ρ_{d-} (kg/m^3) is the water density at the negative half-height distance, and N_z is a non-zero constant to prevent division by zero equal to 10^{-20}.

For vertically layered schemes, the outflow rate from each layer is calculated by first calculating V_z for each layer and then using the following equation:

$$Q_k = Q\left(\frac{V_k}{\sum_{i=LL}^{UL} V_i}\right) \tag{3-18}$$

where Q_k (m^3/s) is the outflow from layer *k*, *Q* (m^3/s)is the outflow rate through the withdrawal structure, V_k (m/s) is the normalized velocity for layer *k* calculated with equation (3-16), and *UL* and *LL* are the layers coincident with d^+ and d^-, respectively. Q_k is substituted for each layer into equation (3-1) as Q_{out}.

Sample outputs from the algorithm for three cases that differ only in the elevation of the outlet are shown in Figure 3-17. Note that as the density gradient at the elevation of the outlet becomes smaller (the density gradient becomes smaller with depth for the cases shown in

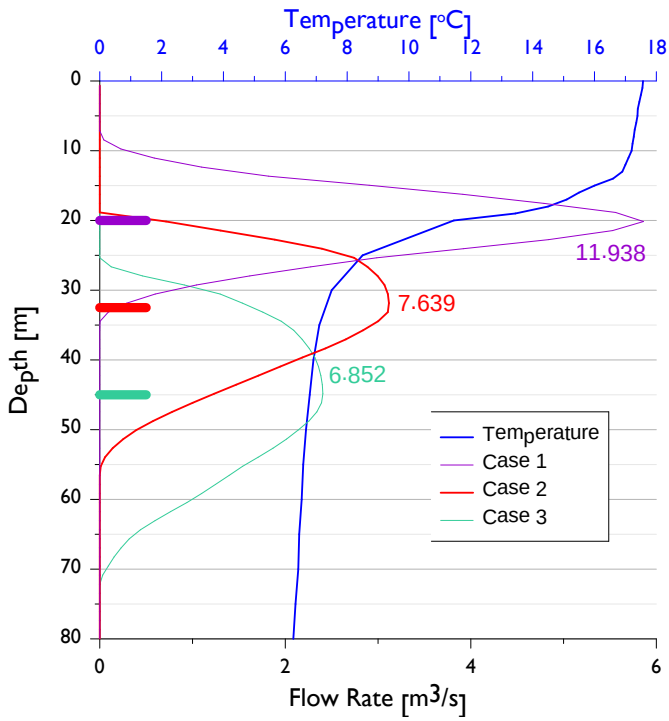

FIGURE 3-17 The Cone of Withdrawal for Three Outlet Structures at Different Elevations in a Stratified Water Column (The blue line shows the temperature profile with depth while the colored bars along the *y*-axis show the depth of the outlet structure. The numbers indicate the outflow temperature for the like-colored withdrawal cone in degrees Celsius.)

Figure 3-17) the cone of withdrawal widens. For the temperature calculation, the water from the withdrawal cone is assumed to fully mix as it passes through the outlet structure, resulting in a volume-density-temperature averaged outflow temperature. The outflow temperature for each case are also shown in Figure 3-17.

3.3.3.4 Inflow Dynamics

Inflow mixing is calculated similarly to outflow mixing, with the width of the mixing zone expressed as a function of the inflow rate, the inflow temperature, and the density gradient within the reservoir. The center line of the inflow zone is assumed to be the mid-elevation of the computational layer that most closely matches the inflow density. That layer is called the inflow layer, and is the layer from which the upper and lower boundaries of the mixing zone are calculated. Specifically, the inflow width is expressed as follows:

$$d = 1.35\left(\frac{Q_{in}L}{A}\frac{1}{\sqrt{g\frac{\Delta\rho}{\rho_o}}}\right)^{2/3} \tag{3-19}$$

where Q_{in} (m^3/s) is the inflow rate, L (m) is the effective length of the reservoir, A (m^2) is the area of the reservoir at the elevation of the inflow layer, $\Delta\rho$ (kg/m^3) is the difference in the density of the inflow layer and the upper (or lower) boundary layer, and ρ_o (kg/m^3) is the density of the inflow layer. The effective reservoir length is described as the length of the old river channel, since mixing tends to be confined in the horizontal direction to the width of the inflow stream. Thus, the value $Q_{in}L/A$ is the effective inflow per unit width.

Equation (3-20) is solved iteratively until the calculated half-zone width, *d*, is within one-half of an average layer thickness of the distance from the inflow elevation to the upper (and lower) boundary layer. If the density of the inflow is less than that of the surface layer, the inflow is assumed to enter the top layer. If the density of the inflow is greater than the bottom layer, then inflow is directed into the bottom layer. If there is no significant gradient in the reservoir, the inflows are assumed to enter all layers on a layer-volume weighted basis.

Once the upper and lower boundaries of the inflow mixing zone are calculated, the flow is distributed on a layer-volume weighted basis using the following:

$$Q_{in}(k) = Q_{in}\frac{V_k}{\sum_{kbot}^{ktop} V_i} \tag{3-20}$$

where $Q_{in}(k)$ (m^3/s) is the inflow to layer *k*, *kbot* and *ktop* are the upper and lower boundary layers, respectively, and all other terms are described above.

Since the stream reaches are vertically averaged, inflows to a stream reach are assumed to fully mix with the water in the receiving reach.

3.3.3.5 Density Mixing

Density mixing is handled using a density gradient approach (Environmental Laboratory 1995). Once the change in temperature (and thus the density) from the surface heat exchange and dispersion are calculated, the density of each layer is compared to the one beneath it. If the density difference between the upper layer and the next layer down is greater than a user-supplied stability factor, then those two layers are mixed and a new density and temperature for each of those layers is calculated. This is done iteratively until the density gradient is monotonically increasing (within a tolerance equal to a user defined minimum gradient) from top to bottom.

3.3.3.6 Diffusion Mixing

Diffusion mixing is a combination of molecular and turbulent diffusion and convective mixing that tends to dominate mixing in the epilimnion of most reservoirs. Mixing within the reservoir layers from diffusion is calculated using a combination of the "stability method" and the wind mixing method (Davis 1986). The stability method has shown to be suitable for deep, well-stratified lakes or shallow lakes where wind mixing is not the dominant mixing force; this is based on the assumption that mixing will be at a minimum when the density gradient (or the water column stability) is at a maximum. The wind mixing method, on the other hand, is appropriate for lakes where wind mixing is dominant and assumes that wind-induced mixing is greater at the surface and diminishes exponentially with depth. HydroSCOPE utilizes both methods by first calculating the effective diffusion coefficient with the stability method and then adding the wind mixing component to that. In addition, a limit is put on the rate at which the effective diffusion coefficient can change from one layer to the next. The effective diffusion coefficient is then used to solve the 1-D diffusion equation. Each of these steps are detailed below.

The first step calculates an interim diffusion coefficient using the following relationship between it and the stability:

$$\begin{aligned} D_c' &= A_2 S_c^{A_3} \quad \text{when} \quad S_c > C_s \\ D_c' &= A_1 \quad \text{when} \quad S_c \leq C_s \end{aligned} \tag{3-21}$$

where D'_c [m^2/s] is the interim diffusion coefficient, S_c (m^{-1}) is the stability, C_s (m^{-1}) is the critical stability, A_1 (m^2/s) is the maximum diffusion coefficient, and A_3 is a calibrated constant. The A_2 coefficient is calculated as a function of A_1, A_3, and C_s using the following:

$$A_2 = \frac{A_1}{C_s^{A_3}} \tag{3-22}$$

The stability is calculated as follows:

$$S_c = \frac{1}{\rho}\frac{\partial \rho}{\partial z} \tag{3-23}$$

where ρ (kg/m^3) is the water density, and z refers to the vertical dimension. The interim diffusion coefficient is limited to a minimum value of D_{min} (m^2/s).The second interim diffusion coefficient is calculated by adding the wind mixing component to D'_c using the following:

$$\begin{aligned} D_c'' &= D_c' A_4 w_s \exp(-A_5 z) \quad \text{when} \quad A_4 w_s \exp(-A_5 z) > 1 \\ D_c'' &= D_c' \quad \text{when} \quad A_4 w_s \exp(-A_5 z) \leq 1 \end{aligned} \tag{3-24}$$

where D''_c (m^2/s) is the second interim diffusion coefficient, w_s (m/s) is the wind speed 2 m above the water surface, z (m) is the depth of the water column, and A_4 (s/m) and A_5 (m^{-1}) are calibrated constants.

Lastly, the effective diffusion coefficient is limited to a user-defined rate of increase with depth using the following:

$$\begin{aligned} D_c(z) &= A_6 D_c(z+\Delta z) \quad \text{when} \quad D''_c(z) > A_6 D_c(z+\Delta z) \\ D_c(z) &= D''_c(z) \quad \text{when} \quad D''_c(z) \le A_6 D_c(z+\Delta z) \end{aligned} \tag{3-25}$$

where D_c (m^2/s) is the effective diffusion coefficient at elevation z (m) in the water column, Δz (m) is an offset to the elevation, and A_6 is a calibrated constant.

The effective diffusion coefficient is calculated for each layer in a reservoir at each time step and is used to calculate the heat transfer between layers in the vertical direction due to molecular and turbulent diffusion, convective mixing, and wind shear on the surface of the reservoir, and as input to the second term on the right-hand side of equation (3-1).

3.3.3.7 Volume/Area/Elevation Curves

The model uses the elevation of the water surface as the default variable when describing the initial state of a reservoir. For calculating the water balance, the volume and surface area of the reservoir (and each of the computational layers) is calculated from the elevations using a three-column lookup table where the first column represents the elevation, the second column the volumn, and the third column the area. The lookup tables are constructed as a pre-processing exercise and entered into the model via the GUI interface.

3.3.3.8 River Routing

The model uses an exponential material delay estimation (EMDE) for routing water through the river reaches. The EMDE (Richardson and Pugh 1981) is based on the volume of water within each river reach. It is semi-empirical because it ignores the physics of water transport in its calculation (however, the parameterization of the EMDE input variables due to account for the physics). Initially developed for modeling supply-chain inventories using system dynamics (Kirkwood 1998), the EMDE provides a basis for routing water through large river reaches in a manner that is quick to execute. The method requires input values for a time lag and a delay order where the time lag is defined as the river reach length divided by the wave celerity, and the order refers to the number of internal levels or sub-reaches that are used in the calculation.

The EMDE method uses a relationship between the flow rate, residence time, and storage, defined as follows:

$$\frac{dS}{dt} + \frac{S}{t_{lag}} = Q_{in} - Q_o \tag{3-26}$$

where S (L^3) is the storage in the reach, t_{lag} (t) is the time lag, Q_{in} (L^3/t) is the total inflow rate, and Q_o (L^3/t) is the initial inflow rate. The time lag, t_{lag}, is defined as follows:

$$t_{lag} = aS^b + c \tag{3-27}$$

where a, b, and c are power-law coefficients and the other terms are described above. The time lag can be calculated from representative rating tables for each stream reach. This is explained in more detail for the demonstration sites. The storage, S, is the net storage and represents the difference in storage between the current time step and the initial condition. Thus, it is always initialized to a value of zero. Equation (3-26) is solved using finite differences, which results in the following equation:

$$S(j)^t = [Q(j-1)_{in} - Q(j)_{out}]\Delta t + S(j)^{t-1} \tag{3-28}$$

where S^t is the storage at the current timestep, S^{t-1} is the storage at the previous timestep, Q_{out} (L^3/t) is the outflow from the reach, Δt is the simulation timestep, j (unitless) is the order number where $j = 1$ to n where n is the delay order. All other terms are described above. For cases where $n = 1$, $Q(j-1)_{in} = Q_{in}$ = inflow at the top of the reach. For cases where $n > 1$, $Q(j-1)_{in} = Q_{in}$ for the first sub-reach ($j = 1$), and the outflow from the previous subreach for $j = 2...n$.

The outflow from each reach is calculated using the following:

$$Q(j)_{out} = {S(j)^{t-1}}/{t^*_{lag}} + Q_o \tag{3-29}$$

where t^*_{lag} is the time lag divided by the delay order (t_{lag} /n) and all other terms are described above. To ensure numerical stability, $t^*_{lag} \geq \Delta t$.

The application of the EMDE method to river routing is a new and unique approach; as such, the method was tested against three different scenarios and compared to the same scenarios as modeled using HEC-RAS (Brunner 2010). The first test case considered a trapezoidal channel of constant slope and geometry. The second case uses a trapezoidal channel with three distinct reaches, each with their own hydraulic properties, slopes, and geometries. The third case is based on the Gunnison River below the Crystal reservoir, which uses the rating curves from three USGS gauges to distinguish three distinct reaches. For the first and second cases, a truncated sinusoidal function is used as the inflow hydrograph at the top of the uppermost reach (Reach 1). The third case uses hydrograph data from the USGS gauges from 1999 and includes side flows

from major tributaries (the North Fork of the Gunnision River, the Uncompahgre River, and combined unaged inflows). The third test case is compared to both a HEC-RAS model and to a USGS gauge hydrograph. The reach geometries and hydraulic properties for each test case are listed in Table 3-2. Plots of the reach geometries are shown in Figure 3-18.

The EMDE simulations were run for 1 year starting January 1, using a 6-hour timestep. The HEC-RAS simulation timesteps ranged from 20 minutes to 1 hour and were set to the maximum value that avoided oscillations in the solution. The discharge hydrograph from Reach 3 is compared for each of the test cases. For Test Case 3, comparisons are also made to hydrograph data from USGS gauges 09152500 and 09144250. The results for all of the comparisons are shown in Figures 3-19 through 3-22. In all cases, the EMDE method matches the HEC-RAS simulations very well with the average errors equal to 0.11%, 0.08%, and 0.74% for Test Cases 1, 2, and 3, respectively. In addition, a good match is seen between the USGS hydrographs and both the HEC-RAS model and the EMDE method.

3.3.3.9 Inflow Temperatures

For cases where temperatures are not available for inflows to the system, the temperatures are calculated using a regression equation developed by Bartholow (1989). The regression calculates the temperature of the inflows as a function of the inflow rate, air temperature, and time of year. The regression coefficients are established by fitting nearby data where flow, temperature, and air temperatures are available. The regression equation is given as follows:

$$T_j = A_0 + A_1 T_{aj} + A_2 ln(Q_j) + A_3 sin\left(\frac{2\pi j}{365}\right) + A_4 cos\left(\frac{2\pi j}{365}\right) \quad (3\text{-}30)$$

where j is the Julian day from the beginning of the year, T_j (ºC) is the inflow temperature at time j, Q_j (m^3/s) is the inflow rate at time j, and A_0, A_1, A_2, A_3, and A_4 are calibrated coefficients.

3.3.3.10 Optimization

HydroSCOPE couples the simulation model with an optimization software package called DAKOTA, which is an open-source software toolkit created at SNL (Adams 2010) to perform optimization, uncertainty quantification, and sensitivity analysis. The framework for communication between the simulation model and DAKOTA is shown in Figure 3-23. The optimization objectives are defined in the user interface, as are any operational rules and/or constraints. At each optimization iteration, a release schedule is fed to the model, which simulates the effects over time of implementing that particular schedule. Once the simulation is complete, DAKOTA compares the model output to the objectives and constraints and then creates a new set of release schedules to be fed to the model. The process is repeated until the improvement between iterations becomes smaller than a user-defined threshold. While

TABLE 3-2 Listing of Reach Geometries and Hydraulic Properties for Testing the EMDE River Routing Method

	Test Case No.						
	1	2			3		
Property	Reach 1	Reach 1	Reach 2	Reach 3	Reach 1	Reach 2	Reach 3
Reach N9o.	1	1	2	3	1	2	3
Slope (%)	0.0001	0.009	0.002	0.001	0.00933	0.00197	0.00122
Length (km)	150	50	50	50	50	50	50
Mannings n	0.050	0.085	0.050	0.030	0.085	0.037	0.040
Left Slope (W/H)	12	1	1	12	6	8.5	10.5
Right Slope (W/H)	24	1	12	24	6	8.5	10.5
Bottom Width (m)	12	57.660	52.290	17.880	22.600	46.250	50.00
EMDE							
t_{lag} a coef.[a] (s)	4.193×10^8	6.265×10^8	8.586×10^8	8.089×10^7	3.044×10^8	2.284×10^8	4.832×10^8
t_{lag} b coef.[a]	–0.418	–0.675	–0.683	–0.542	–0.666	–0.680	–0.603
t_{lag} c coef.[a] (s)	80025.41	1748.54	8162.50	12458.51	9228.43	4397.49	11225.91
EMDE Order	5	1	1	1	1	1	1

[a] See equation (3-27).

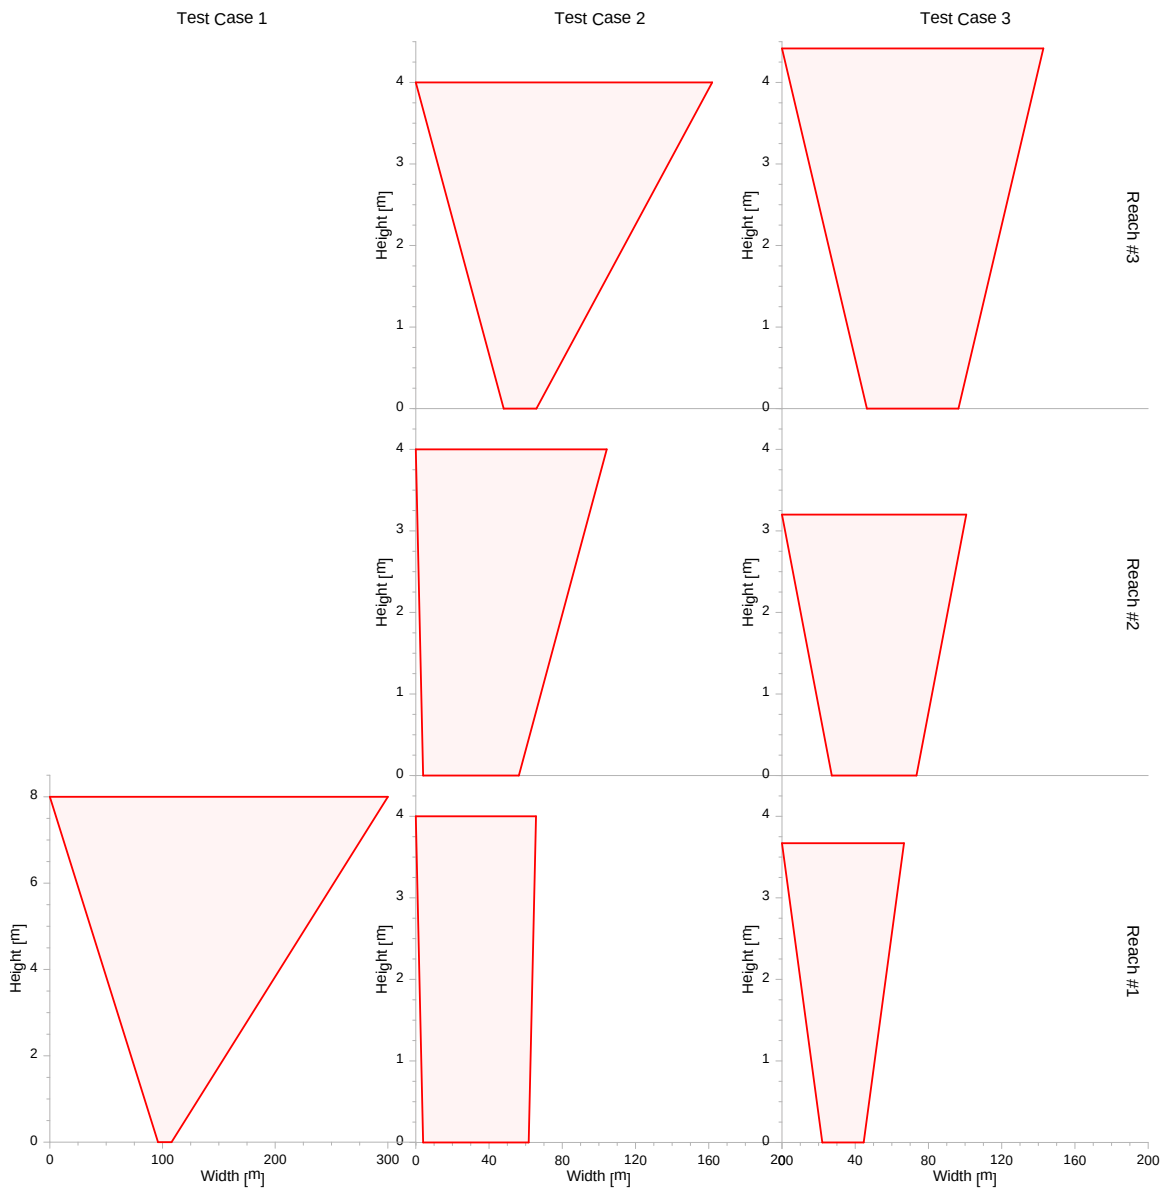

FIGURE 3-18 Geometric Configuration for Each Test Case and Reach (Reach 1 is at the top end of each river.)

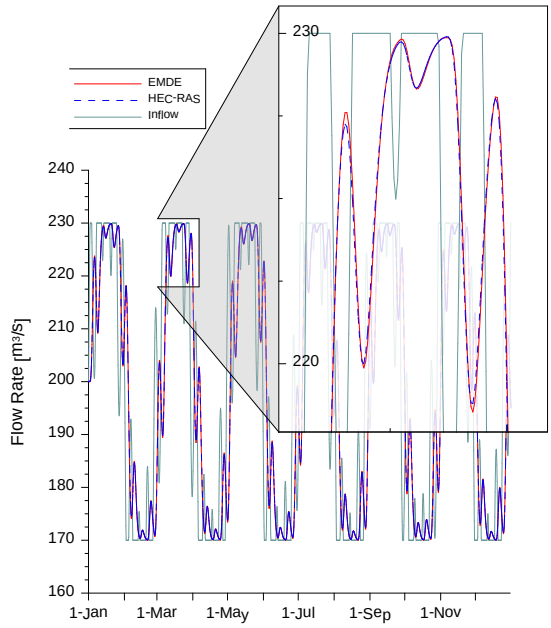

FIGURE 3-19 Test Case 1 Comparison Plot of the Inflow Hydrograph and the Discharge from Reach 1, Calculated by the EMDE Method and a HEC-RAS Model (Test Case 1 has only one reach)

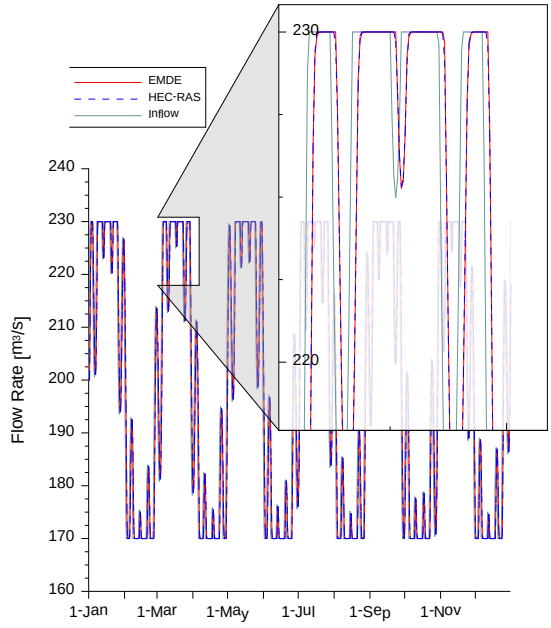

FIGURE 3-20 Test Case 2 Comparison Plot of the Inflow Hydrograph and the Discharge from Reach 3, Calculated by the EMDE Method and a HEC-RAS Model

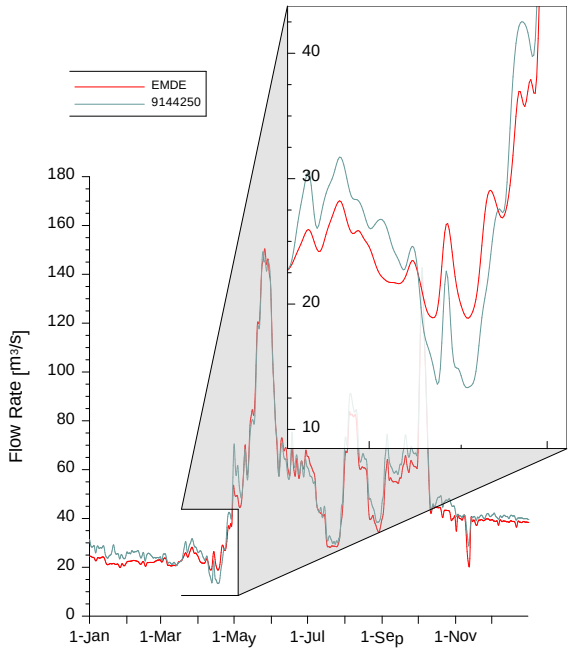

FIGURE 3-21 Test Case 3 Plot of EMDE Calculated Discharge from Reach 2 Compared to the USGS Gauge 09144250 Hydrograph

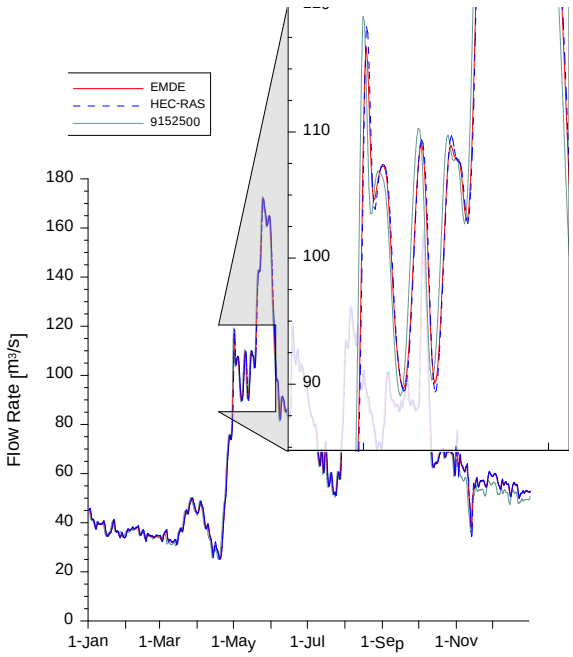

FIGURE 3-22 Test Case 3 Plot of EMDE Calculated Discharge from Reach 3 Compared to the HEC-RAS Model and the USGS Gauge 09152500 Hydrograph

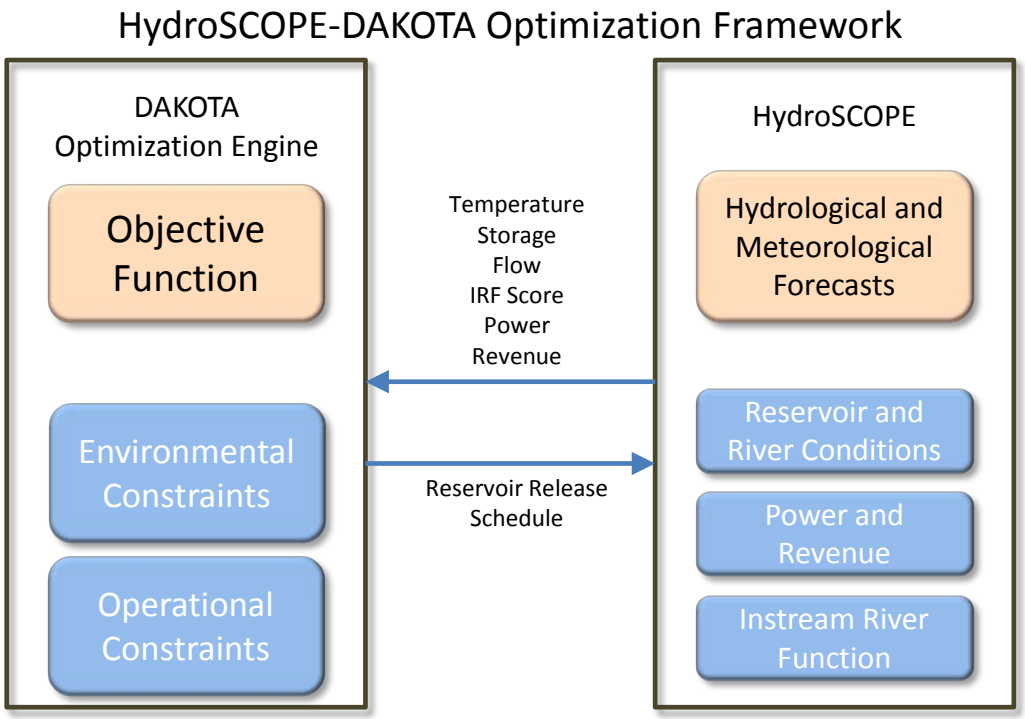

FIGURE 3-23 Communication Structure of the HydroSCOPE-DAKOTA Framework

DAKOTA can implement many different optimization algorithms, its implementation here uses single- and multi-objective genetic algorithms (SOGA and MOGA).

HydroSCOPE is designed to work with any of the various optimization schemes and is set up to solve the following general optimization problem:

$$\begin{aligned} &\text{minimize} && F(\underline{s}, t) \\ &\text{subject to} && \underline{h}(\underline{s}, t) = 0 \\ & && \underline{g}(\underline{s}, t) > 0 \\ & && \underline{S}^{L} < \underline{s} < \underline{S}^{U} \end{aligned}$$

where $F(\underline{s}, t)$ is the objective function or functions at time t, $\underline{s}$ is a vector of bounded state variables, $\underline{h}(\underline{s}, t)$ is a set of equality constraints, and $\underline{g}(\underline{s}, t)$ is a set of inequality constraints. The state variables for the hydropower problem are release schedules and can be bounded by both maximum and minimum flows and ramp rates.

The objectives and constraints used within the optimization tool depend on the needs of the site, and are likely to include maximizing revenue, maximizing power production, minimizing environmental damage, or maximizing the environmental performance value (i.e., the Index of River Functionality [IRF]) score. The optimization scheme can include multiple objectives, one objective, or a single objective that combines multiple concerns into one equation. For example, an objective function that seeks to maximize revenue might also

include a component that penalizes the result if an environmental objective is not met. Constraints in HydroSCOPE might include operational or environmental constraints such as minimum flows, maximum ramp rates, and flood control requirements.

To run an optimization, the problem is defined in the user interface and automatically populates the four files that run the optimization problem. The optimization formulation, including methods and decision variables, are written to a DAKOTA input file in the required format. A set of "wrapper" files are used to kick off the optimization simulations, handle the necessary communication requirements between HydroSCOPE and DAKOTA, and do any post-processing that may be needed to calculate the objective function.

The decision variables in the current configuration are a set of multipliers on the default release schedules for each reservoir in the system. Daily multipliers are used for the first week of the time horizon, weekly multipliers are used for the next 3 weeks, and monthly multipliers are used after that. For example, a 3-month time horizon would include 12 decision variables for each reservoir: day 1, day 2, day 3, day 4, day 5, day 6, day 7, week 2, week 3, week 4, month 2, and month 3. An exaggerated illustration of the multipliers is shown in Figure 3-24.

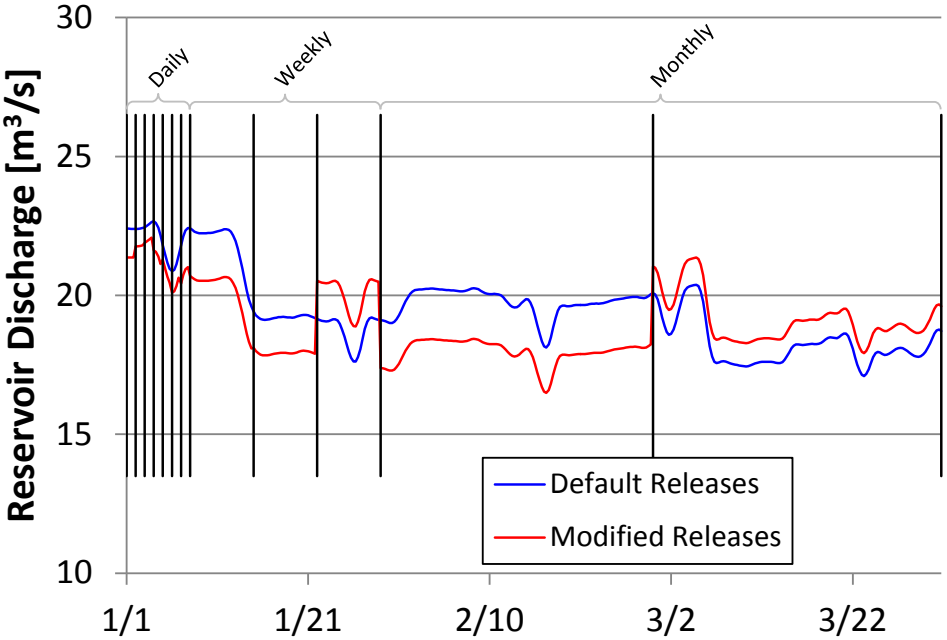

FIGURE 3-24 Illustration of the Daily, Weekly, and Monthly Multipliers Used for Optimizing the Reservoir Release Schedules (The example is exaggerated in that it shows much higher displacements from the default case than are typically found through the optimization.)

3.3.4 Data Requirements

The data requirements for HydroSCOPE are listed in Tables 3-3 and 3-4.

TABLE 3-3 Listing of Time Variant Input Variables to HydroSCOPE

Category	Description
Meteorological	Air temperature[a] Dew temperature[a] Cloud cover[a] Solar radiation[a] Wind speed[a] Rainfall[a]
Inflows	Tributaries[a] Outfalls[b]
Inflow Temperatures	Tributaries[c] Outfalls[b]
Outflows	Withdrawal structures[b] Diversions[b]

[a] Data may be dynamically supplied by models other than HydroSCOPE within the larger toolset.

[b] Data common to HydroSCOPE and at least one other tool in the WUOT.

[c] Data specific to HydroSCOPE.

3.3.5 Model Output

For each of the reservoirs in a simulation, HydroSCOPE output includes time series of total discharge rate and temperature, vertical temperature profiles, storage, elevation, evaporation rate, power generation, and power sales. The total discharge rate is further refined to include power and non-power generating discharge. For the rivers, HydroSCOPE output includes the flow, stage, and temperature at the bottom end of each river reach. Output for the optimization runs are generated from DAKOTA and include the objective function and decision variable values for each iteration of the optimization. In addition to the above, HydroSCOPE is able to return the IRF score as calculated by the environmental assessment tool (see Section 3.5 for a description of the environmental assessment tool) for any single simulation or optimization run.

The model was calibrated against data from 2009 for the Aspinall Cascade (described in Section 4.4.1.2) and 144 km of the Gunnison River below the cascade. Simulations from January 1, 2009, through December 31, 2009, were run for the full system and compared to data from the same time period. Figure 3-25 shows the calibrated vertical profiles from the three reservoirs as compared to field data collected at selected dates throughout the year. The calibration was conducted using the DAKOTA functionality of HydroSCOPE with an objective of minimizing the root mean squared error (RMSE) between the field data and the simulation

TABLE 3-4 Listing of Static Input Variables to HydroSCOPE[a]

Category	Description	Examples
Physical Characteristics	Dam/Reservoir	Spillway crest elevation
		Outlet capacities and elevations
		Plant configuration/capacities
		Storage/elevation/area curves
	River	Reach lengths
		Channel bathymetry
		Riparian vegetation density and height
		Channel bathymetry
		Flow rating curves
		Orientation, slope
Operating Characteristics/Constraints		Reservoir rule curves
		Pool elevations
		Unit power curve(s)
		Ramp rates
		Downstream flow and temperature requirements[b]
		Other environmental constraints[b]
Adjustable Parameters		Light extinction coefficients[b]
		Surface absorption percentage[b]
		Wind function coefficients[b]
		Ice parameters[b]
Initial Conditions		Reservoir vertical temperature profile[b]
		Reservoir elevation[b]
		River flow[b]
Calculated Variables[c]		River and reservoir temperature
		Reservoir storage/elevation/area
		Reservoir evaporation rates
		Power generation
		Power sales
		River elevation/flow
		IRF Score

[a] Example variables exemplify the type of variable needed for each category and may or may not apply to a particular site.

[b] Items specific to HydroSCOPE.

[c] Calculated variables are functions of time and space. Reservoir spatial variation is in the vertical direction while river spatial variation is in the longitudinal direction (i.e., in the direction of flow).

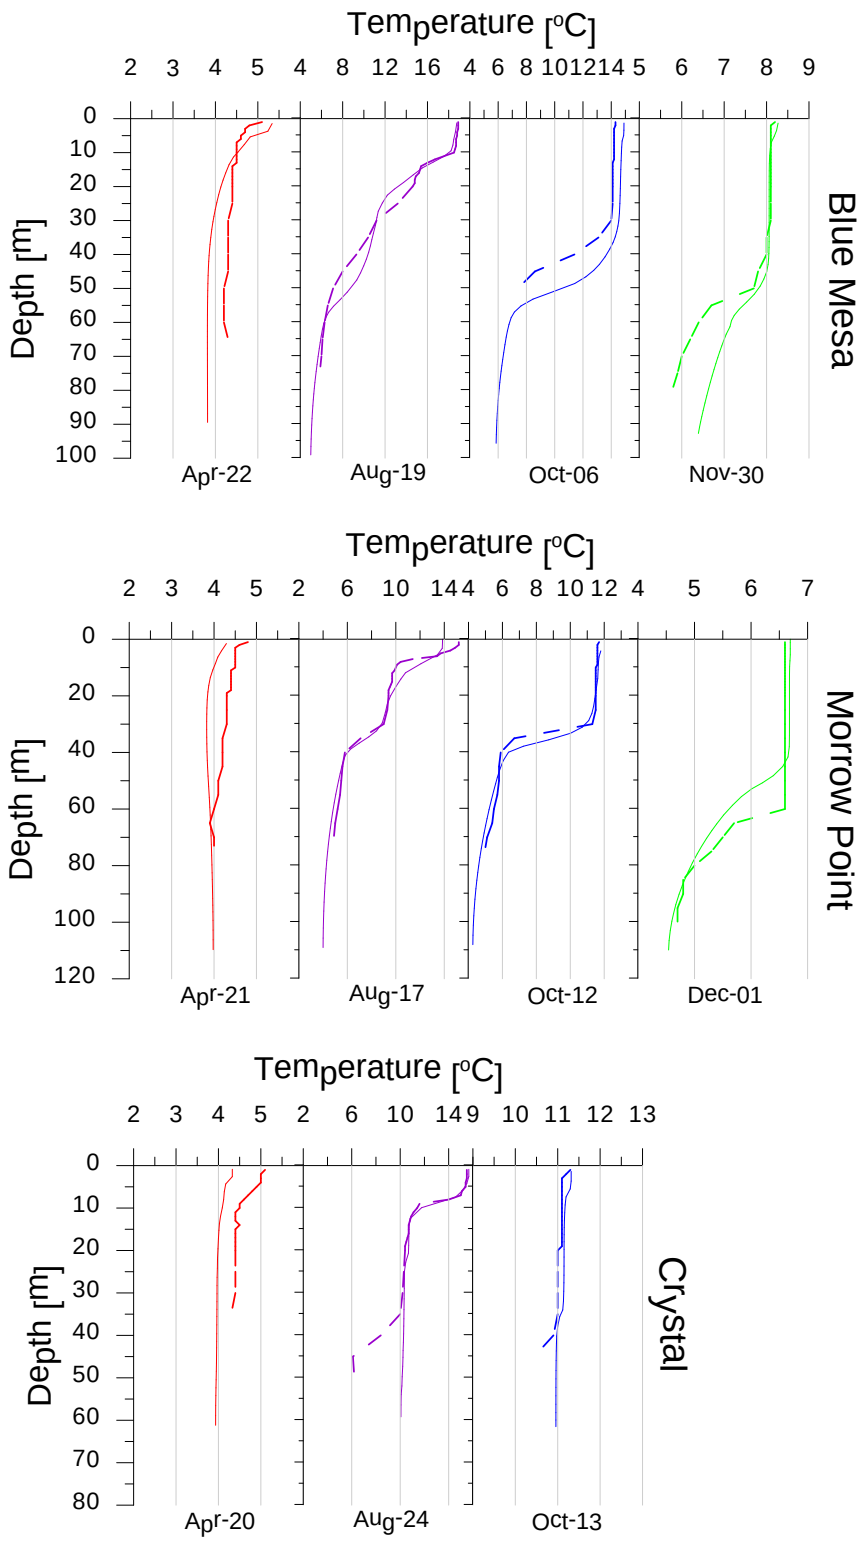

FIGURE 3-25 Calibrated Vertical Temperature Profiles for the Three Reservoirs of the Aspinall Cascade on the Gunnison River (The plot compares model output to field data collected in 2009.)

results. The decision variables (i.e., the calibration parameters) are the *a*, *b*, and *c* coefficients of the wind function (equation [3-8]), the solar extinction coefficients (equation [3-5]), and the minimum and critical diffusion coefficients and $A_{1\text{-}6}$ coefficients used to calculate the effective diffusion (equations [3-21]–[3-25]).

The river reaches were calibrated to USGS hydrograph data and, where available, temperature data. The hydrograph comparisons are shown above for 2009 in Figures 3-21 and 3-22. Temperature data were only available at the Whitewater gauge (09152500) for the daily maximum and minimum values. Comparisons of the gauge data to the model simulation are shown in Figure 3-26.

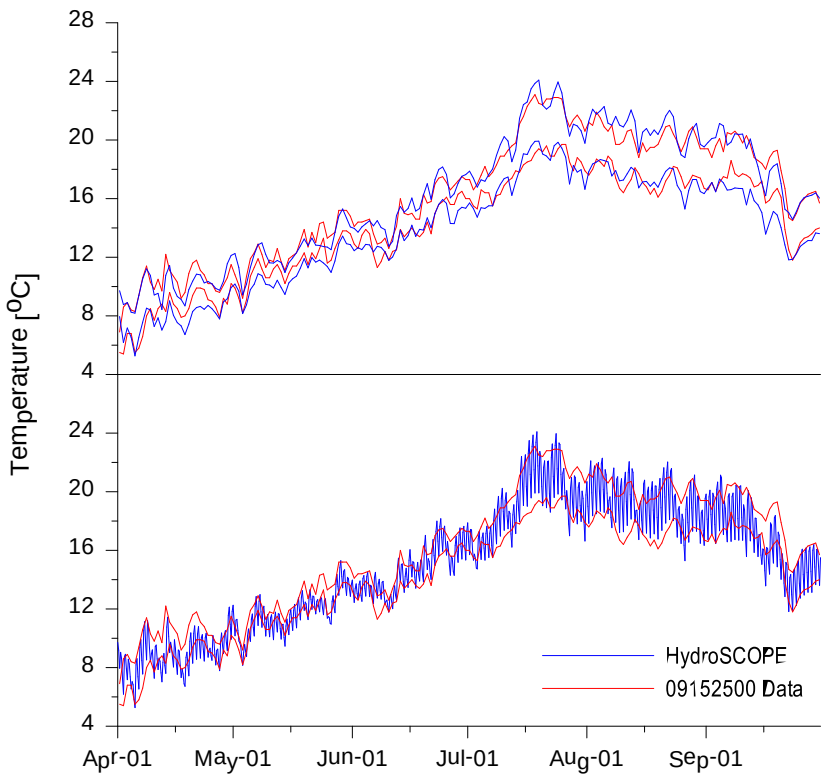

FIGURE 3-26 Plot of Maximum and Minimum Daily Temperatures (top) for HydroSCOPE and Data from USGS Gauge 09152500 for the 6-month Period from April through September 2009; Plot Showing the Daily Variation Produced from HydroSCOPE (bottom) as Compared to the Daily Maximum and Minimum Temperatures from the Gauge

3.4 ENVIRONMENTAL PERFORMANCE

3.4.1 Concept

This component of the WUOT, the IRF, is a methodology to evaluate how well particular discharge and thermal regimes resulting from hydropower operations meet desired ecosystem management goals. IRF provides an objective accounting tool for evaluating the environmental performance of rivers by considering site-specific environmental concerns, and it also facilitates the development and evaluation of tradeoff relationships between environmental and economic objectives. As a result, integration of the IRF tool with the overall WUOT can allow the development of optimized discharge regimes that take both environmental performance and power-generation economics into consideration during planning of reservoir releases. The intent is to provide the ability to explicitly co-optimize environmental performance and power revenue, rather than using environmental objectives solely as threshold boundary conditions. This would go beyond the status quo of minimally setting regulatory thresholds by helping to identify alternative operations that can meet power obligations while protecting and promoting environmental sustainability. The integration of the IRF with reservoir operation tools is intended to assist facility operators, natural resource managers, and other stakeholders with decision making and management of shared water resources.

3.4.2 Design

The IRF methodology (Saha and Hayse 2012) is based on the principles of river restoration, whereby a flow regime is evaluated (or developed for planning reservoir operations) by identifying components of in-stream flow (i.e., magnitude, timing, duration, frequency, and rate of change) suitable for accomplishing pertinent environmental objectives in a particular section of the river (Petts 1984; Ward and Stanford 1995; Richter et al. 1996; Poff et al. 1997; Postel and Richter 2003; Poff et al. 2010). This methodology requires the identification of applicable environmental concerns for a particular regulated section of a river. It requires users to input information describing the environmental needs by specifying the time period during which achieving each specific environmental objectives is relevant, and the frequency with which that objective needs to be accomplished. The water that will need to be allocated for accomplishing an objective within the specified time period is governed by defining relationships between water discharge and metrics used to evaluate each environmental objective.

The IRF tool was designed such that influencing factors within the context of flow management (both discharge and thermal regimes) could be specified, monitored, and regulated. An IRF software tool was developed that calculates an overall environmental score (IRF score) based on user-defined environmental objectives (Figure 3-27). Desired site-specific objectives such as spawning by an endangered fish species and seed germination in riparian areas are identified by users, together with the metrics necessary for characterizing the degree to which those objectives are met. For example, the amount of floodplain inundation specified in terms of discharge magnitude might be an applicable metric related to hardwood seed germination. Even

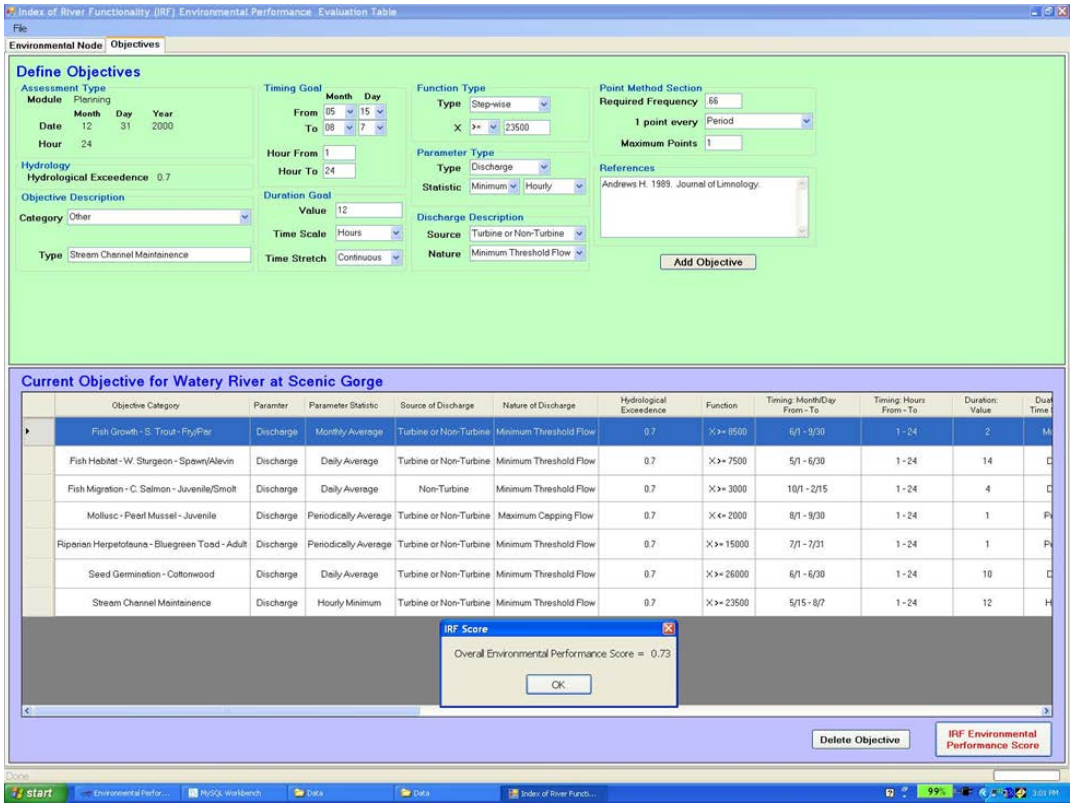

FIGURE 3-27 Screenshot of the IRF Tool Depicting the "Objectives" Tab, Which Allows Users to Define Environmental Objectives

though an environmental monitoring site is represented as a point location on the stream channel, it is meant to represent a reach of the river for which specific objectives are applicable.

The environmental objectives are user defined using relationships between the selected metrics and hydropower-influenced flow characteristics (e.g., discharge or water temperature) and consideration of seasonal timing, duration, and return frequency requirements for the targeted ecosystem parameters. Developing a relationship with water discharge is based on macrohabitat requirements (such as depth, water temperature, and/or substrate) for an environmental objective derived from scientific literature and/or information about the channel morphology (cross-sectional geometry, gradient, etc.) for that section of the river. Alternatively, if a model suitable for evaluating the relationship between an environmental objective and a flow parameter already exists (e.g., fish habitat suitability model, juvenile fish growth model), the results of the model can be used to identify relationships with discharge conditions at the monitoring location. The relationships between a flow parameter and the performance of an environmental objective can be described as either continuous (single or multiple linear) or a non-continuous function. While linearizing reduces the computational time during identification of the optimized solution, defining multiple linear relationships can be used to maintain the integrity of an original nonlinear relationship.

The example shown in Figure 3-27 indicates user-defined environmental objectives for a hypothetical site along with results displayed in the lower panel of the screen. Once an objective is identified, the user then specifies the appropriate metric to evaluate environmental performance of that objective. Multiple objectives that are relevant to a site can be explicitly defined in the IRF tool.

The IRF methodology allows performance scores for each objective to be calculated in various ways, which could include a cumulative score, a mean score, or a minimum score during the appropriate time period; the selected scoring aggregation method should be based on the nature of the environmental objective. The evaluation of overall environmental performance takes scores for each defined objective into account and designates environmental performance using an overall environmental score or IRF that can range from zero to one (zero being non-functional or not meeting objectives, one indicating maximal achievement of objectives, and intermediate values representing a partial [but potentially satisfactory] achievement of objectives). The overall IRF score updates at every time step for which flow conditions are input, and changes are based on the degree to which individual environmental objectives are being addressed. The IRF score represents environmental performance at a particular assessment time, although past achievement of environmental objectives during previous time periods is also taken into consideration by the scoring methodology.

The IRF methodology has the ability to accommodate a variety of environmental objectives, ranging from habitat requirements for a certain species and life stage of fish, to the maintenance of river channel conditions. For example, at a specific site, environmental objectives appropriate for addressing the requirements of various life stages of endangered, threatened, or migratory species of management concern that utilize the river channel could be identified and considered alongside objectives for riparian habitat maintenance through floodplain inundation (Eaton 1998; Hayse et al. 2005; Rood et al. 2005) as long as the appropriate limits on timing, duration, magnitudes, and rates of change in flow conditions are defined for each objective. Objectives pertaining to ecological resources affected by reservoir elevation can also be defined with the IRF tool. In addition, the source of the water release needed for accomplishing a specific environmental objective can be specified in the tool. For example, accomplishing an environmental objective such as out-migration of salmonid smolts may preclude a release of water through the turbines (because mortality would be increased) and would, instead, favor releases through other routes, thereby affecting power generation and optimization (Raymond 1979; Gibson and Myers 2002). As another example, the effect of selective withdrawal from the reservoir to meet downstream temperature requirements can also be investigated in the IRF tool; downstream temperature moderation and power generation reduction due to hydraulic head loss could be optimized in the WUOT. Users of the IRF tool would also have the ability to assign higher relative weights to comparatively more important objectives.

In addition to the information about environmental assessment nodes and the physical flow information, users of the IRF tool would need to provide input information to define each of the environmental objectives. This includes defining the following:

1. The relevant timing (e.g., seasonal time frame) during which an objective needs to be met;

2. The desired interannual return frequency for meeting the objective;

3. The frequency limit designating the maximum number of consecutive years that an objective can be ignored (after which its score goes to zero);

4. The duration (e.g., number of days or months) for which the appropriate conditions for the objective should be provided, an indication of the minimum and maximum durations needed, and whether the duration for which appropriate conditions need to be provided must be met continuously over consecutive time periods or as a cumulative duration within a longer time period;

5. The relative weight signifying the importance of that objective relative to other defined objectives;

6. A description of the relationship between a suitability metric for the objective and flow and/or temperature conditions; and

7. Text information about the objective (e.g., references and notes about assumptions).

The IRF tool provides a user interface that guides users in providing the appropriate information (Figure 3-27).

At designated time steps, the tool calculates the status of the user-defined metrics for each of the identified downstream environmental objectives based on flow and/or temperature at the evaluation points and identifies the degree to which environmental objectives have been accomplished on a scale from zero to one. Scores for multiple environmental objectives are combined into an overall score that can be used to evaluate relationships between environmental performance and power generation goals. Because the desired inter-annual return frequency for objectives is taken into consideration, objectives do not necessarily need to be met each year. However, minimum requirements for objectives must be achieved at least to some extent, even if the desired frequency is not fully met. Complete failure to meet any single objective (based on the specified frequency limit) results in an overall score of zero. The overall IRF score represents the degree to which the user-defined objectives are met by the hydrological regime.

3.4.2.1 IRF Modes

The IRF tool is designed to work in two modes: (1) historical mode or (2) planning mode. In the historical mode, the IRF tool can perform a standalone analysis doing retrospective evaluation of environmental performance based on user-defined objectives and historic flow

conditions. Planning mode, on the other hand, allows evaluation of environmental performance based on predicted releases, and would be the typical mode for integrated use with the other WUOT components.

The method used to calculate the overall IRF score differs depending upon the mode, with the planning mode using a prioritization routine based on desired frequency for meeting objectives, the amount of time that has passed since each objective was last achieved, and the degree of completion for each objective. Besides providing a performance score, planning mode provides the ability to improve future performance scores by providing incentives first to achieve the minimum target set for environmental objective scores and then to improve them by optimizing with non-environmental objectives under suitable hydrological conditions. Moreover, in planning mode users also have the choice to set a frequency limit for objectives that, if violated, would cause that objective to become moot. In the historical mode of analysis, degree of completion and time since last completion does not influence environmental performance. Because the past cannot be undone, prioritization does not occur in the historical analysis mode. Environmental performance scores in the historic mode are determined based on desired frequencies and actual observed frequencies, while ignoring any limits that may have been set on the frequency. While planning mode allows the comparative evaluation of multiple future discharge operations at any given assessment time, retrospective mode is used for evaluating historical environmental performance (objective and overall) as a function of time.

3.4.2.2 Algorithm Components of IRF Tool

There are several components that are involved in the calculation of the IRF score in planning mode. The major components along with the algorithm logic are described below.

Overall Environmental Score. The final or overall environmental score (IRF score) represents the degree to which a flow (discharge and thermal) regime meets the defined objectives within a specific river reach. The IRF score is measured on a relative scale from zero to one (zero being non-compliant [one or more objectives are not being met], one being fully compliant [all objectives maximally accomplished], and intermediate values representing partial accomplishment of objectives). The higher the individual objective scores are, the higher the overall score becomes. The IRF score is calculated as the geometric mean of the individual scores of all environmental objectives divided by the geometric mean of maximum scores of all environmental objectives (equation [3-31]). The formulation ensures that if even a single environmental individual objective score is zero, the overall score becomes zero. The intention of the formula was to indicate that a zero score for one environmental objective cannot be compensated for by a maximum (or relatively high) score of another objective. Ultimately, the IRF score is used to optimize environmental performance with power generation objective(s) using an optimization solver that is embedded within the WUOT. The overall IRF score is calculated as follows:

$$IRF = \left(\prod_{i=1}^{n} OS_i^{RW_i} \right)^{1/\sum_{i=1}^{n} RW_i} \tag{3-31}$$

where OS (Objective Score) is {OS_1, OS_2, ….., OS_n}; their corresponding RW (Relative Weights) is {RW_1, RW_2, ….., RW_n}; and IRF output range is [0, 1], 64-bit double precision.

Relative Weight. Relative weight allows users to assign a higher weight to an objective in comparison to other objectives. Under default conditions, all objectives have a relative weight of 1, signifying equal importance. The values of relative weights range from 1 to 5, where a value of 5 makes an objective five times more important than the objective with a base value of 1. For example, where environmental objectives related to stocked recreational fisheries and native endangered fish overlap, a higher relative weight for the latter objective could be applied.

Objective Score. The score for an individual objective is calculated based on the maximum score it received within the period specified by its required frequency from the assessment time. If a score is not obtained within that period, and the objective is not in term (T = 0), then the objective score is calculated by dividing the weighted score with the maximum score assigned for the objective. However, if a score is not obtained within the period specified by the desired frequency, and the objective is in term (T = 1), then the objective score is incentivized with the term points to achieve the minimum points required by the objective to receive a score. The objective score provides an indication of how well objectives are being addressed individually, as well as how well they are being addressed relative to other objectives at any given time. It has been formulated such that it cannot exceed the maximum possible score assigned for that particular objective. Objective score is calculated as follows:

$$OS_i = (WS_i / MS_i) + [Ti * \{1 - (WSi / MSi)\} * (TP_i / MnP_i)], \text{ when } IF_i = 1$$
$$\text{and } OS_i = (MXNS_i / MS_i), \text{ when } IF_i = 0, \tag{3-32}$$

where WS_i is weighted score of an individual objective; MS_i is maximum score assigned for that objective; Ti is in term or out of term (binary); TP_i is points accumulated by an objective while it is in term; MnP_i is minimum points required by an objective to get a score; $MXNSi$ is maximum nonzero score observed within the period starting from the assessment time back to the time as specified by the required frequency; IFi is importance factor (binary); and OS range is [0, 1], 64-bit double precision.

Weighted Score. Weighted Score for an individual objective is its adjusted score divided by its corresponding priority factor. The weighted score has been formulated such that the adjusted score of an objective is weighted by its priority factor, which assigns the objective its current score or a fraction of its last score, based on how long ago the objective was achieved (LO), how often a score is desired (RF), and its degree of completion (DC). LO, RF, and DC are

all components of priority factor (see below). Weighted score varies directly with adjusted score and inversely with priority factor. Weighted score is calculated as follows:

$$WS_i = (AS_i / PF_i), \tag{3-33}$$

where AS_i is adjusted score of an individual objective and PF_i is the priority factor value for that objective.

Maximum Score. Maximum score indicates a theoretical maximum possible score for a particular environmental objective, and is calculated based on the maximum possible points that can be accrued within the assigned period specified by the user for particular objectives under consideration. Alternatively, a user can specify a lower score that is less than the theoretical maximum possible score if it is deemed sufficient to impart the maximum desired value toward accomplishing the target objective. Maximum score is based on a ratio of maximum and minimum points and is represented relative to a minimum baseline score of one. Maximum score is calculated as follows:

$$MS_i = Min\ (A \vee B), \tag{3-34}$$

where A is [number of time-steps (e.g., hours, days) required to achieve maximum points for an objective] / [number of time-steps (in same time units as numerator) required to achieve a minimum point for that objective]; and B is user-specified maximum points (expressed as positive values).

Adjusted Score. The adjusted score for an environmental objective assigns the most recent nonzero score, then adjusts the score to allow hard constraints to be addressed, and provides insight into accomplishing a mandatory objective. Environmental concerns with hard constraints are defined here as those objectives that are deemed mandatory during all periods of relevance without failure and with a desired frequency of one. An example of a mandatory hard constraint would be a fish-passage requirement on a certain section of a regulated stream for out-migrating salmon smolts that has to be addressed every year, regardless of hydrological condition. The adjusted score prevents noncompliance with those mandatory constraints. The adjusted score has been formulated to take into consideration the last nonzero score achieved by an environmental objective, and assigns a zero score if it is not feasible to meet a hard constraint. Adjusted score is calculated as follows:

$$AS_i = (MRNS_i * HF_i), \tag{3-35}$$

where $MRNS_i$ is the most recent nonzero score of an individual objective; and HF_i is the hard constraint factor value for that objective, which is a binary coefficient where $HF = 0$, only if all three conditions are met (i.e., frequency limit – last occurrence = 0, term score = 0, and feasibility within term = 0), otherwise $HF = 1$.

Term. Term indicates the timing during which an environmental objective is relevant. It represents the window of opportunity for scoring during the year for that objective, as specified by relevant timing mentioned by the user. Term is a binary variable, where $T = 1$ for within term, and $T = 0$ for outside term.

Term Score. Term score indicates the score of a particular environmental objective during an ongoing term. Term score is calculated from term points and is represented in respect to the maximum score for a particular objective. In effect, this represents the portion of the maximum score (*MSi*) that has been achieved at the current moment:

$$TS = [(TP / MxP) * MS], \tag{3-36}$$

where *TP* is term points, *MxP* is maximum points, and *MS* is maximum score.

Priority Factor. There is a prioritization routine that calculates the priority factor for individual objectives by taking into account how often an environmental objective needs to be accomplished (a user-defined input), when the objective was last accomplished, the projected availability of water to accomplish that objective, and the degree to which an objective has already been completed during the current time period. Prioritization in the IRF provides a method for causing completion of an objective to become more pressing as the achieved frequency falls behind the desired frequency. The priority factor is intended to provide a means for prioritizing individual objectives to enhance the overall IRF score during optimization. The forcing for an objective is further increased if a significant amount of water has already been committed toward addressing that objective. Priority factor is calculated as follows:

$$PF_i = (RF_i * LO_i * DC_i)^{IF_i}, \tag{3-37}$$

where *RF* is required frequency, *LO* is last successful occurrence (in same time scale as RF), *DC* is degree of completion, and *IF* is importance factor (binary).

Required Frequency. The required frequency (RF) denotes how often an environmental objective needs to be accomplished. This is a user-specified value (range: 0.01–1; precision: two decimal places) that indicates an annual return frequency by specifying how often a score is desired for a particular objective. In most cases, the RF should be seen as depicting a desire for the environmental objective to be met with some overall frequency from a historic perspective. For example, if it is considered ecologically important that an event occur with a frequency of 0.2, this should be viewed as a desire to have this objective met in 20% of years over a longer multiyear period of record. An objective with a RF of 1 will also have a frequency limit of 1. This case (RF = 1 and FL = 1) should be considered to be a rule or hard constraint with discrete binary function, rather than an objective with continuous function and the possibility of incremental improvement or worsening in condition (refer to the Adjusted Score section for calculation of adjusted score for objectives with RF = 1). Even though these hard constraints (i.e., objectives with RF = 1) can be considered by the IRF tool, ideally these constraints need to

be applied in the seasonal and day-ahead real-time optimization tools so that only feasible solutions from the global solution space are generated from a fixed number of iterations to develop the pareto front for better identification of tradeoffs. Although using a RF value of 1 may be required for some objectives, using a value of 0.99 may be more desirable; this would only lower the overall score rather than resulting in an overall score of 0 during a noncompliant term, while still recognizing that the objective should be addressed in nearly every year (i.e., 99 out of 100 terms).

Last Occurrence. Last occurrence (LO) indicates the amount of time, in terms, that has passed since a particular environmental objective was last accomplished (i.e., minimum required points were obtained). LO is denoted by positive integers 1 through *n*. For example, if an objective was accomplished during the last possible term, then LO = 1; if it last got a score two possible time periods back, then LO = 2, and so on. LO is calculated by counting back from the present or most recent term to the last term when a nonzero score (minimum required points) was obtained. LO is expressed in the same time-scale units as RF. For an unvalidated objective (i.e. an objective that has never been accomplished, even to its mimimum level, during the period of record), users are provided the option to assume that the objective was fully met before the record to prevent the objective score (and, in turn, the overall IRF score) from being set to zero. This option is only available for objectives without a frequency limit and should be exercised with caution, but may be practical in cases where data exists only for the regulated period, even though it is certain that the objective was accomplished in the pre-regulated period. In this case, the objective is assumed to have been accomplished to its maximum level at time T_0, which refers to the time right before the period of record; LO in this case equals the number of terms observed in the period of record plus one. The larger the product of RF and LO is, the larger the PF value (which will result in a smaller WS for that objective) will be. In the overall formulation, this causes completion of an objective to become more pressing as the achieved frequency falls behind the desired frequency.

Degree of Completion. The degree of completion of an objective is calculated based on the volume of water that has already been used for accomplishing that objective. It designates extra weighting based on how close an objective is to being accomplished during the assessment time. The degree of completion is denoted by the ratio of the volume of water consumed to the amount of water required to accomplish the objective. The degree of completion has been formulated such that when an environmental objective is relevant and approaches completion, the degree of completion value provides an additional incentive to complete an objective once water has been used to start addressing that objective; as the objective gets closer to being accomplished, the incentive becomes greater. Degree of completion is calculated as follows:

$$DC_i = \left[1 + \left(\frac{U_i}{R_i}\right) - S_i\right]^{T_i} \tag{3-38}$$

where DC = degree of completion, U = volume of water that has already been used toward accomplishing an objective (as indicated by term points), R = volume of water required to accomplish an objective (as indicated by minimum points for obtaining objective score), S = state of completion for an objective (S = 0 for incomplete, and S = 1 for complete; if S = 1, then $U = R$), and T = term. The maximum range of U is R. The DC range is [1, 1.99], with precision up to two decimal places.

Importance Factor. The importance factor for an objective is a binary decision variable that indicates whether it is important to address an objective in the most recent ongoing term. If an objective has been accomplished and is within the desired frequency, based on how long it has been since it was last accomplished, then it is not important to address that objective in the current term. Importance factor is calculated as follows:

$$IF_i = 0 \text{ if } [RF_i * LO_i] \leq 1 \text{ and } IF_i = 1 \text{ if } [RF_i * LO_i] > 1 \qquad (3\text{-}39)$$

where IF is importance factor, RF is required frequency, and LO is last successful occurrence (in the same time scale as RF).

Frequency Limit. The frequency limit indicates the number consecutive years an objective can be ignored. This is also a user-input value (range: 1–∞; positive Z values). For example, an objective with a RF of 0.5 and a frequency limit of 7 indicates that even though the user desires that the objective be met once every other year, after 7 back-to-back years of not meeting it, its score would go to zero. Alternatively, if an objective is intended to occur at least once every 2 years without fail, then a frequency limit of 2 needs to be specified. It is not mandatory to specify a frequency limit for all objectives. Users have the discretion to use their professional judgment to decide whether a frequency limit is required for an objective. Without a frequency limit, the score reduces decrementally each year (after passing the desired frequency) the objective is not addressed, but it never goes to zero (even though it can become miniscule). Because of the continuously shifting time horizon in the future, objectives could be postponed forever, since it is impossible to calculate frequency without having finite bounds on the time. A frequency limit is required for some objectives, such as recruitment in an imperiled species, which if continuously ignored for an extended period of time may cause its extinction. Even though the desired frequency could be met by accomplishing that objective in several back-to-back years after an extremely long gap, the species would be lost forever during the long period of unfavorable conditions.

State of Completion. The state of completion indicates whether an objective has been accomplished during its assigned period. State of completion is denoted as a binary function, where it will be considered complete only when an environmental objective is relevant and has been accomplished (i.e., a score has been obtained for that objective). Alternatively, if an environmental objective is not relevant or if it has not been accomplished, then the state of completion will be considered incomplete. Essentially, a Kronecker delta function in the formula

causes the state of completion to designate a "complete" status, when the current score equals the maximum possible score for that objective. State of completion is calculated as follows:

$$S_i = [(K\delta_{X_i Y_i})\, T_i] \quad , \tag{3-40}$$

where S is state of completion, $K\delta$ is Kronecker delta function, X is minimum points required to get a score for an objective, Y is term points, and T = term (T = 1 if within term, and T = 0 if outside term).

Feasibility within Term. The feasibility within term indicates whether an environmental objective can still be accomplished during its period of relevance. It takes into account if an objective needs to be addressed on a continuous or cumulative basis during its period of relevance and then determines whether it is still possible to accomplish that objective within the period if it is not addressed in the current time step. It is a binary variable indicating the feasibility of accomplishing the objectives within the remaining period of the term. The feasibility within term factor also provides a small score (a score of 0.01 is maintained for all timesteps) to objectives with a RF of 1 that are feasible within the remainder of the term. This prevents a constraint from being met at the first possible timestep in its period; instead, a more conducive time can be identified to meet that constraint as long as it is feasible within the allotted term.

3.4.2.3 Calculating IRF Score Using IRF Planning Mode

The IRF planning tool denotes the objective and IRF scores at any future assessment time, and it also provides direction to improve the objective and overall environmental performance. Initially, points for an objective are calculated for each term from the assessment date going back in time for the period of time indicated by desired frequency based on observed historical records and the future planned operation. Points are calculated based on the user-defined relationships between the metric used to evaluate each objective and flow variable (such as discharge, stage, and water temperature) and other scoring conditions. Outside of the desired frequency, points are calculated back in time for each term up to the frequency limit, if one is defined for an objective; however, if a frequency limit is not defined for that objective, then the points for each previous term are calculated until the minimum points have been observed. In the planning tool, objectives are prevented from falling behind their frequency limit. The "feasibility within term" factor returns a non-viable (zero-value) solution even before the frequency limit is violated.

If an objective has been accomplished to the maximum level within the period indicated by its desired frequency, then the "importance factor" provides no incentive for water being spent in pursuit of that objective. However, if the maximum point is not achieved at the desired frequency, the highest available point is chosen within that period to compute the objective score; objectives with less than the minimum specified points do not get a score. Alternatively, if

an objective has never been accomplished within the period indicated by its desired frequency, the scoring algorithms in the planning tool provide incentives at each assessment time step within its term to accomplish the objective at its minimum level. If an objective falls behind its desired frequency (as indicated by last successful occurrence), its score is weighted by a "priority factor" that takes into account the objective's desired frequency and how long has it been since the objective was last accomplished. The more the objective falls behind its desired frequency, the more severe the score reduction. Moreover, if that objective is also within term, the priority factor further weights the score by taking into account the portion of the objective that has been completed. While the final objective scores are obtained by normalizing the weighted score of each objective with its maximum score, the overall IRF score is computed from the geometric mean of all objective scores. Objectives with higher relative weights contribute proportionally higher toward the overall IRF score.

3.4.2.4 Integration of the IRF Tool in the WUOT

Integration of the IRF with the other WUOT components can benefit many hydropower projects as it quantifies linkages between flow conditions and environmental responses that would result from planned hydropower operations. The general idea is to shape the future hydrograph to increase the joint-value function for environmental performance and power revenue (or power generation) over a finite time horizon (assessment date) within the constraints of best estimates of water availability and the physical operational limits, and then to incrementally shift the time horizon to the next future time step and repeat the process for a specified number of time steps. In the integrated WUOT, IRF helps identify tradeoffs by developing a Pareto frontier that is based on the relationship between environmental and power generation objectives (Figure 3-28).

The IRF integrated toolset could allow evaluation of the economic and environmental values of volumes of water at specific times; identification of tradeoffs will involve optimal shaping of that volume in terms of a release pattern. At any future time step, the optimization solvers of the WUOT can converge on an optimized solution that yields the maximum joint-value for an operation selected from a set of potential future operations, also considering the effect of historic operation (Figure 3-29). Thus the IRF-integrated WUOT provides a benefit-cost evaluation tool to assess the value of power generation and its associated environmental impacts for various types of operational scenarios.

The flexibility of the integrated toolset allows the users to assign different relative weights to various environmental and non-environmental goals, which can facilitate decision making regarding water use among multiple stakeholders. The IRF component of the WUOT also provides adaptive management capability for managing environmental resources based on a better understanding of the relationships between ecological needs and river regulation as a result of hydropower operation. The IRF integrated toolset should not be considered a prescriptive means for designating reservoir operation, but rather a way to comparatively evaluate the level of accomplishment of facility objectives and an indication of the potential environmental impacts from different types of hydropower operations. The data needs and sources for the IRF tool and their exchange with the other WUOT components are presented in Figure 3-30.

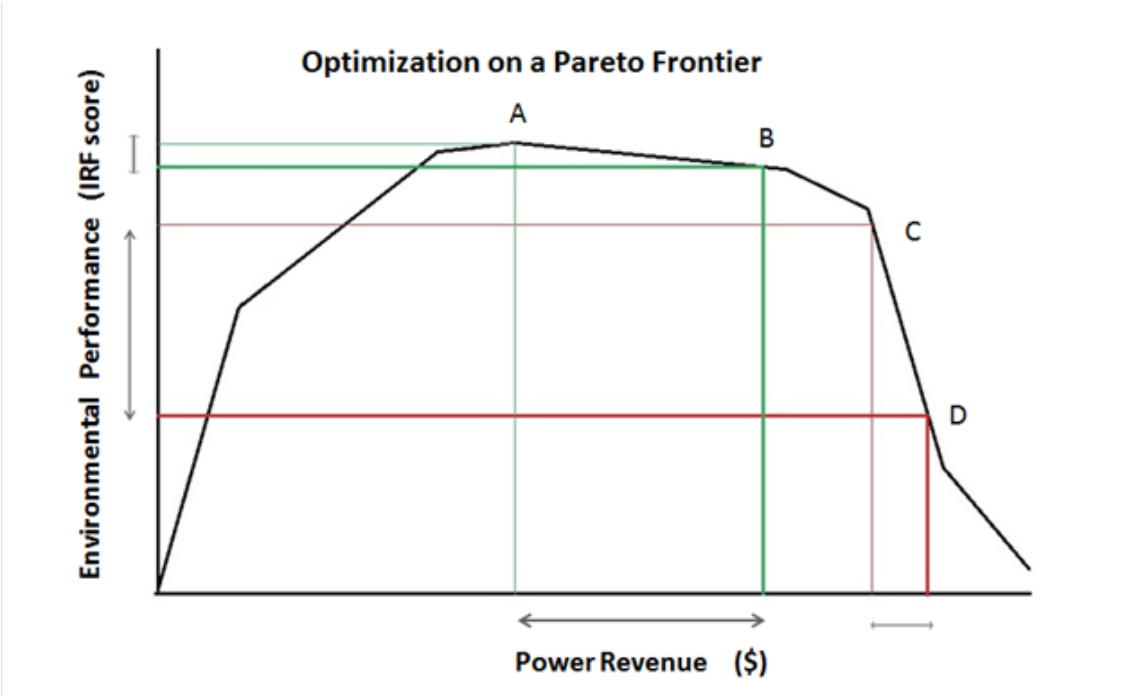

FIGURE 3-28 Hypothetical Example to Illustrate the Relationship between Environmental Performance and Power Revenue (The goal of water use optimization is to maximize joint benefits. So in this case, traversing from point A to point B is acceptable as power revenue significantly increases without greatly impacting environmental performance. However, it is undesirable to go from point C to point D because revenue is not significantly increased, even though it greatly compromises the environmental performance. This figure is an over-simplification of the relationship; in practicality, the relationship is much more complex, because the shape of curve would vary temporally and with specified objectives. Initially power revenue and environmental performance values are zero for both because in a hypothetical scenario where no water is released from the reservoir, it is likely that environmental objectives will not be met to any significant degree and that no power would be generated. This overgeneralized figure does not take into account aspects such as existence of flow bypass routes, cascade systems with pump-back operation, or ancillary services.)

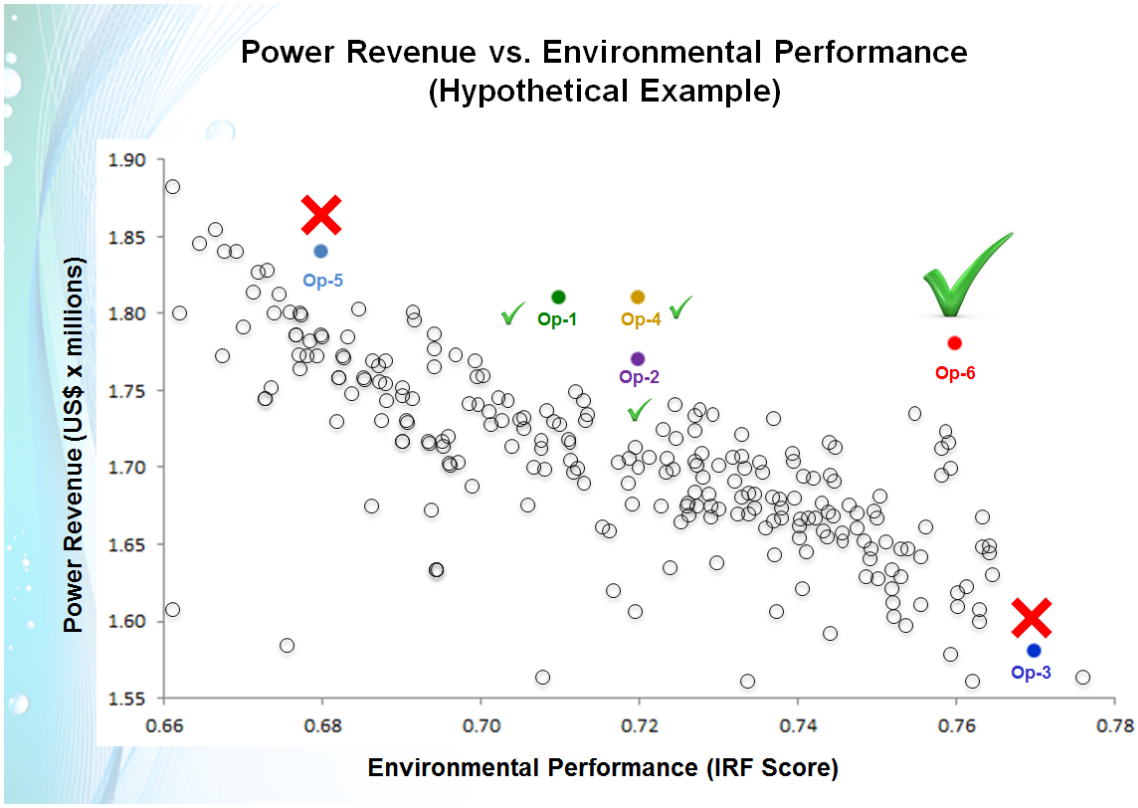

FIGURE 3-29 Hypothetical Example Showing the Results of Optimization after 1,000 Simulations (Results that were out of range of consideration are not shown here. Operation-6 [Op-6] provided the most optimal value and all other operations produced sub-optimal results. Operations 1, 2, and 4 [even though not optimal] could also be investigated further for potential operation, if desired.)

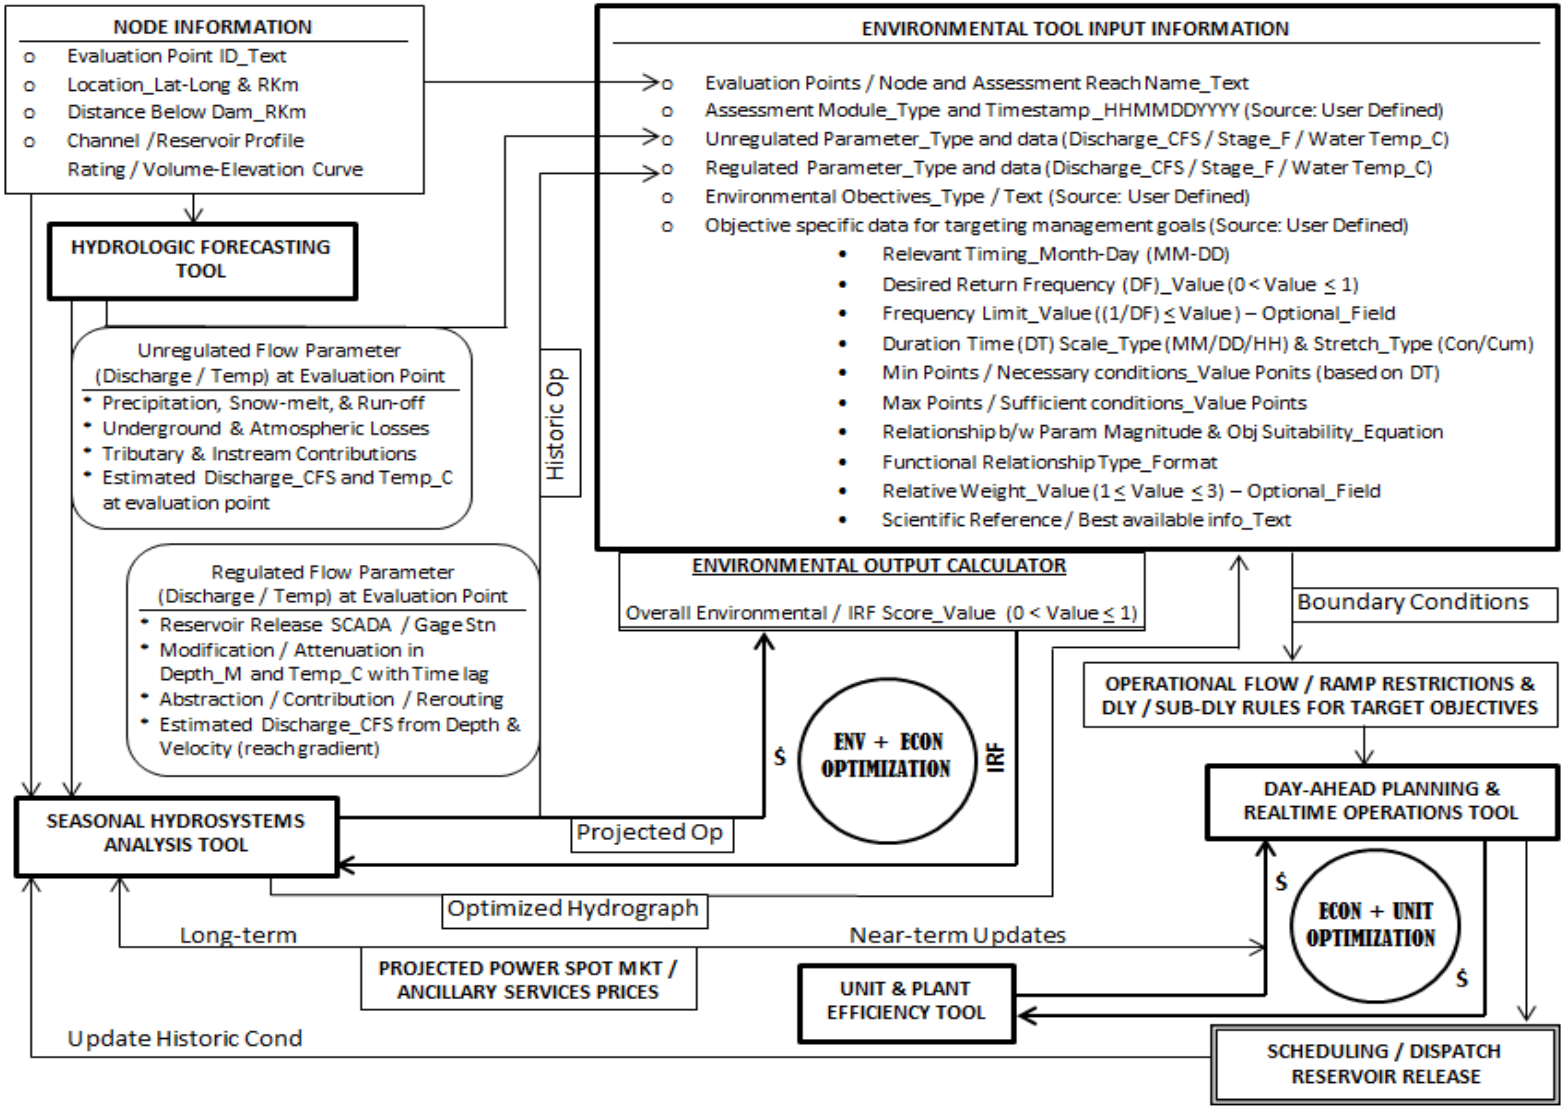

FIGURE 3-30 Layout Depicting Data Needs and Sources for IRF Tool and Their Exchange with the Other WUOT Components during Optimization

3.4.2.5 Calculating IRF Score Using IRF Historic Mode

The IRF tool can be run in a historic mode to calculate the environmental performance of past operations, whether for regulated or unregulated periods, or for calculating environmental performance of static operational scenarios. In this mode, the IRF tool runs in alone and is not integrated with other toolset components. It simply provides evaluation of environmental performance based on defined objectives; unlike the planning mode, in the historic mode there is no optimization option to improve future performance. The basic procedure of calculating the IRF historic score is similar to that for the IRF planning score, where a geometric mean is used to depict the overall performance of all identified environmental objectives. Similar to the IRF planning tool, the objective scores in the IRF historic tool are calculated for each term based on observed historical records and user-defined relationships between the metrics used to evaluate each objective and flow variable (such as discharge, stage, and water temperature). In the case of partially missing flow data, the points used to calculate objective scores are prorated based on observed data for the term. The scores for objectives that do not meet the minimum points required are readjusted to zero, and objective scores that exceed the maximum required level are capped at the maximum value. Next, based on the desired frequency defined for the objective, the numbers of years to be analyzed are calculated. For example, if the desired frequency for an objective is 0.25 and there are 30 years in the period of record being analyzed, then the 8 years with the highest scores are chosen (7.5 is rounded off to the next highest integer value) for analysis (the highest-scoring years may include some years with zero points). Points for these years are divided by the maximum points that could be obtained to calculate a normalized score. The arithmetic mean of the normalized scores for an objective provides the objective score. Finally, the geometric mean of all objective scores is calculated to obtain the IRF score. Objectives with higher relative weights proportionally contribute more to the overall IRF score.

3.4.3 IRF Tool Development

The IRF tool has been developed using the C# programming language and has been compiled into an executable application for the end user (Figure 3-31). The tool provides an interface that guides the user to enter appropriate input parameters and allows for downloading discharge and temperature data from text files or the Internet (e.g., USGS gauge station Web sites). A mechanism to evaluate the suitability of downloaded gauge station data has also been built into the IRF tool. After data entry is completed, the IRF tool stores entered information in a local MySQL database. The stored information is used, together with information provided by other tools in the WUOT, to calculate an overall IRF score based on the identified set of environmental objectives and flow-related information for desired downstream evaluation points. The resultant IRF score is an indicator of the environmental performance for one or more specified stretches of a regulated river. Reservoir operations can then be revised to optimize the overall IRF score. The scores output from the IRF tool are stored in the common database and its integration with the seasonal and day-ahead/real-time operational tools occurs through the same database. Modifications in the interface features are still ongoing to make it more operationally helpful and user friendly, to enhance its functional capabilities, and to streamline data exchange and computational efficiency.

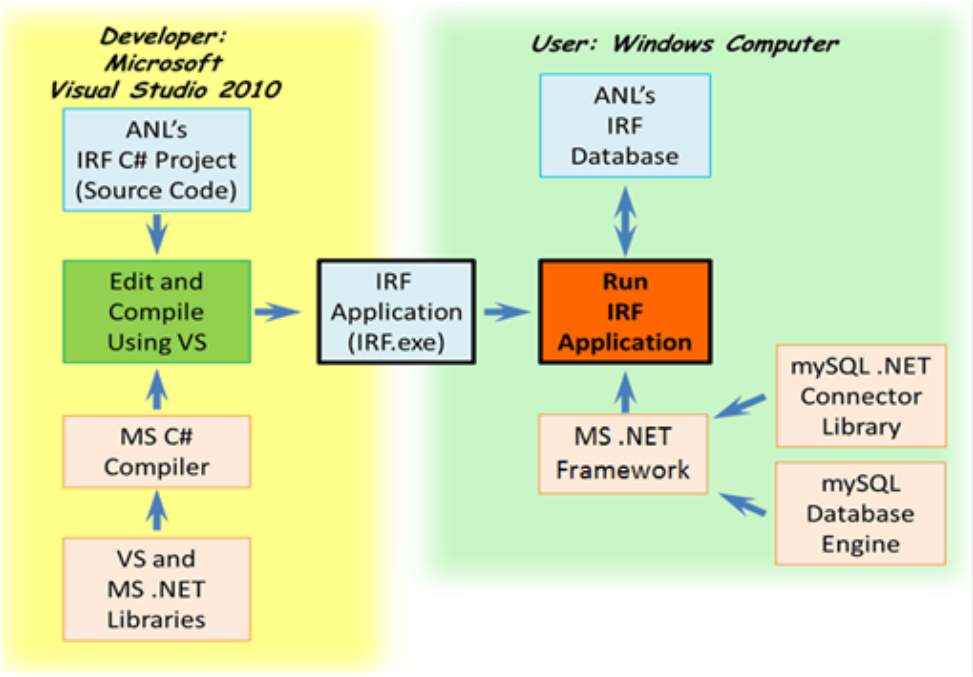

FIGURE 3-31 The IRF Tool Is Coded Using the C# Programming Language and Interacts with an SQL Database (The computer code is compiled into an executable program that runs on a Dot-Net platform.)

3.5 DAY-AHEAD AND REAL-TIME SCHEDULING

The modeling of hydropower plant day-ahead and real-time scheduling requires an accurate representation water and power system components, processes, and relationships. This section discusses the basic physics of hydropower plant operations and outlines key relationships between key model variables. It also describes some of the challenges involved in optimizing complex physical processes that occur in the cascade's reservoirs and power plants. Key functions describe relationships between (1) reservoir storage volume and elevation curves, (2) water flow rate and tail water elevation, and (3) power generation and head and power release rate. All these relationships are nonlinear and in some cases are discontinuous. Discussions in this section describe the basic principles that are used as the foundations for modeling hydropower plant operations and power production.

3.5.1 Overview of Physical Processes and Modeling Challenges

Hydroelectricity generation is based on the principle of water falling under the force of gravity to turn the blades of a turbine that is attached to a generator by a large shaft. As turbine blades turn, magnets in the generator rotate past copper coils to produce alternating current by moving electrons. The water that is contained in reservoirs represents stored energy. When a dam's control gates are opened, some of the water in the reservoir flows through penstock tubes to the power plant's turbines. The kinetic energy (energy of motion) of the flowing water is

converted to mechanical energy as it turns the turbine blades. The attached generator converts the mechanical energy into electricity. As shown in equation (3-41), the amount of electricity produced by this process is based on two main factors, which include the head (height of the falling water) and the water flow rate through the turbines. This equation is referred to as the *universal power equation*.

$$PQ_{dh}^{T} = \frac{\gamma \times R_{dh}^{T} \times TE_{dh}^{T} \times PF_{dh} \times H_{dh}}{CF^{kW} \times 1000} \qquad \forall \quad d,h,T \qquad (3\text{-}41)$$

where

PQ_{dh}^{T}	=	real power production from turbine *T* of dam ***d***, during hour ***h*** (MW);
γ	=	62.4, the specific weight of water at 50 degrees Fahrenheit (lbs/cubic foot);
R_{dh}^{T}	=	water release rate through turbine ***T*** of dam ***d*** during hour ***h*** (cfs);
TE_{dh}^{T}	=	efficiency for turbine ***T*** of dam ***d***, during hour ***h*** (fraction);
PF_{dh}	=	power factor for the power plant at dam ***d***, during hour ***h*** (fraction);
H_{dh}	=	head (ft) or forebay minus tailwater elevation at dam ***d***, during hour ***h***; and
CF^{kW}	=	737.5, horsepower to kilowatt conversion factor (kW/ft-lbs/sec).

Note: There are 1,000 kilowatts (kW) in a megawatt

TE_{dh}^{T} is also determined by water release rate and head.

The dam regulates water inflows from upstream sources and stores water in a reservoir. Total inflows into a reservoir consist of water released from upstream reservoirs ***RD***, and/or uncontrolled main stem river inflows and from side flows ***SF***. A side flow is a source of water that enters the stream in between two reservoirs or a water source that feeds directly into a reservoir. Water is released from the reservoir (referred to as outflows) through the penstock and past turbine blades ***RT*** to produce power or through bypass tubes and spillways in which no power is produced ***RW***. Both the timing and routing of water releases (i.e., power and non-power releases) are decision variables solved by an optimization model and are the most critical variables for maximizing the economic value hydropower resources and determining downstream flows that enhance environmental quality. Values for ***RT***, ***RW***, ***RD***, and ***SF*** are expressed in terms of total water release during time period ***t***.

The volume of water stored in a reservoir ***V*** at time ***t*** is based on a linear water balance equation. Variables included in the mass balance equation are the initial reservoir volume at time t_o (i.e., given state of the reservoir the instant before the first modeled period), reservoir inflows from time ***t***-1 to time ***t***, and reservoir outflows from time ***t***-1 to time ***t***. The timestep can be any length of period specified by users.

The elevation of the water surface in the reservoir ***E***, called the forebay, is a function of the amount of water that it stores. Unless the reservoir has the shape of a cube or rectangular

prism, the relationship between reservoir elevation and volume is nonlinear, as illustrated in Figure 3-32. The reservoir elevation level is typically constrained such that it remains within a given range.

Computations of the reservoir elevation are not only important for ensuring compliance with reservoir operating rules, but they are also key factors that determine the power output from the plant. The hydraulic head variable measured in terms of feet in the universal power equation (equation [3-41]) is the vertical distance in elevation between the reservoir water surface and the elevation of the water in the tailrace ***T*** at the base of the dam. The tailrace elevation is a non-linear function of the total water release rate from the dam as shown in Figure 3-33.

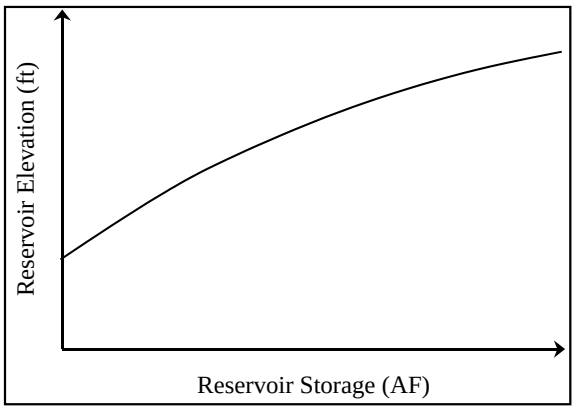

FIGURE 3-32 Reservoir Elevation Level as a Function of Water Storage

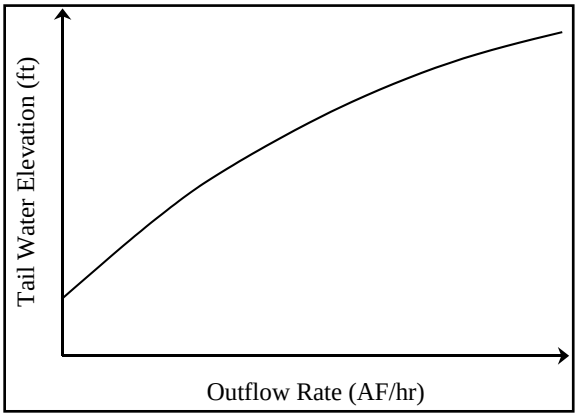

FIGURE 3-33 Tailrace Elevation as a Function of Outflow Rate

The hydraulic head may also be a function of a lower reservoir elevation in situations where the lower reservoir water is contained by the base of the upper-level dam.

The head limits the maximum turbine flow rate putting an upper bound on the maximum power output from the plant. As depicted in Figure 3-34, this non-linear function is usually concave. The upper bound often becomes a binding constraint on power production levels at lower reservoir elevations. At higher reservoir elevations turbine and generator constraints are typically the binding factors that limit power output.

The CHEERS model estimates the forebay elevation, head, tail, and upper bound on the turbine flow rate and generation capability. Holding all other variables constant in the universal power equation, output increases with higher turbine flow rates. In reality, however, all other variables in the universal power equation are not constant. For example, the tailrace elevation increases with higher turbine releases, thereby lowering the head and power output. Also, as more water is released, the reservoir's surface water elevation decreases, which also reduces both the head and power output. It is also interesting to note that for a given turbine flow, a non-power release will increase the tailrace and/or lower reservoir elevation level, thereby lowering generating unit's production level.

Another factor that complicates the modeling of a hydropower plant is that output from each individual turbine is based on an efficiency factor as expressed in the universal power equation by variable the turbine efficiency variable *EF*. It accounts for the friction of water flowing through trash racks, valves, and penstocks, as well as mechanical friction and energy conversion losses. As shown in Figure 3-35, turbine efficiency varies as a function of both turbine output measured in horsepower (HP) and head. As the turbine output increases from a minimum level to a higher one, the turbine efficiency increases. Typically, the efficiency then begins to decline to some point that is less than the turbine output capability. As shown in Figure 3-35, this efficiency curve tends to shift toward the left (i.e., lower horsepower output) with lower heads. Also, the maximum efficiency point on the curve may vary as a function of head.

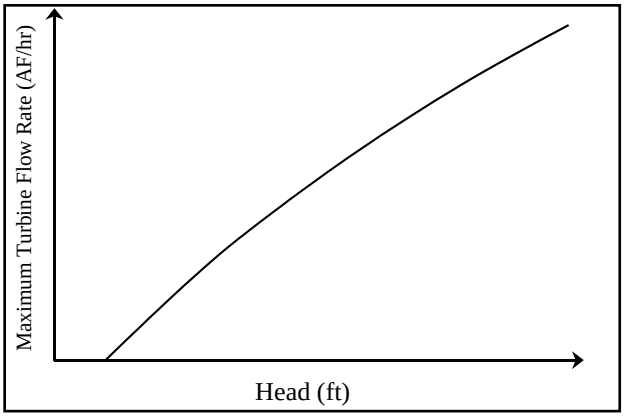

FIGURE 3-34 Maximum Turbine Flow Rate as a Function of Head

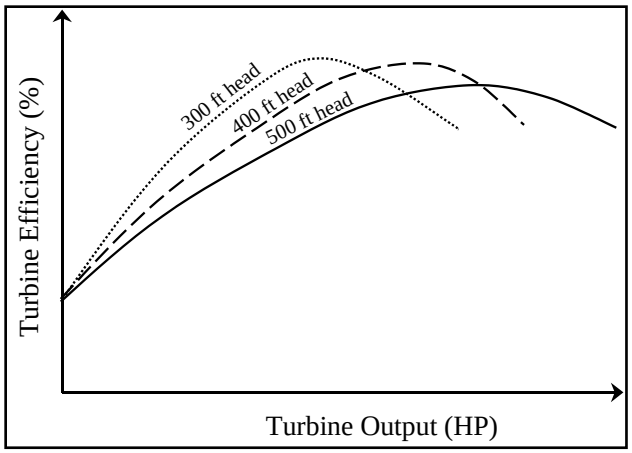

FIGURE 3-35 Turbine Efficiency as a Function of Both Turbine Output and Head

Modeling challenges are compounded by factors that create discontinuities in power production curves. These include minimum turbine generation levels and turbine rough zones.

When there are multiple turbines at a power plant, optimal plant operation must take into account not only the efficiency curve for each individual turbine, but also the additive effects of water releases from all turbines on the head. In some situations, it is sub-optimal to release the same amount of water through each turbine. Therefore, for a given scheduled generation level, an optimization model must determine the allocation of flows through each turbine to minimize total reservoir water releases. This allocation process maximizes the efficiency of the power plant in terms of water release per unit of electricity generated.

The amount and value of the real power that is sold to the market is a function of all the physical properties of hydropower plants discussed above and hydrological characteristics of the basin that feeds water into reservoirs. It is also a function of electricity market conditions. From an economic viewpoint, the most important variables that influence the value of hydropower plants are the demand for electricity, the cost to produce power from supply-side power resources and deliver it to load centers, and the capacities of the transmission grid and distribution systems.

Under most hydrological conditions, water resources are scarce in the sense that there is an insufficient quantity of water to run all unit generators at a plant at all times. Given the physical, environmental, and institutional limits of a system, CHEERS uses these limited water resources to maximize the value of hydropower resources. At times, this power objective conflicts with environmental goals.

3.5.2 Concept

The CHEERS model creates day-ahead and real-time schedules. It provides solutions that suggest when, where, and how much water to release from reservoirs, power to generate from individual units, and capacity to reserve for ancillary services to fulfill reliability criteria and for sales to the market. Modeled ancillary services include regulation up and down, spinning reserves, and non-spinning reserves. Schedules are driven by multiple objectives that consider energy and environmental goals simultaneously.

Whereas the other tools in the WUOT might be concerned with only certain types of network objects, the CHEERS model will always use in its optimization every node and link defined in the GUI by the user. Furthermore, other tools by their nature may deal with only the water commodity, but CHEERS has the greatest level of flexibility to work with user-defined commodities and attributes. Over simulated time, CHEERS calculates attribute values as the commodities flow over links and enter, exit, and are stored within nodes.

The CHEERS modeling framework is designed to apply to a wide range of systems operating under a diverse set of hydrological conditions. It can be applied to back-casting studies, used in actual run-time settings as part of a utility's daily business practice and, when combined with other tools, used to explore longer-term futures. For all of these applications, the user can specify one or more cases describing alternative operating criteria and/or proposed changes to system components.

3.5.3 Components and Modeling Processes

The CHEERS tool incorporates several components that are typically applied in succession with information flows between each routine. For demonstration projects, we used these CHEERS routines and algorithms:

1. Data Validation and Screening routine,
2. Curve Fitting algorithm,
3. Interpolation routine
4. Piecewise Linear Optimizer,
5. Forecast Error Analyzer,
6. Mathematical Formulation routine,
7. Optimizer,
8. Non-linear Post-processor, and
9. Results Report Writer routine.

In essence, the first five routines assist the user in processing raw data so they can be transformed into functional forms for use by other CHEERS components. Each routine is stored in a separate Excel workbook that communicates with the WUOT database via the GUI. In the current modeling process development stage, we performed data transfers manually.

The Mathematical Formulation routine is written in C#. It communicates directly with the database to translate user inputs into a problem statement that can be understood by the Optimization Routine that solves it. The last two routines post-process Optimizer outputs and summarize model results in the form of tables and graphs. The nonlinear post-processor accesses optimized values and coefficients for model functions from the database. This routine, also written in C#, resolves differences between (1) functions used in CHEERS that "approximate" a process and (2) nonlinear equations and lookup table that more accurately describe a process or function.

The Results Report Writer is also an Excel workbook. It uses optimization results to compute economic and financial values for the system. This routine also summarizes results in the form of graphs and tables.

The Formulator and Optimizer components are the only essential model components. Other CHEERS components are not used if the functions performed by these components are not required for a model application or can be obtained by an alternative means. All routines are nonetheless available to WUOT users.

Section 3.5.4 describes the CHEERS design and how the routines described above were applied for demonstration site applications. Although there are similarities among WUOT models, other tools use the WUOT topology and data input differently. Therefore, the discussions below are specific to the CHEERS model. The discussions are also hydropower-centric and focus on areas that are key to demonstration site applications.

3.5.4 Design

The basic design of CHEERS is very versatile and applicable to many different types of network and material/energy flow problems that extend beyond hydropower and power system applications. This versatility makes it unique among other hydropower models. In the context of the WUOT toolkit, the objectives of the CHEERS are more narrowly defined for the current demonstration site applications. However, the representation can be expanded to include other components and aspects that go beyond hydropower. For example, the model can represent thermal powerplants, battery storage, and fuel supply chains with energy extraction, refining, processing, transport, distribution, and useful energy demand.

At its core, CHEERS consists of a routine that translates WUOT input data and information into a set of mathematical equations (i.e., Mathematical Formulation routine) that are solved by an optimization algorithm (i.e., Optimizer routine). CHEERS equations are formulated as a network problem that describes the operation of one or more interdependent systems. It also includes equations that represent physical processes and time-sensitive scheduling problems.

The Optimization routine determines values for "decision variables" that maximize a given function, called the "objective-function," while satisfying given inequalities of other

functions called "constraint-functions." In its most general form this mathematical problem is conveniently written as:

$$\begin{aligned}&\max_x f(x)\\&\text{subject to: } g_i(x) \le 0, i = 1, \dots, m\\&x_i \in \mathbb{Z}, i \in I,\end{aligned}$$

where m, I, f, g are given. In general, the sense of the inequality of any constraint may be reversed. Also, a constraint may be expressed as an equality formula.

Core CHEERS routines require that equations involving both f and $g_i, i = 1, \dots, m$ be linear functions. In addition these two routines designate some variables as "integers"; that is, solutions for these variables must be discontinuous integer number values. Solutions for other variables, referred to as "real variables" in the mathematical problem, can be continuous real number values. These two requirements place CHEERS core functions into a classification that is commonly referred to as a "mixed-integer linear program (MILP)" problem.

MILPs are non-deterministic polynomial-time (NP)-hard problems (Papadimitriou 1981); in the worst case, the time required to solve MILPs by any known algorithm grows exponentially with the number of integer variables. However, the behavior is usually much better in practice, particularly when the branch-and-cut algorithm is used to solve MILPs. Several software products, both open-source (Forrest 2006; Achterberg 2009; Ralphs et al. 2011) and proprietary (IBM undated; GUROBI Optimization undated; FICO 2012), are available that implement this algorithm. These solvers have been shown to solve real-life problems with tens of thousands of variables and constraints.

On the other hand, if either f or any $g_i, i = 1, \dots, m$, is not linear, the problem is called a mixed-integer nonlinear program (MINLP). In general, MINLPs are much more difficult to solve than MILPs, and the solvers available are much slower and less robust. Many difficulties arise if the nonlinear functions are not convex or have discontinuities or poor numerical stability. While there are a growing number of solvers available to solve MINLPs (Sahinidis 1996; Mahajan et al. 2012; Belotti 2009; Bonami et al. 2008; Abhishek et al. 2010), they often have difficulties solving problems with more than 100 variables and constraints.

Water, hydropower, and environmental systems are inherently nonlinear and discontinuous, seriously challenging the limitations of leading optimization software packages and the computational capabilities of our most advanced computer systems. For example, the elevation of water in a reservoir changes nonlinearly as a function of water storage volume. Such changes, in turn, affect the amount of electricity produced by a hydropower turbine, which is a nonlinear function of both the flow rate and the hydraulic head. In addition, hydropower turbines have minimum technical production levels and rough operating zones that operators attempt to avoid to prevent equipment degradation.

While an MINLP can most accurately represent the problems modeled by CHEERS, the current state of MINLP technology limits its use to tiny models only. CHEERS problems

typically require thousands of variables and constraints. Therefore, the two core CHEERS routines are restricted to using only linear equations with both real and integer decision variables. Nonlinear functions are approximated by piecewise linear functions combined with integer variables to provide solutions. A non-linear post-processor routine applies the solution obtained from the MILP to solve the full set of nonlinear equations, correcting for differences between the piecewise linear approximations and the original non-linear model.

Although the CHEERS optimization model possesses several aspects of a network flow optimization problem (Ahuja et al. 1998), there are some fundamental differences from a pure network problem. Unlike traditional network flow models, CHEERS contains more complex features. Properties of a commodity may be altered as it flows through a network and/or over time and there are numerous temporal constraints, which are more common in sequencing and scheduling problems. Considering the complex structure of the model, we use a general-purpose MILP solver to find the solution.

For demonstration site applications the CHEERS Optimizer routine uses the LINGO™ (LINDO Systems Inc.) optimization solver to solve network/scheduling MILP problems. Other commercial state-of-the-art mixed integer programming solvers include solvers such as CPLEX version 12 (IBM) and GUROBI version 5 (Gurobi Optimization). Alternative open-source solvers such as COIN-OR CBC have been at least moderately competitive with the state-of-the-art commercial solvers in the past few years. COIN-OR solvers are freely available under the Eclipse Public License (EPL). GLPK is another freely available optimization solver under the GNU Public License (GPL).

Standard optimization solver interfaces will be used to enable the seamless use of a variety of optimization solvers. Optimization solver interfaces are either based on ASCII files or programming APIs. Existing file/text based interfaces can be categorized as algebraic modeling languages or simple/”lightweight” ASCII interfaces. Algebraic modeling languages include the Abstract Modeling Language for Mathematical Programming (AMPL) and General Algebraic Modeling System (GAMS). Simple ASCII interfaces include the MPS and LP file formats; these require the models to be written in expanded form and are typically involve significantly larger text files. Finally, standard programming APIs include the C++ based COIN-OR Open-Solver Interface (OSI).

The CHEERS Mathematical Formulation routine currently creates problem statements and equations in both the LINGO format and in the AMPL (a mathematical programming language) format. The AMPL modeling language allows model developers to write optimization problems in a format similar to the mathematical notation needed to express them. Several solvers—including CBC, SCIP, CPLEX, and GUROBI—are capable of solving AMPL models. This approach thus provides flexibility in selecting a solver and avoiding solver “lock-in.” CPLEX and GUROBI may be required for larger and more complex problems than those currently posed by the demonstrations.

3.5.4.1 Objective Function for Demonstration Sites

The CHEERS model, by design, has wide-ranging applicability. The description below focuses on the constrained water and hydropower optimization problems posed at the current demonstration sites. Objectives are based on maximizing the sum of one or more commodity attributes multiplied by objective function scalars. For example, one objective may be to maximize revenues generated from the sale of a commodity such as electricity. In the model's objective function, the solved value for the electricity quantity decision variable is multiplied by a given scalar that represents the price of electricity. Additional terms can be added to the objective function to include the economics of regulation services and spinning reserves. Another potential future objective scores habitat suitability for an endangered fish species by controlling water attributes such as temperature and dissolved oxygen levels.

Equation (3-42) shows the demonstration site objective function used to maximize the economic value of the hydropower resource. This value includes power grid benefits associated with electricity generation from water releases through a water turbine T at plant p over all time steps t during a modeled timeline. A vector of electricity market prices P is used as a surrogate for its economic value. In addition to the economic benefits of energy, the objective also accounts for revenue streams from the sale of ancillary services. This includes both up- and down-regulation services, along with spinning and non-spinning reserves. On the cost-side, the objective function includes both direct and indirect expenses associated with unit startup and shutdown.

$$\text{maximize } Z = \sum_{p}\sum_{t}\sum_{T}(P^t G_T^t + S_{uT}^t RS_u^t + S_{lT}^t RS_{lT}^t + A_{sT}^t D_{sT}^t + A_{nT}^t D_{nT}^t - W_{uT}^t C_{uT}^t - WC_{lT}^t) \quad (3\text{-}42)$$

where

Z	= net economic value of system operations over the optimization period;
G_T^t	= power generated at turbine T at time t;
P^t	= market price per unit of energy at time t;
S_{uT}^t and S_{lT}^t	= market value of regulation up and down services, respectively;
RS_u^t and RS_{lT}^t	= level of regulation services;
A_{sT}^t and A_{nT}^t	= market value of spinning and non-spinning reserves, respectively;
D_{sT}^t and D_{nT}^t	= level of reserves provided;
W_{uT}^t and W_{lT}^t	= turbine startup and shutdown cost, respectively; and
C_{lT}^t and C_{uT}^t	= whether a unit is turned on or off, respectively.

When multiple objectives are specified, the user must place weights on each one. A commodity attribute has either positive or negative value. CHEERS input forms allow the user to specify objective function components and scalars.

The maximization objective is subject to a set of constraints that restrict the operation of power plants and reservoirs in a system by the following means:

1. A "balance" is maintained at every node in the network;

2. Total water release volume from each storage node over the optimized time period must be within minimum and maximum levels;

3. Reservoir elevation (and thus, operational volumes) is limited to a specified range and is restricted by a maximum change over specific time spans (e.g., single and/or multiple days); and

4. Changes in water flows (hourly for demonstration site applications) through hydropower turbines and from reservoirs must be limited to a specified range.

3.5.5 System Networks and Building Blocks

A WUOT network comprises two or more node objects and one or more link objects. For each object, CHEERS solves for the value of one or more attributes associated with one or more commodities. These attributes serve as decision variables in the CHEERS formulation and are solved for at either the time step end state or for the average state over a time step interval. As described in previous sections, WUOT nodes types include (1) boundary, (2) storage, (3) conversion, and (4) junction. Links represent connections between nodes and facilitate commodity flows throughout the network.

In the CHEERS context, model object types are generic, enabling an object to represent a wide range of physical objects and processes. For example, there is no model object type that represents a hydropower plant. Instead, one or more model objects along with input data are used to define the properties and functionality of a specific hydropower plant. The simulated operation of the plant and its parts are dictated by the network configuration and the rules that govern its processes.

Theoretically, the number of link and node objects in a CHEERS network is unlimited. However, there is a practical limit to the network size because the complexity of the system and model run times increase as a function of the network size, the number of time steps, and the form of attribute and commodity functions and relationships. A brief description of each WUOT object type and how it is applied by CHEERS to demonstration site problems follows.

3.5.5.1 Boundary Node

A boundary node represents a point on the edge of a network in which a single commodity with associated attributes either enters (source) or exits (sink) the network. The commodity enters or exits the boundary node over time via one or more links that connect it to other network nodes.

For demonstration site applications, water both enters and exits a network at boundary nodes. The timing of water flows can be obtained or derived from one of many sources. These include, but are not limited to, PNNL's Hydrological Forecasting tool, persistence forecasting, or USGS gauge data and predictions. Water exits the system at boundary nodes that represent evaporation, withdrawals for irrigation and other water demands, and main stem flows that represent the lowest point of interest in a hydrological system.

Boundary nodes also represent electricity sources and sinks. Sources include purchases of electricity from external systems via long-term, short-term, and spot markets. Sinks represent internal system energy consumption and the delivery of electricity to points where it is sold to entities outside of the modeled system.

In situations where commodity levels entering or exiting a system are not given, amounts are provided by CHEERS for each simulated time step based on model objectives, constraints, and the system topology.

3.5.5.2 Storage Node

As the name implies, a storage node accumulates a commodity in a containment facility and releases it over time. A commodity enters the storage node via flows on one or more connected input links. Likewise, the commodity exits the storage node through one or more output links. The storage node contains only one homogeneous commodity. Therefore, all links attached to the storage, both incoming and outgoing, transport the same commodity, which is the commodity contained at the node.

A mass balance equation computes the storage level at the end of each simulated time step. It is based on the following:

- The quantity that is stored at the end of the previous time step plus;
- The sum of all quantities entering the storage on incoming links during the current time step *minus*; and
- The sum of all quantities exiting the node on all outgoing links during the current time step.

Since the CHEERS model does not solve for the state of the storage node prior to the first optimization period, initial state information is exogenously input into the tool. This is typically historical/supervisory control and data acquisition (SCADA) information stored in a demonstration site database.

For demonstration site applications, the storage node is used to represent a reservoir that stores water behind a dam structure. Water sources include uncontrolled river channels, releases from upstream reservoirs, and side flows that either pour directly into a reservoir or into a water

channel that connects two cascaded reservoirs. Water exits a storage facility through penstocks leading to turbines and non-power channels such as bypass tubes and spillways. Water also exits a reservoir through evaporation and is diverted for crop irrigation and to serve municipal demands.

A storage node may contain one or more stratification layers (i.e., storage levels or zones) as defined by the user. For example, stratification in a reservoir may include "dead storage," "active," and "flood control." Although the commodity stored in a reservoir is assumed to be homogeneous, attribute values may differ by stratification zone as computed by a user-defined function.

Storage nodes may have one or more attributes. For example, demonstration site reservoirs have storage volume and forebay elevation attributes. The forebay elevation is computed by a function relating water storage volume to elevation. As shown in Figure 3-36, this function is normally nonlinear and varies according to the geometry of the reservoir. Typically, it will increase relatively quickly as a function of water storage volume when the reservoir is empty (forebay elevation is at a low point) but slowly increase in elevation when storage is nearly full (forebay elevation near the top of the crest elevation).

Since CHEERS is formulated as an MILP problem, all functions must be linear. Therefore, the non-linear equation is approximated with a piecewise linear function. In order to obtain a piecewise linear approximation, two or more breakpoints are selected where different sloped pieces meet. The user specifies the number of segments used, and may specify breakpoints based on boundaries between reservoir layers or some other criteria. Alternatively, the user can utilize the Piecewise Linear Optimizer, which determines the slopes and intercept of line segments

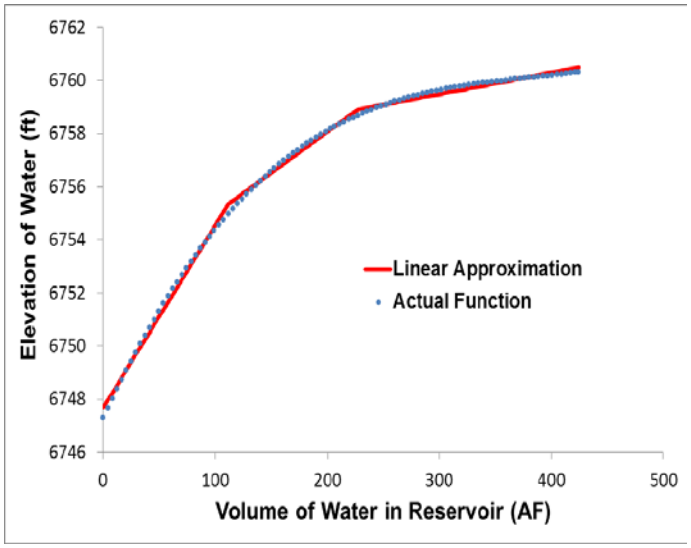

FIGURE 3-36 Relationship between Reservoir Volume and Elevation

yielding the minimum sum-of-square error between the non-linear function and the piecewise linear approximation. In general, the sum-of-square error for piece-j is given by equation (3-43):

$$e_j = \sum_{i=1}^{k} \left(a_j x_{ij} - b_j - \hat{y}_{ij}\right)^2, \tag{3-43}$$

where x_{ij} is the volume associated with the point i lying in piece j, and $y_{ij}$is the observed height. If the breakpoints are fixed, then a_j, b_j are obtained by solving a convex unconstrained problem. For demonstration site applications, the breakpoints are obtained by the routine employing an evolutional algorithm that intelligently iterates this procedure for numerous sets of breakpoint locations and reports on the set that yields the smallest error, e_j. A more advanced optimization routine developed in this project finds better solutions faster than the evolutionary algorithm, and will add it to CHEERS as an alternative option to the evolutionary algorithm in the next release. Details of this new algorithm and illustrative applications are provided in Appendix F.

At many locations, only a small portion of a reservoir's volume-elevation curve is needed for day-ahead and real-time optimization. Therefore, only that portion of the curve that applies for a simulation period *x* is used to construct the piecewise linear function. This approach yields a significant reduction in the error between the piecewise relationship and the non-linear function. Typically the "top layer" of a reservoir can accurately be described with very few line segments, thereby increasing computational efficiency.

For applications with very large storage volumes and a nearly constant elevation change function over the daily storage range, it is sufficient to describe elevation/volume changes using a single line segment. The slope of this line is the first derivative of the volume-elevation function that is expressed as a polynomial equation. The derivative is computed for the point on the curve that represents the average anticipated next-day reservoir storage volume. Alternatively, the slope may be based on two points of the elevation/volume that represents the operational storage range of over the simulated period.

The WUOT GUI enables users to input CHEERS simple linear relationships among commodity attributes (i.e., decision variables) with user input forms. A simple relationship sets a variable (left-hand side of a function) equal to a constant plus the sum of one or more variables multiplied by coefficients (right-hand side of the equation). In the situation described above, the reservoir elevation in terms of feet (ft) is set equal to a constant (in ft) plus the reservoir water storage volume in terms of acre-feet (AF) times a coefficient (in ft/AF). This coefficient relates the change in reservoir elevation per AF of water stored.

3.5.5.3 Conversion Node

A conversion node uses one or more incoming commodities and converts these inputs into one or more outgoing commodities. To define a conversion process, the user constructs functions in a CHEERS input form that relates attributes of one or more nodal input commodities (e.g., quantity of water) to the attributes of one or more output commodities (e.g., quantity of electricity). Input forms enable the user to express a conversion function as either a simple linear

equation or as a piecewise linear function. Piecewise linear functions are expressed as a vector of *x* and *y* points that identify linear segment intersections points.

For demonstration site applications, a conversion node representing a power turbine/generator utilizes water as an incoming commodity producing both water and electricity as outgoing commodities with associated quantity attributes. Turbine electricity production is represented as a function of both the rate of water flow through turbines and head. Since all CHEERS functions must be linear, a piecewise linear function approximates non-linear relationships. Piecewise linear function break points are determined by the Piecewise Linear Optimizer.

Specially-Ordered-Set of type 2 (SOS2) constraints are used to represent the piecewise linear function in CHEERS. SOS2 constraints are compact and computationally efficient representations of piecewise linear functions. Such representations ensure that the MILP solver recognizes the piecewise-linear structure, which may otherwise be lost. This representation also enables the solver to use techniques specially designed for these structures (Forrest and Beale 1976).

The SOS2 representation describes a two-dimensional surface. Because the power function is bivariate involving both head and water flow rate, additional variables are required to represent this three dimensional surface. An SOS representation having higher dimensions (Tomlin 1981) has been proposed to solve a function of this type, but it is not recognized by most solvers and modeling languages. Model users would need to implement customized routines specific to the optimization solver in order to use SOS extensions, an approach that is impractical for a general framework like CHEERS. The CHEERS model instead uses an alternative method to create a linear approximation of the function. This approach keeps the framework independent of the optimization solver used, while preserving the accuracy of the formulation.

In the first step, the power-turbine flow function reference is determined for a fixed head. This fixed head level is typically based on the anticipated lowest head during the modeled timeline. The Piecewise Linear Optimizer is then used to obtain a piecewise-linear representation of the power-turbine flow function. Expressed in terms of a set of SOS2 constraints, this function is illustrated by the dashed red line in Figure 3-37. Keeping line segment break points at the same turbine flow rates, the power-turbine piecewise function is also determined for a higher head. Based on the computed head, the CHEERS model uses a linear interpolation to determine unit output levels from the two curves.

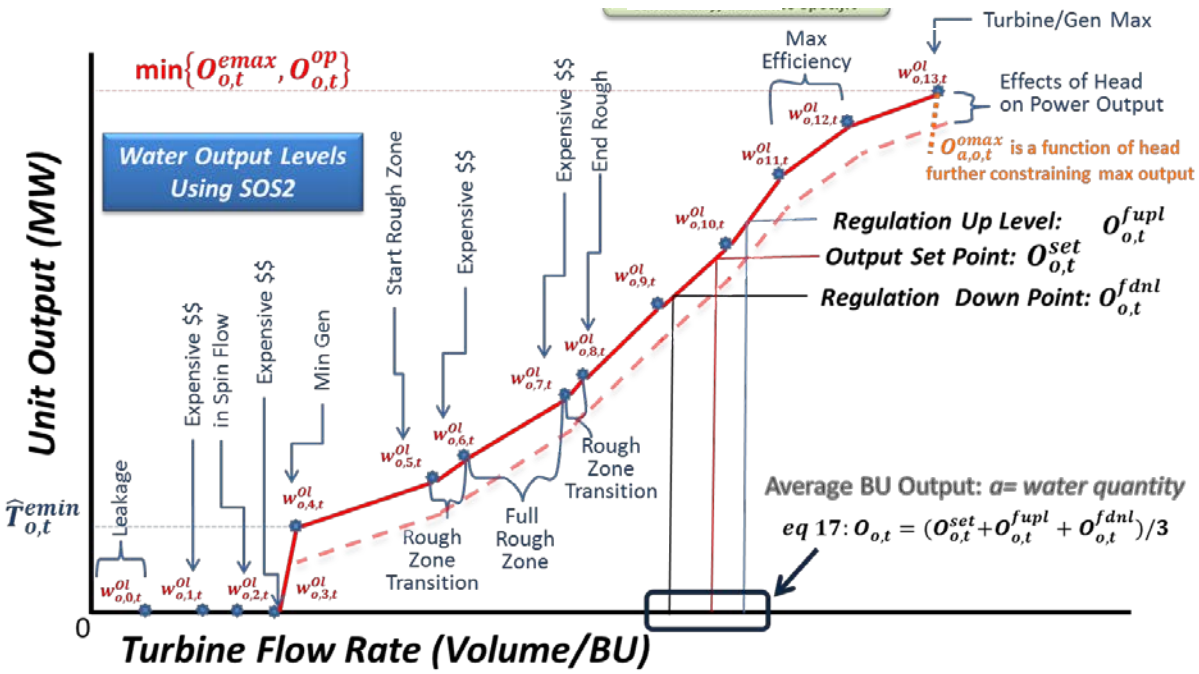

FIGURE 3-37 Hypothetical SOS2 Curve Relating Generating Unit Power Output to Turbine Flow Rate

Turbine level generation is computed with equation (3-44). It multiplies the water flow rate RT_p^T through turbine *T* at point *p* by the weight w_p^T for turbine *T* at point *p*, where *np* is the number of break points in the set a SOS2 *xy* points define a piecewise linear curve:

$$G_T^t = \sum_{p=1}^{np} \left(w_p^T \text{x}\, RT_p^T \right),\ p = 1, \dots, np\ . \tag{3-44}$$

SOS2 weights sum to 1.0 (equation [3-45]) and all weights are real numbers between one and zero (equation [3-46]):

$$G_T^t = \sum_{p=1}^{np}\ w_p^T,\ p = 1, \dots, np \tag{3-45}$$

$$0 \geq\ w_p^T \geq 1.0\ \forall\ p \text{ and } T. \tag{3-46}$$

Furthermore, no more than two adjacent members in the set may be non-zero.

Weights are multiplied by scalars in the objective function to prevent units from generating below a minimum level and to discourage the operation of turbines in rough zones. Weights can also be multiplied by other variables in the model. For example, a weight can be used to define the maximum ancillary service level associated with operating points.

Certain constraints require discrete choices. For instance, there is a cost associated with starting a turbine, but the cost is incurred only if the turbine was off during the previous time

period and on during the current period. This cost and other such discrete phenomena are modeled using binary variables. A binary variable is one that is required to be either 1 or 0 in a feasible solution. To model the turbine startup cost, we use the binary variable S_T^t, which is assigned 1 if the turbine T is starting up in time period t. The startup cost is represented by equation (3-47):

$$\pi = \sum_t \sum_T \pi_T N_T^t, \tag{3-47}$$

where π_T is the cost of starting up turbine T, and N_T^t indicates the status change from off to on. We use binary variable O_T^t to denote if the turbine T is on at time t. Then, variables N_T^t and O_T^t are related by the constraints shown in equations (3-48) and (3-49):

$$O_T^t \geq N_T^t \tag{3-48}$$

$$N_T^t \geq O_T^t - O_T^{t-1} \tag{3-49}$$

The first constraint ensures that O_T^t is 1 whenever N_T^t is 1. The second constraint ensures that N_T^t is 1 whenever O_T^t is 1 and O_T^{t-1} is 0. More binary variables are similarly used to model other constraints, like those related to rough zones of the turbine frequency that must be avoided when it is generating electricity.

Conversion nodes may have limits placed on attributes associated with either the conversion node itself or other nodes or links anywhere in the network. For example, the maximum release of water through a turbine conversion node during a time interval may be a function of the forebay elevation of its reservoir (i.e., storage node attribute) and water density (i.e., commodity attribute).

3.5.5.4 Links

The link object connects two nodes and simulates the flow of a commodity. Links are unidirectional and transport one and only one commodity. There may be multiple links flowing out of a node, and multiple links flowing into a node. Two links may be connected to the same pair of nodes to represent the flow of the same commodity in different directions. Demonstration site applications use links to transport water and electricity throughout the network.

By default, all of a commodity that exits a source node arrives at the destination node in the same time step. By applying a Travel Time Distribution (TTD) function as shown in Figure 3-38 to a link, a commodity may take more than one time step to arrive at a destination node from a source node. Furthermore, all of the quantity released from an origin node may not arrive at the destination node at the same time; a portion of a commodity may travel faster or slower than other portions. To simulate this situation, CHEERS discretizes the released amount into smaller quantities (fractions of the total), whereby each fractional amount takes a different amount of time to travel from beginning of the link to the link's ending point.

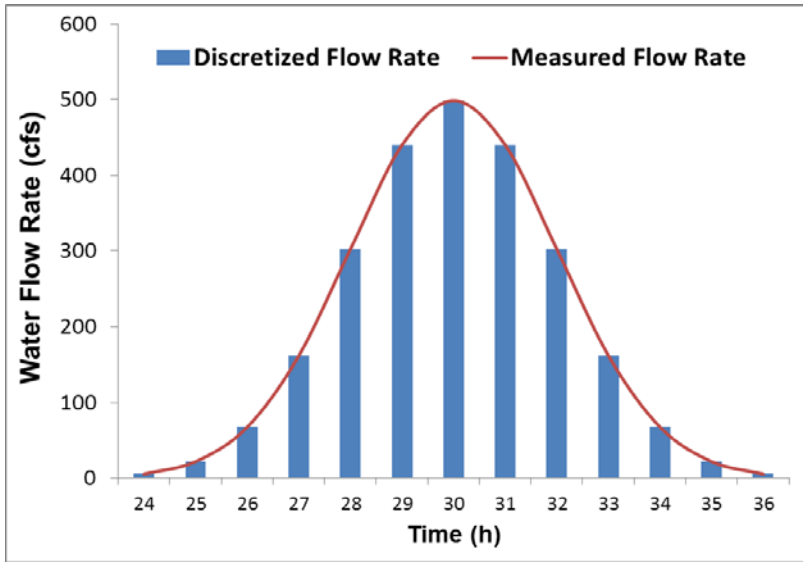

FIGURE 3-38 Hypothetical Water Travel Time Distribution Function

For hydrological applications, applying a TTD simulates a wave of water that attenuates as it travels down the reach of a river. If the geographic distances between nodes in a network are significant, then some portion of the water entering a river reach will arrive at a downstream point more quickly or more slowly than the rest of the water. In general, the quantity of water reaching the downstream point follows a continuous distribution as shown by the hypothetical smooth curve in Figure 3-38. Discretizing this curve per base time-unit (hourly, in this case) to conserve the quantity of water causes the discrete function to closely resemble the measured flow rate. The shape of this distribution usually resembles the shape of a normal distribution, but it may vary according to the characteristics of the river channel, travel distance, and the volume of flow.

Mathematically, the quantity of commodity flowing from node i measured at node j at time t is given by equation (3-50):

$$Q_{ij}^{t} = \sum_{\tau=1}^{K_{ij}} Q_i^{t-\tau} f_{ij}^{\tau}, \tag{3-50}$$

where the quantity Q_i^t of water is released from node i at time t andf_{ij}^{τ} is a parameter obtained from discretizing the time-travel distribution function, and K_{ij} is the number of time periods within which the discharge from node i affects distribution at node j.

Where applicable, a loss factor is used to account for commodity loss that may occur between the entry and exit points on a link. For example, some water in a river channel may evaporate as it travels downstream.

The example shown in Figure 3-39 illustrates this TTD approach. A relatively large amount of water is released from a hydropower plant during the third simulation hour. All of the

water eventually reaches the downstream gauge, but the peak flow attenuates and arrives at some downstream location over a three-hour span during hours 9, 10, and 11. Typically, the sum of hourly TTD fractions equals 1.0; however, if losses or gains occur en route, the fractions may sum to a value less than 1.0. In the example shown in Figure 3-39, there may be water evaporation from the river channel (sum <1.0) or a water gains from groundwater sources (sum >1.0).

When the user applies a TTD function it implies that the link not only transports a commodity, but also contains it. In this respect, the link functions as a storage object. The quantity retained in a link is estimated by a mass balance equation. The TTD functions are not limited to water flowing through a channel; they may also be applied to other commodities (oil traveling through a pipeline, for example).

3.5.5.5 Junction Node

The intent of this node type is to spatially aggregate and disaggregate a commodity. It is also used to measure, modify, and constrain a commodity attribute at specific geographical point. One or more links enter a junction node and one or more links exit it. No commodity conversions or storage occurs; therefore, the total amount of a commodity that enters it instantaneously exits it in an unaltered state.

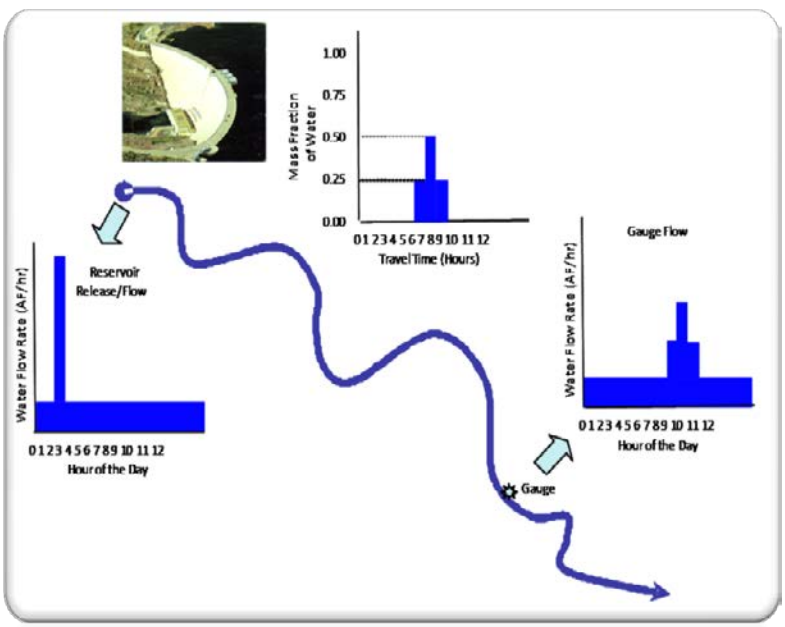

FIGURE 3-39 Illustration of a Wave Function Modeled as TTD

For demonstration site applications, junction nodes are used for several different purposes. For example, tailrace points are used to estimate water elevations below a dam based on one or more inflows from power and non-power releases. Junction nodes are also placed at key points in a basin where environmental restrictions are specified and IRF scores are computed by the Environmental tool.

3.5.6 Simulation Timeline

The CHEERS model determines decision variable values for all chronological time steps during the model timeline. This time chronology feature of CHEERS allows it to simulate a host of operational aspects, such as generator ramp-rate limitations, unit commitment schedules, unit starts and stops, and multi-time-period constraints, including maximum and minimum water releases from a reservoir over the simulated period. Simulated time is also critical for solving network interdependency problems, such as cascaded reservoirs, because CHEERS recognizes reservoir connectivity and the time it takes water to flow between two reservoirs. When optimizing, the model recognizes that the solved state of the system at any point in time affects operations at all other points in the optimized time period.

To define a CHEERS timeline, a user first specifies a starting point in time and the time interval (i.e., time step) at which network decision variables (also known as commodity attribute states) are solved. The simulation time interval is referred to the time Base Unit (BU). The user enters the both the timeline start and BU information into WUOT using the CHEERS "Run Settings" input form. The Run Settings menu also allows the user to specify sub-periods in the CHEERS timeline. Listed in the required sequential order, these sub-periods are classified as historical, real-time (i.e., now), scheduled, projected, and extension. One or more sub-periods can be omitted from a model run by assigning it a zero time span; however, there must be at least one sub-period with a non-zero time period and all sub-periods must be contiguous in time.

After the user disaggregates the run timeline into sub-periods, each sub-period is designated by the user in the Run Setting input form as being either *specified* or *optimized*. During a specified sub-period, the model is constrained to produce values for selected (not all) key decision variables that equal those entered by the user. Also, values assigned to these variables do not contribute toward the objection function.

The first sub-period that has a non-zero time span must be designated as a specified period since it provides CHEERS with initial conditions for model decision variables. Its duration must be sufficiently long to accommodate all model constraints. For example, if a requirement restricts a reservoir elevation change of over a 24-hour rolling period, the first specified sub-period must span at least one day. That is, CHEERS ensures that the solution it produces complies with time-dependent constraints during the optimization period with respect to both initial conditions and all time steps that occurred prior to the first optimized BU. CHEERS allows the user define decision variable specifications and constraints that apply to either an individual sub-period or multiple sub-periods.

3.5.6.1 Time Constructs

The model has built-in definitions of standard units of time, such as month, day, hour, and minute, which for most purposes, are comprised of one or more time BUs. Other time periods, such as off-peak and on-peak periods, are defined by the user. Both standard time units and user-defined time designations are referred to as "time constructs." Constraints applied to objects and groups of objects can be specified for individual and/or several aggregated time BUs. Constraints and functions applied to objects and groups of objects can be specified for individual time BUs, standard time periods, and user-defined constructs.

Figure 3-40 shows major CHEERS time construct classifications, including long-term, short-term, and real-time designations. As will be described in greater detail later in this report, these are critical time considerations at demonstration sites. These time concepts are not unique to the WUOT demonstration sites, but are generally applicable across all hydropower projects.

The CHEERS model contains several options that reduce problems associated with temporal boundary conditions. These include an extension period in which CHEERS simulations are performed for a specified number of time intervals beyond the focal time span. This technique is especially useful when conditions after the designated endpoint are expected to differ significantly from the period being modeled. Another option is referred to as "wrap," since it ties constraints for one or more end states to the beginning of the simulation period. In this

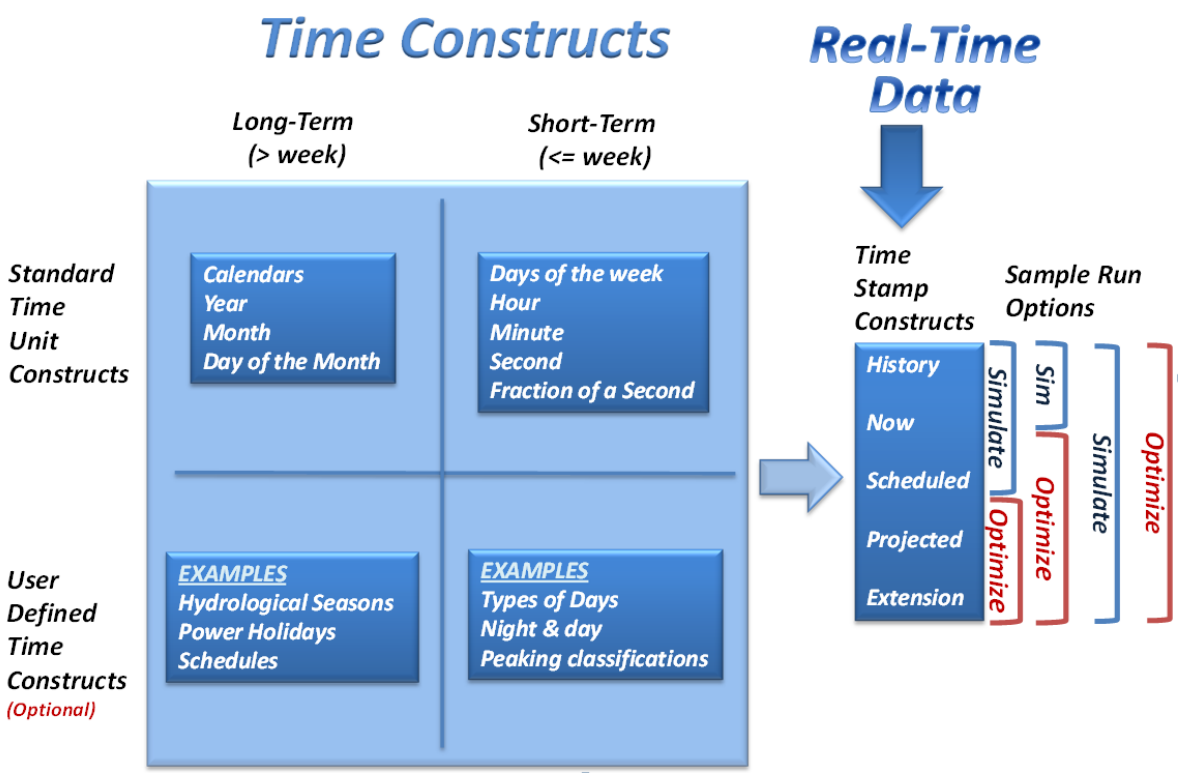

FIGURE 3-40 Major CHEERS Time Construct Classifications and Sub-period Classifications

regard, time becomes a circular continuum. For example, if an object has a down-ramp rate restriction, the difference in operation between the end-state and the end of the first time interval must respect the specified limit. The wrap option is generally applied to situations in which it is anticipated that the current simulation period and the one that follows are similar.

3.5.6.2 Network Balance and Flow

The CHEERS model tracks flows and balances commodity quantities at network nodes and links through time. The quantity of water or electricity in a demonstration site application entering each junction node is assured to equal the quantity leaving it in each time period. It is mathematically modeled as the constraint shown in equation (3-51):

$$\sum_{i:(i,n)\in L} Q_{in}^{t} - \sum_{j:(n,j)\in L} Q_{nj}^{t} = 0, \forall\, t, junction\ node\ n \qquad (3\text{-}51)$$

where Q_{in}^{t} refers to the flow variable in link (i, n), at time t.

The constraint for a boundary node is only slightly different, as shown in equation (3-52):

$$\sum_{i:(n,j)\in L} Q_{nj}^{t} - \sum_{j:(i,n)\in L} Q_{in}^{t} = B_{n}^{t},\ \forall\, t, boundary\ node\ n \qquad (3\text{-}52)$$

where B_{n}^{t} is an input parameter denoting quantity entering the network from the boundary node n at time t.

At storage node, we also need to take into account the quantity in storage that is carried over from the previous time period, as shown in equation (3-53):

$$\sum_{i:(i,n)\in L} Q_{in}^{t} - \sum_{j:(n,j)\in L} Q_{nj}^{t} + S_{n}^{t-1} - S_{n}^{t} = 0,\ \forall\, t, storage\ node\ n \qquad (3\text{-}53)$$

where S_{n}^{t} is the variable denoting the quantity stored in storage node n at the end of time t.

3.5.6.3 Constraints

Many of the decision variables in the model are not allowed to attain values above upper bounds or below lower bounds. These restrictions may arise out of physical constraints, engineering limits, or regulations. An example of a physical constraint is the flow rate in a link, which cannot be less than zero. In the CHEERS model, a user may specify lower and upper bounds for any variable. Mathematically, the constraints take the following form:

$$L \le x \le U,$$

where L and U are the bounds on variable x. If a variable must be forced to a fixed number, the values of both bounds are made equal.

The CHEERS model also constrains the change in decision variable values between two sequential BU time periods. These constraints, referred to as ramping limits, are used in demonstration site applications to limit the hourly change in total water releases from a reservoir. This constraint is modeled using equation (3-54):

$$-R_{SL} \leq Q_S^t - Q_S^{t-1} \leq R_{SU}, \ \forall\, t, T \tag{3-54}$$

where R_{SL}, R_{SU} are parameters denoting maximum ramp-up and ramp-down limits for Storage node S, and Q_S^t is a variable denoting the quantity of water released from Storage S at time t. Similar ramping constraints could be added for restrictions on the total water released from a reservoir and for changes in water flows in river channels.

As illustrated in Figure 3-41, ultimately CHEERS will constrain operations over both time and space. These constraints may be applied to an individual component, groups of components, a single time step, or over a multiple-time step construct. In addition, constraints limit changes between time periods. The user may impose one or more limits on nodes, links, and groupings in a hierarchical manner. Some of these options are still under development. However, we have implemented all features required for demonstration site applications.

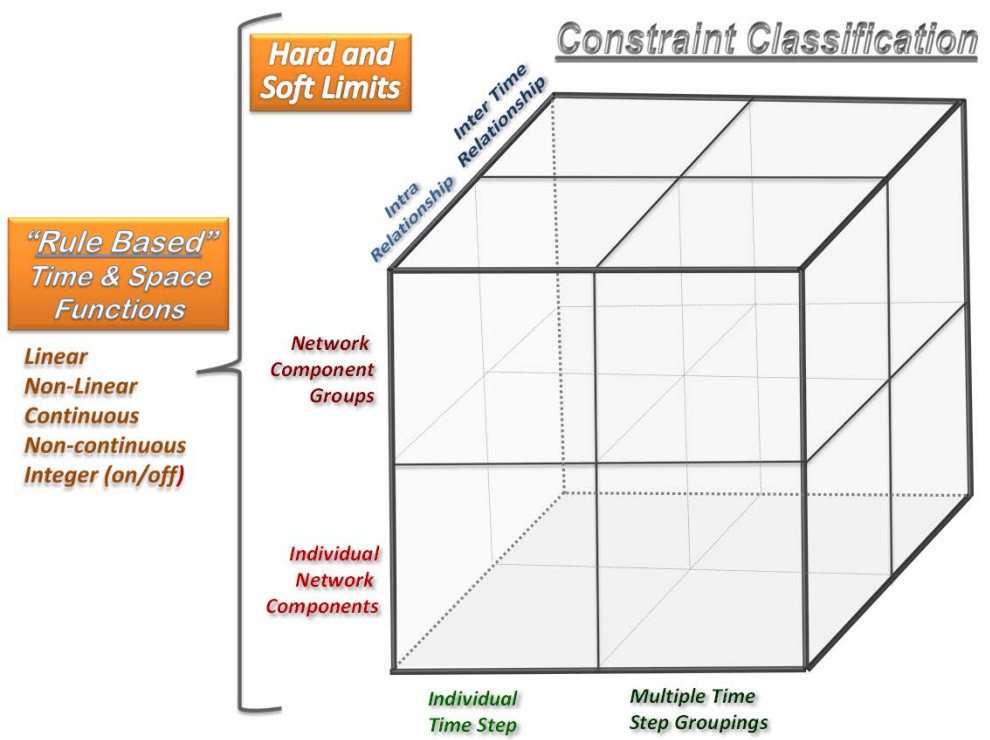

FIGURE 3-41 CHEERS Constraint Classifications over Time and Space

In the current version of the model, the user defines limitations that constrain the solution. The specific limitations may vary across different *cases*. A case is collection of constraint values, relationship equations, and the model result values associated with the use of that particular collection. A user enters constraints into CHEERS via the following three input forms, each of which is a dynamic Excel workbook accessed from the GUI:

1. Base Unit Global Constraints,
2. Base Unit Override Constraints,
3. Sequential Time Pattern Constraints.

Using the first two input forms listed above, limitations are applied to commodity attributes for specific points in time or "timestamps." Input form data are used by CHEERS to apply limits to both node and link objects for all time BUs (Global Constraints) or for individual BUs (Override Constraints). Input form options enable the user to limit maximum and minimum values or apply a fixed set value. Change in values between two consecutive BUs time stamps can also be restricted by applying up and down ramp rate restrictions. Depending on the situation, limitations are specified on total object inflow, outflow levels, or internal values.

Entering constraints with the third form listed above allows the user to apply limits over multiple continuous BUs. These constraints limit maximum and minimum total values over a range of sequential BUs and the range of values. There are also several ramping constraint options. Wrap options are used reduce time boundary end-affect problems.

The next version of CHEERS will include rule-based limitations that are applied to both standard and user defined time constructs, such as the day type (weekdays, weekends, holidays, etc.).

3.5.6.4 Sub-BU Time Considerations

The CHEERS model solves for the scheduling and operation of a system for each BU time step. Although CHEERS can theoretically be run at any time step for hydropower day-ahead and real-time applications, the smallest practical time step is in terms of minutes. Therefore, it cannot solve for movements at sub-BU time intervals such as instantaneous variations in load, continuous changes in generation as a unit ramp between set points, and second-by-second fluctuations in production as units respond to automatic generation control (AGC) signals.

Although CHEERS operates at a BU time interval, it is possible for modelers and analysts to gain some appreciation for the effects of intra-BU movements on overall operations. This section discusses the methodologies CHEERS uses to represent the intra-BU effects of ancillary service and both up and down ramping on operations.

Both spinning and non-spinning reserves are classified as operating reserves. This is idle generating capacity that is available for deployment within a short period of time to fill the

supply void created when a generator operating in the system goes down or by some other grid disruption. Operating reserve requirements are typically set equal to the least of the capacity of the largest generator in service plus a fraction of the peak load (Wang et al. 2005).

Spinning reserve is the spare generating capacity made available by increasing the power output of generators already connected to the grid. Units that provide spinning reserve services must be synchronized to the grid and respond to grid events within a 10-minute time frame.

Non-spinning reserve is off-line generating capacity that can operate after a short delay. This typically equals the power available from fast-start generators such as hydropower and gas turbine technologies. Non-spinning reserve, along with adequate transmission reserve, can also be purchased from an interconnected system. In the Western Interconnect, where two demonstration sites are located, units that supply non-spinning reserves have a 30-minute response time and are not required to be synchronized to the grid.

The most detailed modeling of ancillary services at demonstration site applications occurs at the Thermalito-Oroville Complex. This includes regulation up and down and both spinning and non-spinning reserves. The regulation service is used to compensate for (or react to) very short-term (in the range of seconds) changes in the grid through the use of unit AGCs. A unit providing this service must be synchronized to the grid. Because these changes occur very quickly, the time needed to balance the grid via system re-dispatch is insufficient. In some markets, up and down regulations are sold as two separate services.

Figure 3-42 shows how CHEERS represents ancillary services and its effect on the range of unit operation to serve load and/or for energy sales. In CHEERS regulation up, both spinning and non-spinning reserves are modeled as a committed capacity constraint in day-ahead and real-time scheduling cycles. That is, the sum of those three ancillary services plus scheduled energy production at all times must be less than or equal to the maximum operational capacity of the units. These capacity reserves allow the unit to be deployed as needed by grid operations. CHEERS also ensures that there is adequate water available in the upper reservoir to provide these services if one or more units are deployed. The model also ensures that deployment would not violate any system constraint such as minimum and maximum reservoir forebay elevation limits and maximum water release rates.

At the lower end of the unit operating spectrum, CHEERS increases the minimum unit operating level to accommodate the regulation down service. Therefore, when a unit is deployed to provide this service, technical minimum generation levels will not be violated. The CHEERS model also ensures that deployment will not violate any other system constraint.

In addition to reserving unit capacity at facilities that provide ancillary services, CHEERS also estimates the effects on average hydropower generation and associated turbine water releases. The CHEERS model accomplishes this by computing both set point generation

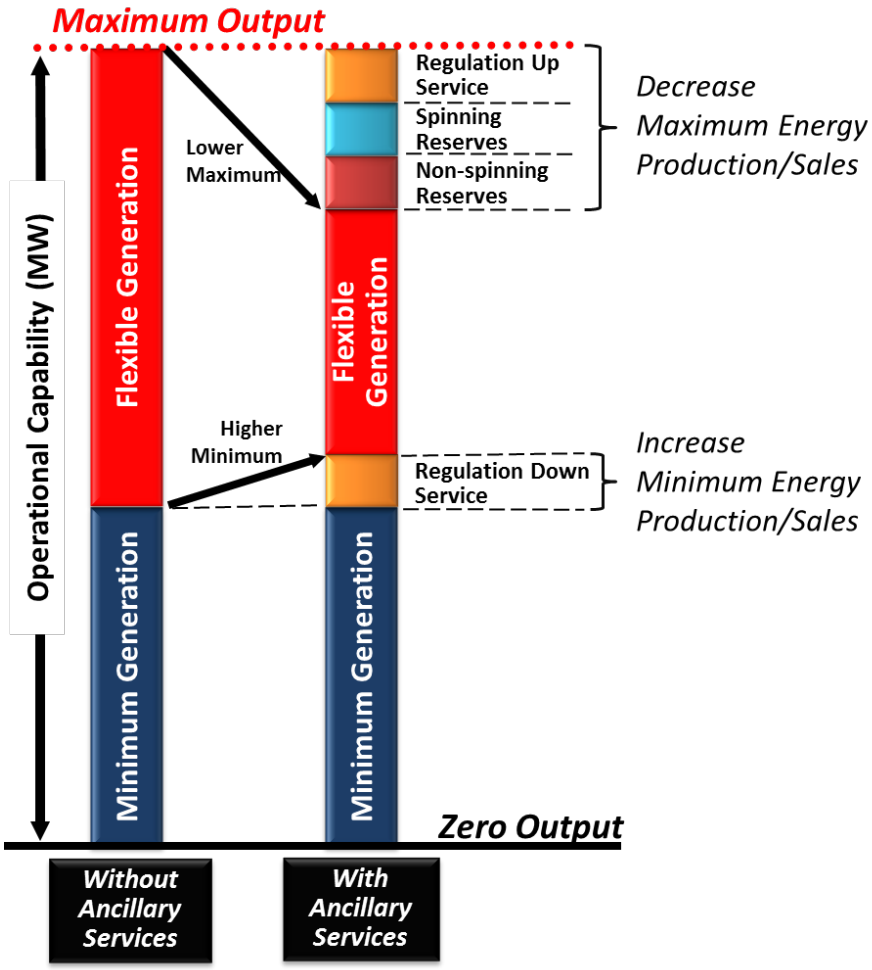

FIGURE 3-42 Operating Range with and without Ancillary Services

and average generation over a BU time period. Figure 3-43 illustrates instantaneous generation from a unit that is providing both up and down regulation. To compute the amount of energy that is generated to support regulation up, CHEERS multiplies a user-defined coefficient by the MW level of regulation service. For example, if 20 MW of regulation service is provided for a 1-hour BU time period and the user enters in a coefficient of 0.5, 10 MWh of energy will have been generated to provide the service. This amount is then added to the set point generation. A similar computation is made for regulation down service. However, the resultant amount is subtracted from the set level generation. Values for these two coefficients may be estimated from historical observations. Detailed analyses of regulation services provided by the Glen Canyon Dam reveal that when identical levels of both regulation up and down are simultaneously provided by the plant, the net effects of regulation on total generation are very near to zero (Brown 2012); that is, regulation up and down energy amounts cancel each other.

Hydropower unit generation movements affect turbine water releases. As shown in Figure 3-37, Special Order Set of Type 2 (SOS2) weights may need to be estimated at multiple

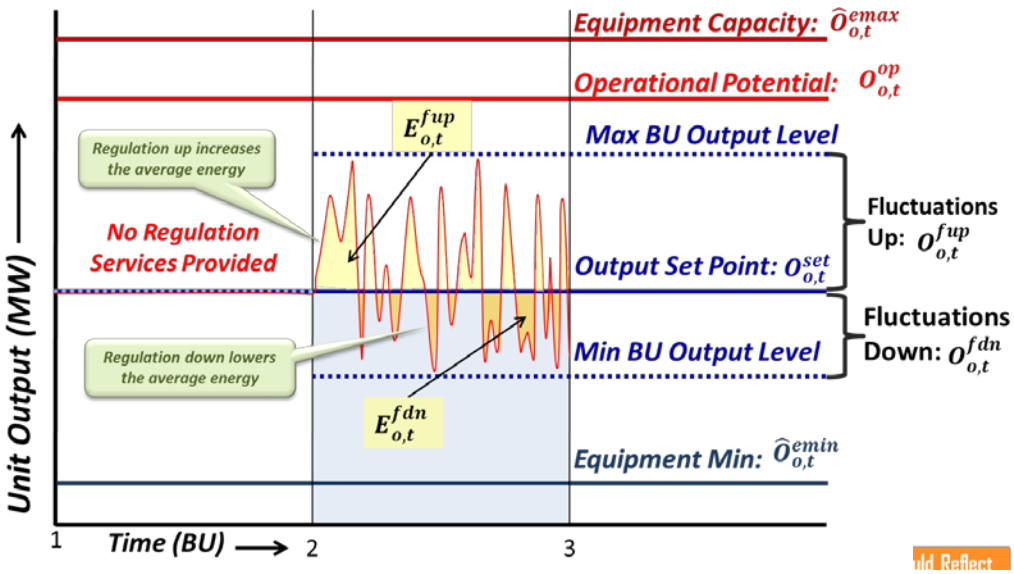

FIGURE 3-43 Regulation Service Impacts on BU Generation Levels

water release rates to obtain an average hourly turbine water release. This includes release rates at the upper and lower ends of the regulation range, along with the flow rate associated with the power output set point

When ramping from one hour's generating set point to a different one, many generating units in the Western Interconnect typically perform this operation from between 10 minutes before the start of the hour through 10 minutes after the end of the hour. Figure 3-44 illustrates how CHEERS represents these effects on generation computations in BU time steps. CHEERS uses simple geometry equations to compute the increase in the average BU generation during the current time step when the next time step has a higher generation set point. Alternatively, when the generation set point is lower in the next time step, average generation in the current BU time is also lower. The CHEERS model also computes the effects of ramping from a previous time step to the current time step on average generation levels. When the previous BU generation set point is lower, average generation in the current BU is also lower. The opposite occurs if the previous BU has a higher operating set point.

Similar to regulation services, SOS2 weights may need to be determined at several release rates in order to estimate the average hourly turbine water releases over a ramping period. These detailed water release computations are not currently made when performing system optimization. It is anticipated that the next version of CHEERS will include this feature. However, these effects are currently computed in the non-linear post-processor routine.

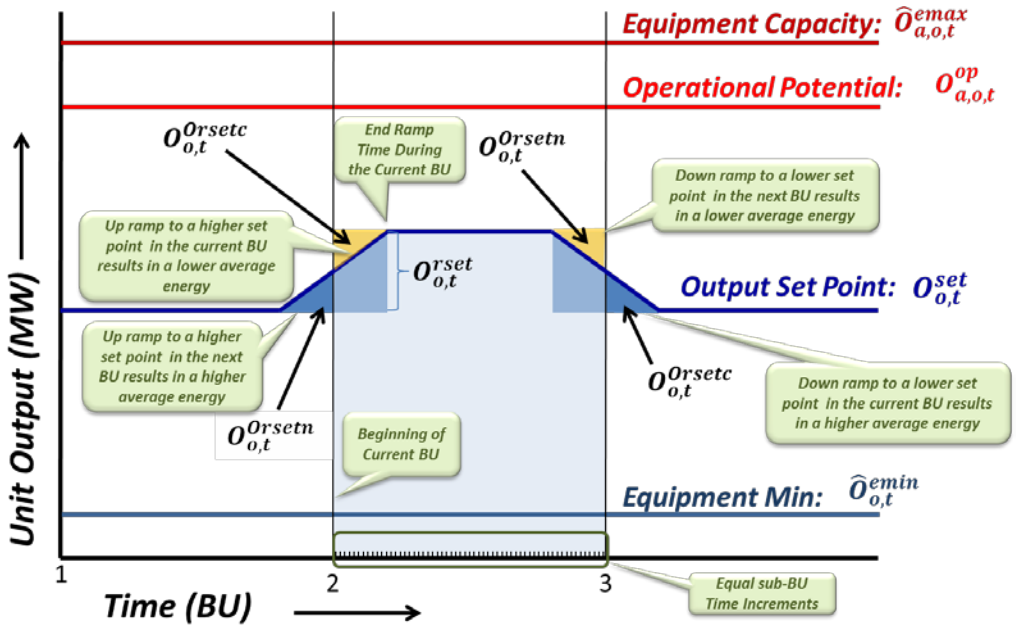

FIGURE 3-44 Ramping Impacts on BU Generation Levels

3.5.6.5 Inflow Forecast Uncertainty

As discussed earlier, water elevation levels in a reservoir are based on a water mass balance equation and a reservoir water volume-to-elevation function. Water release variables in this equation are both predictable and controllable. On the other hand, channel main stem inflows from uncontrolled sources and side flows that either directly empty into a reservoir or into a channel that connects two reservoirs are a function of a large number of variables; many of which are described in the inflow forecasting section and at times are difficult to predict with a high degree of confidence. These uncertainties result in reservoir elevations that differ from predicted levels. Therefore, reservoir elevation operating limits may be violated if schedulers attempt to operate reservoirs when water elevations are at or very near to one of these limits. For example, if a reservoir is at its maximum level and inflows are greater than the projected amounts, insufficient releases may have been scheduled to keep the water elevation below the limit. It should be noted that all reservoirs at demonstration site are constrained by both upper and lower elevation limits.

As illustrated in Figure 3-45, CHEERS uses upper and lower reservoir buffers to reduce, but not eliminate, potential reservoir elevation violations associated with inflow forecast error. At the upper range of the reservoir elevation level, the buffer is used when actual inflows are greater than predicted. The opposite occurs at lower reservoir elevations when inflows are less than expected. Historically, reservoir operating limits have been violated due to inflow forecast error. Appendix K shows several violations that occurred at the Crystal Reservoir during the spring and summer of 2009.

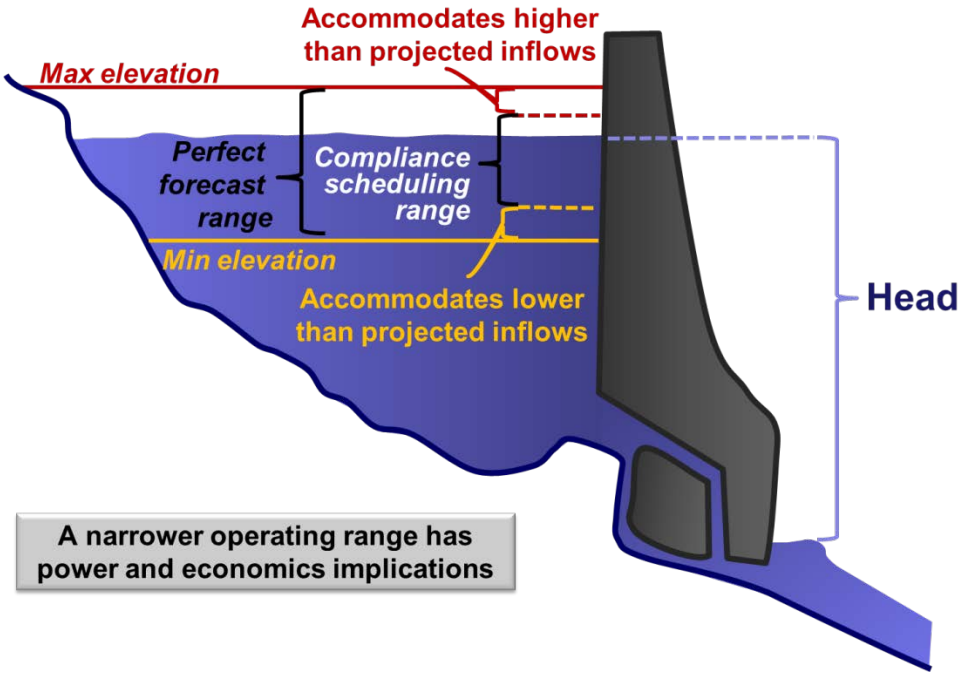

FIGURE 3-45 Illustration of CHEERS Elevation Buffers

The Forecast Error Analyzer routine assists schedulers with decisions regarding reservoir buffer settings. It first computes inflows into the top reservoir and side flows between reservoirs. To illustrate the use of this routine data for the Aspinall Cascade data were downloaded from the CRSP website (http://www.wapa.gov/crsp/opsmaintcrsp/scada.htm). These data contain SCADA data for the 1999 through 2010 time period. Data include hourly reservoir elevations, both turbine and non-turbine water release, generation, and ramping data for each of the Aspinall power plants. Using these data and reservoir elevation-to-water storage volume functions, the water mass balance equation was used to determine hourly net side flows.

Mimicking actual operations and the information available to schedulers, forecast errors were then computed based on the average daily persistence forecasting method. The method assumes that inflows during each hour tomorrow and thereafter will be identical to the average inflow rate that occurred yesterday. More specifically, the estimated inflow forecast for an hour was subtracted from the average daily inflow rate that occurred the day before yesterday. This is identical to actual operations at the CRSP Marketing Office, where each day, the average inflow rate for the previous day is received from the Bureau of Reclamation. The inflow forecast error estimate embeds real-life challenges faced by schedulers including reservoir evaporation impacts, water release rate measurement error, reservoir elevation measurement error, and inaccuracies associated with volume-to-elevation curves.

Historical water inflow forecast errors are used to construct probability exceedance functions such as the one shown in Figure 3-46. This function, along with scheduler's risk tolerance level, is used to set water storage buffer requirements. The figure illustrates upper and

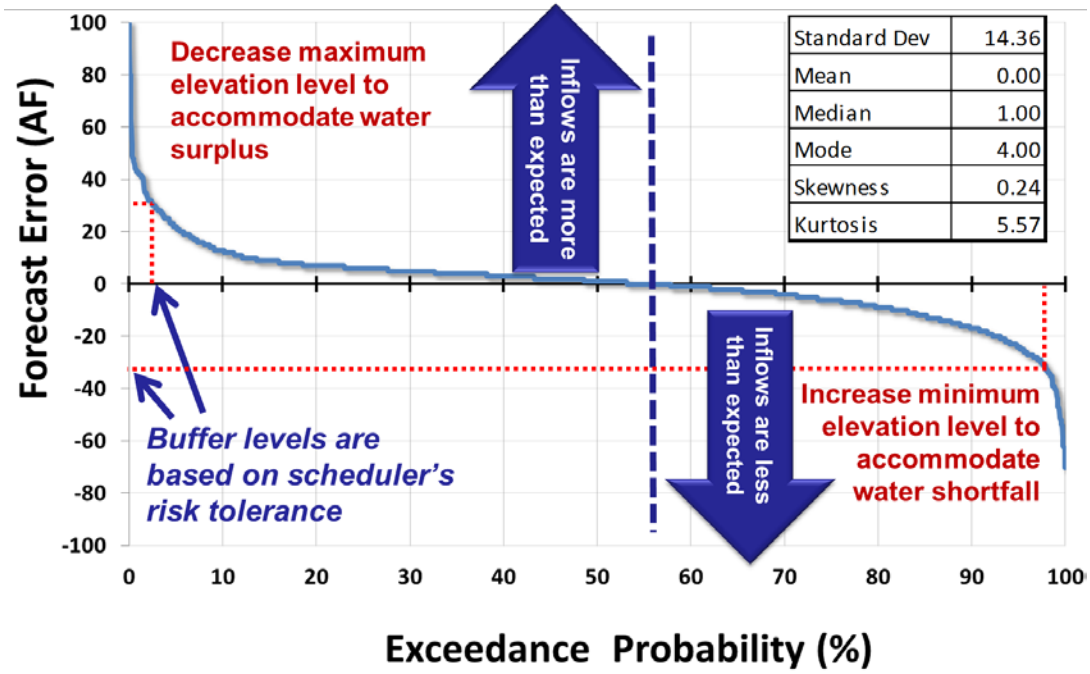

FIGURE 3-46 Forecast Error Probability Distribution Function

lower reservoir buffer levels at a level of 2 percent. This 2 percent risk only applies when the actual reservoir level is at the edge of a buffer zone. The level of risk decreases as conditions move toward the middle of the compliance zone.

The use of reservoir buffers reduces potential reservoir violations; however, the narrower operating range tends to diminish the economic and financial value of power by reducing the operational flexibility power plant operations. Under high reservoir conditions, a lower maximum elevation may necessitate more production at night, when power has a lower value. At the other end of the spectrum—operating under a low reservoir condition—the higher maximum elevation may necessitate less power production during peak load hours, when power has a high value.

A smaller risk level than the one shown in Figure 3-46 will reduce the frequency of potential reservoir violations. However, potential economic costs will increase. Therefore, there are tradeoffs between buffer size and economics costs.

If inflow forecasts were perfect, then the reservoir buffer could be eliminated, leading to a higher economic value of the power resource. In reality, forecasts are never perfect, but reducing the magnitude of forecast errors will lower the buffer requirement needed to achieve a given risk level. The WUOT team is exploring better ways to quantify reservoir buffer amounts, reduce inflow forecast error, and thereby enhance power and economic impacts. One potential improvement is to customize inflow forecasts based on a specific situation or condition. Appendix J provides ones such example, in which inflow forecast error is stratified by time of

day during the rapid snowmelt period. Forecasts may also be improved through the use of a more sophisticated inflow forecasting model such as the WUOT Seasonal Hydro Systems Analysis tool.

3.5.6.6 Solving the Optimization Problem

The CHEERS framework generates the MILP model in the LINGO and AMPL (a mathematical programming language) formats (Fourer et al. 2002). The AMPL modeling language allows model developers to write optimization problems in a format similar to the mathematical notation needed to express them. Several solvers—including CBC, SCIP, CPLEX, and GUROBI—are capable of solving AMPL models. This approach thus provides flexibility in selecting a solver and avoiding solver "lock-in."

Branch-and-cut is the most popular algorithm for solving MILPs, and it is implemented by all of the solvers mentioned above. The method is a combination of two algorithms: the branch-and-bound and cutting-plane algorithms. In both, a linear programming (LP) relaxation obtained by ignoring the integrality restrictions on variables is first solved. If the relaxation is infeasible, then so is the MILP. If the solution of LP relaxation satisfies the integrality restrictions, then it is the optimal solution of the MILP. Otherwise, the two algorithms use different approaches to proceed, as follows. In branch-and-bound, we divide the feasible region into two or more smaller regions and solve their LP relaxations. If the solution value is worse than the incumbent, the problem is discarded. Otherwise, the solvers check for integrality or proceed with more branching. Thus, a solver needs to solve many progressively smaller LP relaxations. In a cutting-plane algorithm, instead of branching, a solver adds a new inequality that cuts off the current LP solution. The new LP, with additional constraints, is solved again, and the process continues until a solution is found or the LP relaxation becomes infeasible.

Different solvers combine the branch-and-bound and cutting-plane algorithms to achieve the maximum efficiency in solving real-life problems. In addition, they use several primal heuristics to find good, feasible solutions quickly. Other advanced techniques like preprocessing, exploiting symmetry, and search-restarts are used to enhance the solution speed. These solvers can also accept a starting solution.

3.5.6.7 Application of Nonlinear Functions and Relationships

In CHEERS optimization modeling of hydropower plant operations, the relationships between reservoir volume and elevation, water flow rate and tail water elevation, power generation and head and power release rate are nonlinear. However, MILP results are largely based on a linearization of the problem. Therefore, CHEERS employs a post-processor routine that uses key MILP results and recomputes values based on either a set of non-linear equations or look-up tables containing raw data. These data tables describe non-linear aspects of system operations. This routine recalculates reservoir elevation (both forebay and afterbay water levels where applicable), tailwater/tailrace elevation, head, power generation, and the economic value

of energy production. Other values may be added to this list of recomputed variables in future releases of CHEERS as additional model requirements are identified.

3.5.6.8 Models of the Nonlinear Relationships and the Validation Process

Optimization results are stored from both the MILP formulation and the post-processor. When significant differences arise, it is an indication that the MILP representation and formulation of the problem need to be revised. Future plans include further developing the post-processor routine to search for better solutions that are in the "vicinity" of the one solved for by the MILP model.

This subsection presents the detailed validation process on the results from linearized optimization model with the nonlinear relationships, to determine the elevation of reservoir, tailwater, and turbine power conversion efficiency. In the validation process, there are two methods used to model the nonlinearity in this section: (1) polynomial regression on data points and (2) data points lookup tables.

3.5.6.9 Reservoir Surface Elevation

As illustrated in Figures 3-32, 3-33, 3-34, and 3-35, the nonlinear relationship is typically concave (i.e., function has a negative second derivative). At a low reservoir elevation, a small change in water storage results in a relatively large elevation change, and when the reservoir is full, a small decrease in reservoir water storage results in a relatively small change in reservoir elevation.

Equation (3-55) expresses this relationship as an nth-order polynomial equation:

$$\sum_{k=1}^{n}(a_d^{k(E)} \times (V_{dh})^k) + b_d^{(E)} = E_{dh}^{Res} \; \forall \, d, h. \qquad (3\text{-}55)$$

Another method is to create a lookup table with historical volume and its corresponding elevation. The table is sorted by the values of volume for better lookup performance. With the lookup table, given a value of volume V_{dh_0}, if it is tabulated, its corresponding elevation can be retrieved directly. If not, a linear interpolation will be applied to estimate the elevation. The algorithm is designed as follows:

Step 1: Search table to find volume V_{dh_0}.

Step 2: If find V_{dh_0}, return its corresponding elevation value E_{dh_0}. Otherwise, go to Step 3.

Step 3: Estimate and return the estimated elevation value E_{dh_0}:

1. Find the smallest volume V_{dh_u} that is greater than V_{dh_0} with elevation E_{dh_u} in the lookup table;

2. Find the largest volume V_{dh_l} that is less than V_{dh_0} with elevation E_{dh_l} in the table;

3. Apply linear interpolation:

$$E_{dh_0} = (E_{dh_u} - E_{dh_l})/(V_{dh_u} - V_{dh_l})*(V_{dh_0} - V_{dh_l}) + E_{dh_l}.$$

3.5.6.10 Tailwater Elevation

As shown in Figure 3-3, the elevation of the tail-water at the base of a dam is a nonlinear function of the total water release rate from the reservoir. This includes both turbine and non-turbine releases. Note that as the water release rate increases the first derivative of the function decreases.

The tail-water function is expressed as an nth-order polynomial equation shown in equation (3-56):

$$\sum_{k=1}^{n}(a_d^{k(TW)} \times (RD_{dh})^k) + b_d^{(TW)} = E_{dh}^{TW} \; \forall \, d, h. \tag{3-56}$$

A lookup table that describes the tailrace elevation can also be used. It is similar as the lookup table method for reservoir elevation except that the data are for release rate and tailwater elevation.

3.5.6.11 Turbine Power Conversion Efficiency

Turbine power conversion efficiency is a nonlinear function of the water flow rate through a turbine and the head of the water, which can be represented as shown in equation (3-57):

$$TE_{dh}^{T} = f(RT_{dh}^{T}, H_{dh}) \tag{3-57}$$

where $f(.)$is a nonlinear function.

Practically, CHEERS does not estimate the surface directly. Instead, it samples H_{dh} and does the regression of the polynomial function of TE_{dh}^{T} given RT_{dh}^{T} and a fixed H_{dh}sample (H_{dh}^{m}, m is the index of the sample). Based on the turbine output power production and head the turbine efficiency is calculated using equation (3-58):

$$\sum_{k=1}^{n}(a_d^{k(TE)T,H_d^m} \times (RT_{dh}^{T})^k) + b_d^{(TE)T,H_d^m} = TE_{dh}^{T} \; \forall \, d, h, T. \tag{3-58}$$

Similarly, another method uses a lookup table. Since the conversion efficiency is determined by two factors: flow rate and head, the tables is different to the previous two. When the value of a combination of flow rate and head is not tabulated, a two-dimensional interpolation method is used to estimate the efficiency. The detail is as follows.

Step 1: Use the lookup table to find flow rate RT_{dh_0} and head H_{dh_0}.

Step 2: If found, return its corresponding efficiency value TE_{dh_0}; otherwise, go to Step 3.

Step 3: Estimate and return the estimated elevation value TE_{dh_0}:

1. Find the smallest head H_{dh_u} that is greater than H_{dh_0} in the lookup table.

2. With head H_{dh_u}, find the smallest flow rate RT_{dh_uu} that is greater than RT_{dh_0} and the largest flow rate RT_{dh_ul} that is less than RT_{dh_0} in the lookup table to get points (H_{dh_u}, RT_{dh_uu}, TE_{dh_uu}) and (H_{dh_u}, RT_{dh_ul}, TE_{dh_ul}).

3. With head H_{dh_l}, find the smallest flow rate RT_{dh_lu} that is greater than RT_{dh_0} and the largest flow rate RT_{dh_ll} that is less than RT_{dh_0} in the lookup table to get points (H_{dh_l}, RT_{dh_lu}, TE_{dh_lu}) and (H_{dh_l}, RT_{dh_ll}, TE_{dh_ll}).

4. Apply two-dimensional interpolation:

$$t = (H_{dh_0} - H_{dh_ll})/(H_{dh_uu} - H_{dh_ll});$$
$$u = (RT_{dh_0} - RT_{dh_ll})/(RT_{dh_uu} - RT_{dh_ll}); \text{ and}$$
$$TE_{dh_0} = (1-t) \times (1-u) \times TE_{dh_ll} + t \times (1-u) \times TE_{dh_rl} + (1-t) \times u \times TE_{dh_lr} + t \times u \times TE_{dh_rr}.$$

3.5.6.12 Iteration Process to Estimate Head Value

This section presents how to estimate with an iterative method the value of head at hour h for a dam system with N turbines, given water release rate and power output of each turbine *n*.

Step 1: set two start values (H_{dh}^{s-1}, H_{dh}^{s}) of the head H_{dh}, commonly, H_{dh}^{s-1} is 0, H_{dh}^{s} is the value of head from previous hour $H_{d,h-1}$.

Step 2: check $abs(H_{dh}^{s-1} - H_{dh}^{s})$, if it is less than a convergence threshold, return H_{dh}^{s} as the value of H_{dh}; otherwise, go to Step 3.

Step 3: For each turbine t within the same dam d,

1. Calculate flow rate ($R_{dh}^{T,s}$) given the value of head H_{dh}^{s} and power output $PQ_{dh}^{T,s}$ in current iteration;

2. Calculate reservoir volume V_{Vh}^{s}based on the flow rate $R_{dh}^{T,s}$, and then get the corresponding elevation E_{Vh};

3. Calculate tailwater elevation TW_{dh} based on the flow rate $R_{dh}^{T,s}$;

4. Get head value of the next iteration $H_{dh} = E_{Vh} - TW_{dh}$; and

5. $H_{dh}^{s-1} = H_{dh}^{s}$; $H_{dh}^{s} = H_{dh}$; return to Step 2.

Note that in Steps 3.1, 3.2, and 3.3, there are nonlinear relationships described in previous section, but in reversed direction. Again, CHEERS applies both polynomial regression and lookup tables to address the nonlinear validation.

3.5.6.13 Nonlinear Validation Process for Ancillary Services

The characteristics of hydropower plants make them well-suited for fulfilling ancillary service requirements including both regulation and spinning reserves. First, a hydropower turbine can start up or shut down in a relatively short period. The second reason is that the output of a hydropower turbine can be changed relatively quickly by adjusting input water flow. In this section, a nonlinear validation post-process considering regulation and spinning reserve services is presented. The first step is to calculate the adjusted flow rate with ramping up/down between consecutive time intervals and regulation service. Then the second step is to check whether there are enough resources to provide spinning reserve services.

Ramping Up/Down between Consecutive Intervals and Regulation Up/Down Services. The power output of a hydropower plant cannot be changed immediately when entering into the next time interval, a hypothetical operation shown by the dashed line in Figure 3-47. The output must be ramped up or down to the scheduled output of the next time interval, shown by the purple line in Figure 3-47, by adjusting the water release rate from the reservoir. In this case, the realized water release will be different than the one based only on scheduled power output. Moreover, a hydropower plant can potentially provide regulation service, and when doing so its realized power output could deviate from its schedule as shown by the dark blue line in Figure 3-47. The water release rate of a hydropower plant is updated considering these two factors.

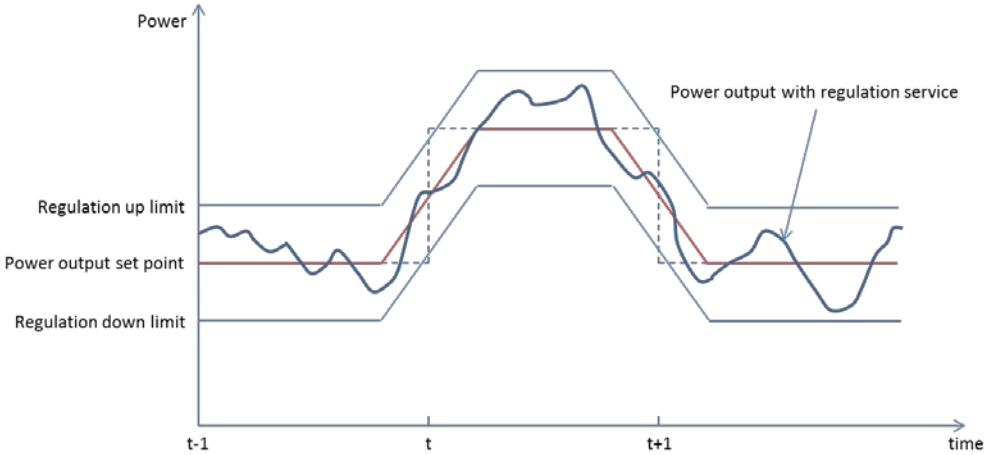

FIGURE 3-47 Power Output Curve with Ramping and Regulation Service

Spinning reserve is additional generation capacity that can be dispatched within a time limit to account for the variability of load uncertainty or some other system contingency. For a hydropower plant, it is important to check whether there is enough capacity to provide spinning reserve service. To provide spinning up reserve, a hydropower plant must potentially send more water flow through turbines. In this case, the validation process checks whether there is enough water in a reservoir and enough space in the tailwater/afterbay to accommodate the additional water potentially released, subject to constraints on release rate and reservoir/afterbay/tailwater volume limits. For spinning down reserve, the water through a hydropower plant is potentially adjusted down to reduce power output. In this case, the validation process checks whether there is enough space in a reservoir to accommodate scheduled inflow water for the next time interval, and is subjected to constraints on release rate limit, reservoir/afterbay/tailwater volume limit.

3.6 REFERENCES

Abhishek, K., S. Leyffer, and J.T. Linderoth, 2010, “FilMINT: An Outer Approximation-Based Solver for Convex Mixed-Integer Nonlinear Programs,” *INFORMS Journal on Computing*, 22: 555–567.

Achterberg, T., 2009, “SCIP: Solving Constraint Integer Programs,” *Mathematical Programming Computation*, 1(1): 1–41.

Adams, B.M., et al., 2010, DAKOTA, A Multileve Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis, Version 5.0+ Users' Manual, Sandia National Laboratories.

Ahuja, R.K., T.L. Magnanti, and J.B. Orlin, 1988, *Network Flows: Theory, Algorithms, And Applications*, Prentice Hall, NJ.

Bartholow, J.M., 1989, *Stream Temperature Investigations: Field and Analytic Methods*, Instream Flow Information Paper No. 13, Biological Report 89(17), U.S. Fish and Wildlife Service.

Belotti, P., 2009, *COUENNE: A User's Manual*, Department of Mathematical Sciences, Clemson University, Clemson, SC. Available at http://www.coin-or.org/Couenne/couenne-user-manual.pdf. Accessed April 23, 2012.

Bonami, P., L.T. Biegler, A.R. Conn, G. Cornuejols, I.E. Grossmann, C.D. Laird, J. Lee, A. Lodi, F. Margot, N. Sawaya, and A. Waechter, 2008, "An Algorithmic Framework for Convex Mixed Integer Nonlinear Programs," *Discrete Optimization*, 5(2): 186–204.

Bras, R.A., 1990, Hydrology, an Introduction to Hydrologic Science. Addison-Wesley.

Brown, C., 2012, personal communication from C. Brown (CRSP Marketing Office, Montrose, CO) to T. Veselka (Argonne National Laboratory), November 7.

Brunner, G.W., 2010, *HEC-RAS River Analysis System. User's Manual v4.1.0*, U.S. Army Corps of Engineers, Institute for Water Resources.

Cole, T.M., and S.A. Wells, 2006, *CE-QUAL-W2: A Two Dimensional, Laterally Averaged, Hydrodynamic and Water Quality Model, Version 3.5*, Instruction Report EL-06-1,U.S. Army Corps of Engineers, Washington, DC.

Davis, D.W., 1986, *HEC-5, Simulation of Flood Control and Conservation Systems, Appendix on Water Quality*, U.S. Army Corps of Engineers, Hydrologic Engineering Center, Davis, CA.

Day, G.N., 1985, "Extended Streamflow Forecasting Using NWSRFS," *Journal of Water Resources Planning and Management* 111(2): 157–170.
Eaton, M.R., 1998, Restoration of Natural Flooding Process on the Cosumnes River Preserve. Overview and Status Report, The Nature Conservancy, San Francisco, CA.

FICO, 2012, "Test Drive the Xpress-Optimization Suite." Available at http://www.fico.com/en/Products/DMTools/xpress-overview/Pages/Xpress-Optimizer.aspx. Accessed April 23, 2012.

Forrest, J.J., 2006, "Coin Branch-and-Cut Solver," July 18. Available at http://dimacs.rutgers.edu/Workshops/COIN/slides/forest2.pdf. Accessed April 23, 2012.

Forrest, J.J., and E.M.L. Beale, 1976, "Global Optimization Using Special Ordered Sets," *Mathematical Programming*, 10: 52–69.

Fourer, R., D.M. Gay, and B.W. Kernighan, 2002, *AMPL: A Modeling Language for Mathematical Programming*, Brooks/Cole Publishing Company.

Gibson, A.J.F., and R.A. Myers, 2002, "A Logistic Regression Model for Estimating Turbine Mortality at Hydroelectric Generating Station," *Transactions of the American Fisheries Society* 131: 623–633.

GUROBI Optimization, undated, "An Easier Way to Better Decisions." Available at http://www.gurobi.com. Accessed April 23, 2012.

Hamill, T.M., J.S. Whitaker, S.L. Mullen, 2006, "Reforecasts: An Important Dataset for Improving Weather Predictions," *Bulletin of the American Meteorological Society* 87: 33–46.

Hayse, J.W., K.E. LaGory, and G.L. Burton. 2005. Consideration of Site-Specific Floodplain Inundation Thresholds in Implementing Peak Flow Magnitude and Duration Recommendations in the Middle Green River, Utah. Final report prepared for Western Area Power Administration, Salt Lake City, Utah (June).

IBM, undated, "IBM ILOG CPLEX Optimizer: High-Performance Mathematical Programming Solver for Linear Programming, Mixed Integer Programming, and Quadratic Programming." Available at http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/. Accessed April 23, 2012.

Kimball, J.S., S.W. Running, and R. Nemani, 1997, "An Improved Method for Estimating Surface Humidity from Daily Minimum Temperature," *Agricultural and Forest Meteorology* 85(1-2): 87–98, ISSN 0168-1923, DOI: 10.1016/S0168-1923(96)02366-0.

Kirkwood, C.W., 1998, *System Dynamics Methods: A Quick Introduction*, Arizona State Univerity, College of Business.

Liang, X., D.P. Lettenmaier, E.F. Wood, and S.J. Burges, 1994, "A Simple Hydrologically Based Model of Land Surface Water and Energy Fluxes for GSMs," *Journal of Geophysical Research* 99(D7):14,415–14,428.

Lohmann, D., R. Nolte-Holube, and E. Raschke, 1996, "A large-scale Horizontal Routing Model to be Coupled to Land Surface Parameterization Schemes," *Tellus* 48(A): 708–721.

Lohmann, D., E. Raschke, B. Nijssen, and D.P. Lettenmaier, 1998, "Regional Scale Hydrology: I. Formulation of the VIC-2L Model Coupled to a Routing Model," *Hydrol. Sci. J.* 43(1): 131–141.

Mahajan, A., S. Leyffer, J.T. Linderoth, and J. Luedtke, 2012, "MINOTAUR wiki page," Argonne National Laboratory, April 12. Available at http://wiki.mcs.anl.gov/minotaur/index.php/Main_Page. Accessed April 23, 2012.

Maurer, E.P., A.W. Wood, J.C. Adam, D.P. Lettenmaier, and B. Nijssen, 2002, "A Long-term Hydrologically-based Data Set of Land Surface Fluxes and States for the Conterminous United States," *J. Climate*. 15: 3237–3251.

Papadimitriou, C.H., 1981, "On the Complexity of Integer Programming," *Journal of the Association for Computing Machinery*, 28: 765–768.

Petts, G.E., 1984, Impounded Rivers: Perspectives for Ecological Management, John Wiley & Sons, Chichester, UK.

Poff, N.L., J.D. Allan, M.B. Bain, J.R. Karr, K.L. Prestegaard, B.D. Richter, R.E. Sparks, and J.C. Stromberg, 1997, "The Natural Flow Regime," *BioScience* 47: 769–784.

Poff, N.L., B.D. Richter, A.H. Arthington, S.E. Bunn, R.J. Naiman, E. Kendy, M. Acreman, C. Apse, B.P. Bledsoe, M.C. Freeman, J. Henriksen, R.B. Jacobson, J.G. Kennen, D.M. Merritt, J.H. O'Keeffe, J.D. Olden, K. Rogers, R.E. Tharme, and A. Warner, 2010, "The Ecological Limits of Hydrologic Alteration (ELOHA): a New Framework for Developing Regional Environmental Flow Standards," *Freshwater Biology* 55: 147–170.

Postel, S., and B. Richter, 2003, *Rivers for Life: Managing Water for People and Nature*. Island Press, Washington, D.C.

Ralphs, T., M. Guzelsoy, and A. Mahajan, 2011, *SYMPHONY Version 5.3 User's Manual*, COR@L Laboratory, Industrial and Systems Engineering Department, Lehigh University, Bethlehem, PA.

Raymond, H.L., 1979, "Effects of Dams and Impoundments on Migrations of Juvenile Chinook Salmon and Steelhead from the Snake River, 1966 to 1975," *Transactions of the American Fisheries Society* 108: 505–529.

Richardson, G.P., and A.L. Pugh, III, 1981, *Introduction to System Dynamics Modelling with DYNAMO*, Productivity Press, MA.

Richter, B.D., J.V. Baumgartner, J. Powell, and D.P. Braun, 1996, "A Method for Assessing Hydrologic Alteration within Ecosystems," *Conservation Biology* 10: 1163–1174.

Rood, S.B., G.M. Samuelson, J.H. Braatne, C.R. Gourley, F. Hughes, and J.M. Mahoney, 2005, "Managing River Flows to Restore Floodplain Forests," *Frontiers in Ecology and the Environment* 3(4): 193–201.

Saha, S., and J.W. Hayse, 2012, An Environmental Evaluation and Optimization Tool for Assessing the Functional Performance of Rivers Regulated for Hydropower Operation. Proceedings of HydroVision International 2012, Louisville, KY, July.

Sahinidis, N.V., 1996, "BARON: A General Purpose Global Optimization Software Package," *Journal of Global Optimization*, 8: 201–205.

Schneider, M.L., S.C. Wilhelms, and L.I. Yates, 2004, "SELECT Version 1.0 Beta: A One-Dimensional Reservoir Selective Withdrawal Model Spreadsheet," U.S. Army Corps of Engineers, Water Operations Technical Support Program, Feb.

Shukla, S., and D.P. Lettenmaier, 2011, "Seasonal Hydrologic Prediction in the United States: Understanding the Role of Initial Hydrologic Conditions and Seasonal Climate Forecast Skill," *Hydrol. Earth Syst. Sci. Discuss. (in review)*, 8: 6565–6592, doi:10.5194/hessd-8-6565-2011.

Thomann, R.V., and J. A. Mueller, 1987, *Principles of Surface Water Quality Modeling and Control*, HarperCollins, New York, NY.

Thornton, P.E., and S.W. Running, 1999, "An Improved Algorithm for Estimating Incident Daily Solar Radiation from Measurements of Temperature, Humidity, and Precipitation," A*g For Met* 93: 211–228.

Tomlin, J.A., 1981, "A Suggested Extension of Special Ordered Sets to Non-Separable Non-Convex Programming Problems," in *Annals of Discrete Mathematics (11) Studies on Graphs and Discrete Programming*, P. Hanson (Ed.), North-Holland Publishing Company, Amsterdam, NY.

Wang, J., X. Wang, and Y. Wu, 2005, "Operating Reserve Model in the Power Market," *IEEE Transactions on Power Systems*, vol. 20, no. 1, Feb.

Ward, W., and J.A. Stanford, 1995, "The Serial Discontinuity Concept: Extending the Model to Floodplain Rivers," *Regulated Rivers: Research and Management* 10(2-4): 159–168.

Wood, A.W., and D.P. Lettenmaier, 2006. "A Testbed for New Seasonal Hydrologic Forecasting Approaches in the Western U.S.," *Bulletin of the American Meteorological Society* 87(12): 1699–1712, doi:10.1175/BAMS-87-12-1699.

4 WUOT DEMONSTRATIONS

4.1 CONCEPT

One intent of this project was that subsequent to the development and demonstration phase, the WUOT be embraced by U.S. hydropower planners and operators and applied across the United States. For that to happen the potential benefits of widespread application needed to be demonstrated successfully in multiple and varied hydropower and environmental situations. As a result, multiple demonstrations of the WUOT were established as a key component of the project.

4.2 DESIGN

One design parameter for the project was that the WUOT be sufficiently flexible to be utilized in different geohydrologic basins and hydropower operations across the United States. Three sites were chosen to demonstrate this flexibility: the Oroville Complex on the Feather River in California, the upper Colorado portion of the Colorado River Storage Project, and the Conowingo Dam complex on the Susquehanna River. The project team canvassed hydropower operators across the United States to identify candidate sites. The three demonstration sites were chosen because they represent a wide range in environmental and hydropower operational conditions and the organizations responsible for operating these systems agreed to provide resources necessary to conduct the demonstrations. Demonstration site operators became members of the project's technical review team, which provided system design information, input data, and review of results, and were integrally involved in all phases of the demonstration process. Conducting the site demonstrations was an iterative process conducted through meetings, webinars, conference calls, and individual communications. Setup and operation of the WUOT requires staff resources to develop site topologies, identify and obtain input and characterization information, populate the common database, make model runs, and evaluate model outputs. The extent and nature of these requirements is dependent on the complexity of the hydropower system under study and the level of analysis to be conducted.

A protocol was established for conducting the site demonstrations. The protocol consists of three phases that engage both study team and demonstration site staff:

4.2.1 Protocol for Conduct of Water Use Optimization Tool Set Site Demonstrations

4.2.1.1 Phase 1: Pre-Demonstration Actions

1. Identify demonstration sites.

2. Identify demonstration site staff (principal, secondary).

3. Identify project team leads (principal, secondary).

4. Set up, test secure Web site (will host data, references, documents, and links to useful information).

5. Conduct information exchange:
   a. Send site staff development documents or related information and demonstration protocol; and
   b. Conduct pre-demonstration webinars exploring tool set demonstration process (expectations, process, schedules, and tool set components to be used for the site demonstration).

6. Demonstration kickoff meeting:
   a. Establish demonstration kickoff date and format (e.g., site visits, conference call/webinar); and
   b. Conduct meeting: toolset overview, data needs and exchange discussion, establish points of contact, and identify major modeling issues.

4.2.1.2 Phase 2: Initial Demonstration

1. Through demo site collaboration with lab staff:
   a. Establish model performance objectives/measures;
   b. Design initial site topologies;
   c. Identify initial baseline and test conditions;
   d. Establish applicable tool set data needs/availability; and
   e. Identify sources/gaps/mitigations and exchange process.

2. Lab staff will then create an initial model design and perform test runs via:
   a. Topology creation;
   b. Organizing information;
   c. Populating database; and
   d. Performing initial model runs.

3. Provide demo sites with results (outputs, issues of concern, proposed issue mitigations).

4. Identify data and analytical gaps with demo site staff.

5. Revise tool set design/software as appropriate.

Iteration 2. Enhanced Demonstration

1. Based on initial model results demonstration site staff will provide labs guidance to make improvements to the following:

   a. Topology design;
   b. Component characterization; and
   c. Operational conditions.

2. Lab staff will then revise tool kit components:
   a. Revise topologies, test conditions, tool set models, modeling;
   b. Review revised results with demo site staff;
   c. Identify gaps and needs; and
   d. Revise tool set design/software as appropriate.

Additional Iterations:

Step 2e repeated one or more times as required to further enhance demonstration site representation.

4.2.1.3 Phase 3: Site Demonstration Reporting

1. Lab staff will prepare a final report that will be reviewed by demonstration site staff. The report will include the following:
   a. Description of the demonstration site;
   b. Model representation of the system;
   c. Description of model results and actual operations; and
   d. Where applicable and identifiable, discussion of potential operational improvements to the system:
      i. Economics of energy and grid services,
      ii. Efficiency gains, and
      iii. Environmental improvements.

4.3 DEMONSTRATION SITE DESCRIPTIONS

4.3.1 Oroville-Thermalito Complex Demonstration Site

The first WUOT demonstration site is the Oroville-Thermalito Complex. This complex is an efficient water storage and delivery system and electricity producer. As shown in Figure 4-1, the complex is located in northern California. It is part of the California State Water Project (SWP), the largest state-built water system and electric power project in the United States. This project includes pumping-generating power plants; reservoirs, lakes, and storage tanks; and canals, tunnels, and pipelines that capture, store, and convey water to 29 water agencies. The Oroville-Thermalito Complex is owned and operated by the CDWR (CDWR undated a, 2012a).

As depicted in Figure 4-2, the Oroville-Thermalito Complex comprises Lake Oroville and the Oroville Dam, Hyatt Pumping-Generating Power Plant (PGP), Thermalito Diversion

FIGURE 4-1 Location of the Oroville-Thermalito Complex in California

Dam (TDD) and Power Plant, Feather River Fish Hatchery, Thermalito Power Canal, Thermalito Forebay, Thermalito PGP, Thermalito Afterbay, and the Lake Oroville Visitors Center. Lake Oroville is the SWP's largest reservoir. The Oroville-Thermalito Complex stores about 3.5 million acre-feet (AF) of water and generates power from water releases through three generating power plants (CDWR 2012b).

Lake Oroville is located in the foothills on the western slope of the Sierra Nevada, 1 mi downstream of the junction of the Feather River's major tributaries. Rain and snowmelt drains from the surrounding mountainsides into waterways that lead into Lake Oroville. Work on the dam site began in 1961, and the reservoir was filled to capacity in 1967. The reservoir has a maximum operating storage of 3,537,580 AF of water at an elevation of 900 ft above sea level. At this elevation, the surface area of the lake is 15,810 acres, and the shoreline extends 167 mi. The lake stores water and releases it into the Feather River to meet SWP needs. The lake also provides pumped-storage capacity, 750,000 AF of flood control storage, and recreation, as well as freshwater releases to control salinity intrusion in the Sacramento-San Joaquin Delta and to provide fish and wildlife protection (CDWR 2012c).

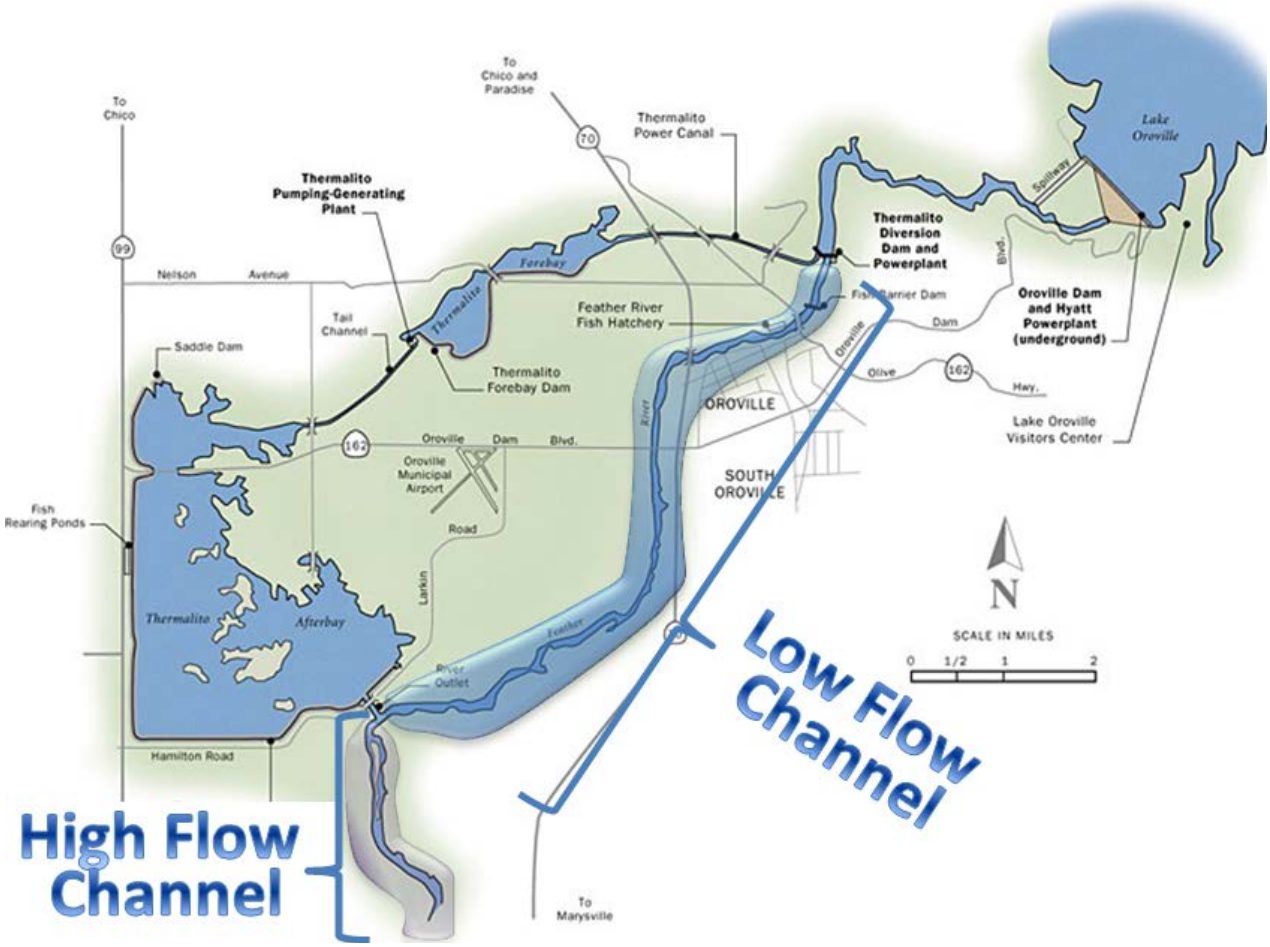

FIGURE 4-2 Map of the Oroville-Thermalito Complex (CDWR 2013d)

Oroville Dam is a zoned, earth-filled structure comprising 80 million cubic yards of material that contains Lake Oroville. At a height of 770 ft, it is the tallest dam in the United States. The dam's crest elevation is 922 ft above sea level. It has a crest length of 6,920 ft and a crest width of 50.6 ft. In 1967, the California Society of Professional Engineers named the Oroville Dam one of the seven engineering wonders of California.

Water is released from Lake Oroville through the Hyatt PGP, river outlets, the Palermo Canal Outlet Works, spillway, and the emergency spillway. Turbine water flows for power generation are controlled by wicket gates and can also be stopped by spherical turbine shutoff valves. Water releases through the dam's two river outlet tubes, which bypass the turbines, are controlled by fixed-dispersion cone valves. Water is also released through the Palermo Canal Outlet Works to serve local water needs. Emergency spillway water releases are controlled by eight radial gates. A sophisticated water intake structure for turbine water releases allows reservoir water to be withdrawn at different reservoir depths to regulate downstream water temperatures in the Feather River.

4.3.1.1 Hyatt PGP

The Edward Hyatt PGP is an underground, hydroelectric, pumping-generating facility. Construction of the plant began in 1964 and was completed in 1967. The Hyatt PGP is the largest power plant in the Oroville-Thermalito Complex, with a total installed generating capacity of 819 MW (CDWR 2012d) and a pumping capacity of 387 MW. The PGP is operated to maximize the value of the hydropower resources. Pumped-storage is one example of this kind of operation: water, released for power in excess of local and downstream requirements, is returned to storage in Lake Oroville during off-peak periods and is used for generation during peak power demands.

Located in the bedrock cavern under Lake Oroville, the power plant facilities include an intake structure, two penstock tunnels, six penstock branch lines, three turbine units, three reversible pump-turbine units, two tailrace tunnels and outlet works, and a switchyard. Water from the lake is conveyed to the units through two penstocks, each with a diameter of 22 ft.

Three of the six generating units are motor-generators coupled to Francis-type, reversible pump turbines that permit the off-peak pumped-storage operations at the Oroville-Thermalito Complex. In generation mode, each vertical Francis turbine can accommodate a water flow rate up to 2,800 cfs, producing 115 MW of electricity. After passing through the turbines, water is discharged into the Feather River via draft tubes to two tailrace tunnels: one free-surface and the other full-flow. Using a maximum of 173,000 horsepower (hp), each of the three units can pump water at a rate up to 1,870 cfs at a rated head of 592 ft.

The three other units at Hyatt PGP are conventional generators driven by vertical-shaft, Francis-type turbines. Each generating unit has a maximum output of 173,000 hp. Water releases for power at each unit are up to 2,850 cfs, discharged into the Feather River via the draft tubes and tailrace tunnels described above. Each unit has a generating capability of about 123 MW (Falaki and Loghmanpour).

4.3.1.2 TDD and Powerplant

The TDD is located on the Feather River, about 4.5 mi downstream from the Oroville Dam. Construction of the 143-ft-high, concrete, gravity diversion dam began in 1963 and was completed in 1968. The dam comprises 154,000 cubic yards of material and has a crest length of 1,300 ft. The crest elevation is 233 ft above sea level. Releases from both the Hyatt PGP and the Kelly Ridge Powerplant drain into a pool contained by the TDD. This tailwater pool serves as a lower reservoir for Hyatt PGP pumping needs. At an elevation of 225 ft above sea level, the pool reaches its maximum water operating storage of 13,350 AF with a water surface area of 320 acres and a shoreline that extends for 10 miles.

The Kelly Ridge Power Plant is a run-of-river hydropower resource that is owned and operated by the Oroville-Wyandotte Irrigation District. It sends water from the Miners Ranch Reservoir into the Kelly Ridge Tunnel, which leads to the power plant penstock. The power plant has a generating capacity of 10 MW and produces an average of 3.8 gigawatt hours (GWh) of

electricity each month (BOR 2003). For day-ahead and real-time operations, water discharges into the pool can be projected with a high level of certainty.

The TDD serves three purposes: (1) to divert water into the Thermalito Power Canal for power generation at the Thermalito PGP; (2) to create a tailwater pool called the Thermalito Diversion Pool for the Hyatt PGP; and (3) to provide headwater for the TDD power plant.

The Thermalito Power Canal, a bidirectional flow canal, links the Thermalito Diversion Pool to the Thermalito Forebay. The canal is designed to convey a maximum of 16,900 cfs of generating flow and 9,000 cfs of pumping flow.

The Diversion Pool serves as a forebay when the Hyatt PGP is in pumping mode. With a maximum elevation of 225 ft, the pool stores a maximum of 13,350 AF with a water surface area of 320 acres and a shoreline that extends 10 mi. The pool elevation affects the height to which water is pumped from the lower reservoir to the upper reservoir (i.e., Lake Oroville). All other factors held constant, the higher the Thermalito Diversion Pool, the lower the amount of electrical energy required to pump water into the lake. On the other hand, the diversion pool elevation typically has no impact on power generation.

Water releases from the diversion dam into the Feather River maintain fish habitat between the diversion dam and Thermalito Afterbay Outlet. The majority of the river water typically flows through the one-unit TDD Power Plant, which came online in 1987. Two slide gates provide isolation of the two penstocks that channel water through the powerplant turbine. This water empties directly into the Feather River via discharge lines.

The unit at the plant has a maximum output capacity of about 3 MW and a maximum turbine flow rate of 615 cfs. Turbine water flows downstream over the Fish Barrier Dam, which is located about a half a mile downstream of the TDD. Water can also be released from the diversion dam directly into the Feather River via a spillway that is controlled by 14 radial gates.

A slide gate at the diversion dam controls the amount of water that flows to the nearby Feather River Fish Hatchery (less than 1 mi downstream) for use within the hatchery itself and for the hatchery fish ladder. This water bypasses the power plant; therefore, it produces no power. Hatchery water releases join the Feather River less than a half a mile below the fish barrier.

4.3.1.3 Fish Barrier Dam and the Feather River Fish Hatchery

The Feather River Fish Barrier Dam is downstream of the TDD and immediately upstream of the Feather River Fish Hatchery. The Fish Barrier Dam diverts fish from moving upstream to the diversion dam, directing them into the Feather River Fish Hatchery via a fish ladder. Fish that are diverted when returning to spawn include salmon and steelhead. When these fish species reach spawning age, and conditions are right, they instinctively swim upstream. Because the barrier dam stretches across the river, the only upstream route is up the fish ladder.

This ladder connecting the Feather River near the fish barrier dam to the hatchery is a 2,150-ft_long, 6-ft-wide concrete structure. When used, the water flow rate in the ladder is 10 cfs with a minimum 2-ft pool depth in each stair. The maximum drop between adjacent stairs is 1 ft. Hatchery facilities also include a spawning building and rearing raceway (CDWR undated b). As described previously, the water source for fish ladder flows is the Thermalito Diversion Pool; flow is controlled by a slide gate.

4.3.1.4 Thermalito Power Canal

The Thermalito Power Canal is a concrete-lined structure extending 10,000 ft from the head-works structure at the TDD to Thermalito Forebay. The canal was constructed between 1965 and 1967. The Thermalito Diversion Pool and the Thermalito Forebay can be isolated by three radial gates at the TDD. Typically these gates are open, allowing the canal to convey water in either direction between Thermalito Diversion Dam and Thermalito Forebay. When operating in generation mode, water released through the Hyatt PGP increases the water volume stored in the Thermalito Diversion Pool, thereby raising the pool's water surface elevation. To reach equilibrium between this pool and the surface water elevation in the Thermalito Forebay, some of the water is conveyed to the forebay through the diversion dam and the canal. In generation mode, the canal's maximum flow rate is 16,900 cfs from the Thermalito Diversion Pool to the Thermalito Forebay. In pump mode, canal water flows in the opposite direction, from the Thermalito Forebay to the Thermalito Diversion Pool, and the maximum water flow rate is 9,000 cfs.

As described in more detail in Section 4.3.1.5, both water releases and pumping at the Thermalito PGP affect Thermalito Forebay water levels and therefore the elevation balance between the forebay and the diversion pool. Although water elevation differences between the afterbay and pool occur, the difference is typically minimal.

4.3.1.5 Thermalito Forebay Dam, Forebay, and PGP

Located about 4 mi west of the city of Oroville, the Thermalito Forebay was constructed between 1965 and 1968. It stores water released through the Thermalito Power Canal from the Thermalito Diversion Pool. Reservoir waters are contained by the Thermalito Forebay Dam on the south and east and by Campbell Hills on the north and west. As described in Section 4.3.1.4, the forebay conveys generating and pumping flows between the Thermalito Power Canal and the Thermalito PGP and provides regulatory storage and surge damping for the Oroville-Thermalito Complex. The reservoir has a maximum operating storage capacity of 11,770 AF. At this storage level, the water surface elevation is 225 ft and the forebay water covers 630 acres with a 10-mi shoreline.

The Thermalito Forebay Dam is a homogeneous, zoned, earth-filled structure with an embankment volume of 1,840 thousand cubic yards. The structure is 91 ft high with a crest length of 15,900 ft. The crest is 231 ft above sea level.

Thermalito PGP construction began in 1964 and was completed in 1969. The first unit came online on February 9, 1968. It currently has four units, three of which have pumping capabilities. Each of the three pumping units can generate up to 28 MW of power. The generation-only unit has a capacity of 36 MW.

Thermalito Forebay Dam releases water into the Thermalito Afterbay through the power plant flow. The penstock that channels forebay water through the generation-only turbine is 24 ft in diameter at the water intake, narrowing to 21 ft near the penstock outlet. Three penstocks convey water through the pumping-generating units: one penstock per unit. The diameter of each penstock ranges from 21 ft at intake to 18 ft near the turbine. In generation mode, the maximum turbine flow rate for the pumping-generating units is 4,200 cfs, while the generating-only unit has a maximum flow rate of 4,800 cfs.

Operations of the Thermalito PGP are coordinated with those of the Hyatt PGP. While in pumping mode, the three units with an installed pumping capacity of 120,000 hp are used to lift water from the afterbay to the forebay at a rate of 9,120 cfs. Water released for power in excess of local and downstream requirements is conserved by pump-back operation during off-peak hours through both power plants into Lake Oroville. This excess water is subsequently released for power generation during periods of peak power demand.

4.3.1.6 Thermalito Afterbay Dam and Afterbay

Constructed from 1965 to 1968, the Thermalito Afterbay Dam and the Thermalito Afterbay control water flows into the Feather River downstream of the Thermalito PGP. Water released into the Feather River is regulated by five radial gates at the dam outlet structure, situated in the southwest corner of the afterbay. The confluence of this Thermalito Afterbay outlet structure and the Feather River is about 10 river miles below the TDD.

The afterbay also provides water storage when the power plant is operating in pump mode; that is, it serves as the lower reservoir for pumping purposes. The afterbay is shallow, and the sun warms basin water for agricultural deliveries to farms located to the east.

The Thermalito Afterbay has a water storage operating capacity of 57,040 AF. When the afterbay is full, it has a water surface elevation of 136.5 ft above sea level, covering a surface area of 4,300 acres with 26 mi of shoreline.

Water in the afterbay is contained by a homogeneous, earth-filled dam comprising more than 5 million cubic yards of material. At a height of only 39 ft, it is a relatively low structure compared with other Oroville-Thermalito Complex dams. However, it is the longest dam in the complex, with a crest length of 42,000 ft. The crest elevation is 142 ft above sea level (CDWR 2012e).

4.3.1.7 Feather River Flow and Reservoir Operating Criteria

Water releases from Oroville-Thermalito Complex operations must comply with a set of Feather River in-stream flow and temperature requirements that vary according to season and hydrology. For regulatory purposes, the Feather River in the Oroville-Thermalito Complex footprint is separated into a low-flow channel (LFC) and a high-flow channel (HFC). As illustrated in Figure 4-2, the LFC reach begins just below the TDD and extends downstream to the point where the Thermalito Afterbay releases water into the river. The HFC begins at that confluence and extends downstream.

Under current operating conditions, the flow rate in the LFC must be at a minimum of 600 cfs. This LFC minimum may change upon the issuance of the new Federal Energy Regulatory Commission license for Lake Oroville. To ensure compliance, the CDWR operates the system so that a slightly higher flow rate is typically attained. Note that this flow rate is roughly the maximum turbine release rate of the 3-MW unit at the TDD Power Plant. Because the flow requirement applies year-round, the power plant is essentially base loaded whenever it operates.

Under flooding conditions, LFC flows are determined on a case-by-case basis. Incremental water releases into the Feather River above the generator's production capacity or beyond the maximum turbine flow rate do not produce additional power. Therefore, there is an economic incentive to divert water down the Thermalito Power Canal into the Thermalito Forebay for power production at the Thermalito PGP. In addition, the Thermalito PGP has a higher dynamic head and therefore typically produces more energy per unit of water released through its turbines than the TDD Power Plant.

In addition to the minimum flow requirement, environmental regulations limit the water flow rate change in the LFC between two consecutive days. When flows are between 600 and 2,500 cfs, the maximum decrease allowed in daily average flows is 300 cfs per day. This maximum daily decrease requirement is 500 cfs per day when flows are between 2,501 and 3,500 cfs. The requirement is further relaxed to 1,000 cfs per day when flows are between 3,501 and 5,000 cfs.

Minimum water flow requirements for the HFC are more complicated. The year is divided into three periods. Period 1 is from October through February; Period 2 is March; and Period 3 is from April through September. The HFC flow is dependent on forecasted elevation and the previous year's April through July unimpaired runoff into Lake Oroville. For this paper, we classify the hydrological states as dry, moderate, and wet, as defined below.

- **Dry**—Lake Oroville's surface water elevation is projected to be less than 733 ft (about 1.5 million AF), and the previous year's April through July unimpaired runoff into Lake Oroville was less than 55% of the historical mean (1,942 thousand acre feet [TAF]), or 1,068 TAF.

- **Moderate—**Lake Oroville's surface water elevation is projected to be greater than 733 ft, and the previous year's April through July unimpaired runoff into Lake Oroville was less than 55% of the historical mean (1,942 TAF), or 1,068 TAF.

- **Wet—**Lake Oroville surface water elevation is projected to be greater than 733 ft, and the previous year's April through July unimpaired runoff into Lake Oroville was greater than 55% of the historical mean (1,942 TAF), or 1,068 TAF.

Table 4-1 lists minimum HFC flow requirements as a function of hydrological classification and time period. When water flows in the HFC are less than 2,500 cfs, channel flows are not permitted to decrease by more than 200 cfs in any 24-hour rolling period.

Further downstream of the Oroville-Thermalito Complex, maximum allowable flows have been established for flood control purposes. The Feather River flows above the Yuba River confluence must not exceed 180,000 cfs; downstream of the Yuba River confluence, the maximum allowable flow rate is 300,000 cfs; and at the confluence with the Bear River, the maximum allowable flow rate is 320,000 cfs.

Operational limitations have also been established for the Oroville-Thermalito Complex reservoirs. As listed in Table 4-2, Lake Oroville has a wide range of reservoir operations. Because Lake Oroville has a large surface area and storage capacity (i.e., maximum operating storage of 3,537.6 TAF), changes in reservoir water surface elevations during day-ahead and real-time operations are small.

TABLE 4-1 Feather River HFC Minimum Flow Requirements

Hydrologic Classification	Time Period	Minimum Flow Requirements (cfs)
Wet	October through February	1,700
	March	1,700
	April through September	1,000
Moderate	October through February	1,200
	March	1,000
	April through September	1,000
Dry[a]	October through February	900–1,200
	March	750–1,000
	April through September	750–1,000

[a] These flow requirements reflect a 25% reduction under certain conditions.

TABLE 4-2 Oroville-Thermalito Complex Reservoir Operating Criteria

Reservoir	Normal Minimum Elevation (ft)	Normal Maximum Elevation (ft)	Absolute Maximum Elevation (ft)
Lake Oroville	640.0	900.0	901.0
Thermalito Diversion Pool	221.0	225.0	225.0
Thermalito Forebay	221.0	225.0	225.0
Thermalito Afterbay	124.0	136.0	136.5

However, the water elevation in the Thermalito Diversion Pool and Thermalito Forebay change comparatively quickly during the course of a day. Under typical operating conditions, these waters—along with those in the power canal—are essentially the same body; therefore, water elevations move up and down in tandem. The combined maximum operating storage of Thermalito Diversion Pool and Thermalito Forebay is 25.2 TAF: roughly 0.74% of Lake Oroville's maximum operating storage.

Additional storage requirements for the Thermalito Afterbay have been established to maintain brood ponds and enhance habitats for nesting waterfowl, including grebes (freshwater diving birds) and giant garter snakes. Reservoir operational requirements change from year to year based on hydrologic conditions; however, typical operational requirements are provided below.

- For waterfowl nesting, a water elevation of 133.5 ft must be reached every 9 days from March 15 through May 31, without exceeding 134 ft.

- For grebe nesting, water elevations must be within a range of 132 to 135 t from the beginning of July through mid-August. This restriction is relaxed to a range of 131 to 136 ft from mid-August through mid-September.

- For the giant garter snake, a water elevation of 133.5 ft must be maintained for a continuous 12-hour period at least once a month from May 1 through October.

As shown in Figure 4-3, Lake Oroville goes through an annual fill-and-drain cycle. On average, the lake is at an annual low point in October. This void allows the reservoir to capture and store winter runoff and spring snowmelt. Electricity demand during this period is typically less than in the summertime, so outflows for power production tend to be lower.

Typically, the reservoir is at its peak storage level in May. After May and through September, precipitation and snowmelt levels are significantly reduced. At the same time, the demand for electricity increases. The result is reservoir draw-down during this period, because water releases for power production significantly exceed reservoir inflows. On average, the annual change in storage throughout the course of a year is about 880 TAF. Data plotted in

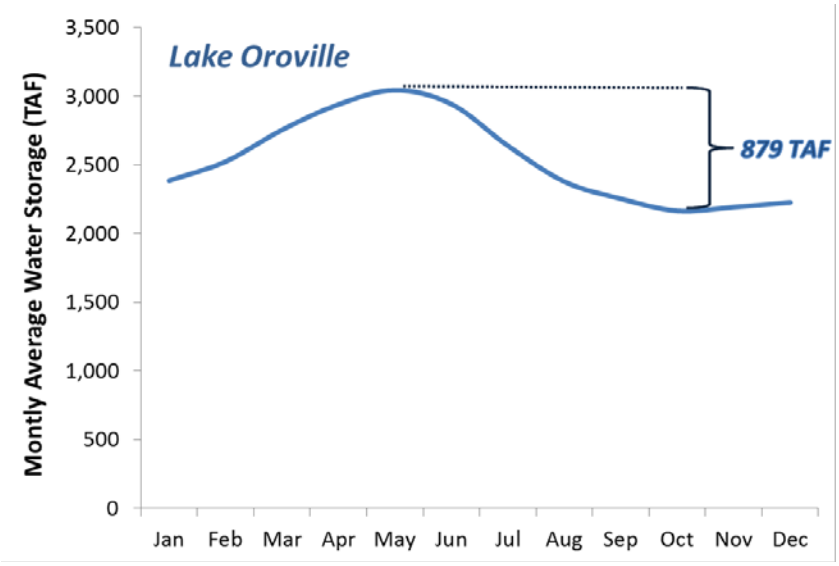

FIGURE 4-3 Historical Monthly Average Water Storage Levels for Lake Oroville

Figures 4-3 and 4-4 were obtained from the CDWR Data Exchange Center. Average values are based on data for 1969 through 2005 for Lake Oroville (CDWR 2010a) and on data from 1968 through 2005 for the Thermalito Afterbay (CDWR 2010b).

As shown in Figure 4-4, monthly average storage at the Thermalito facilities is much smaller and displays significantly less fluctuation throughout the year. As discussed above, some of the storage fluctuations in the Thermalito Afterbay result from environmental operating criteria imposed on reservoir operations.

4.3.1.8 Oroville-Thermalito Complex Resource Overview

This section provides an overview of the Oroville-Thermalito Complex and its components. As illustrated by the blue dots in Figure 4-5, Lake Oroville's maximum operating storage dwarfs those of the other water storage facilities. There is also a very large vertical drop in elevation between the Oroville Dam and the TDD. The TDD and the Thermalito Forebay Dam are at essentially the same elevation and have identical maximum operating storage elevations. These two reservoirs also have roughly the same water storage volumes. At the bottom of the cascade is the Thermalito Afterbay. The drop in elevation from the Thermalito Forebay to the Afterbay is on the same order of magnitude as the drop from the TDD to the Fish Barrier Dam, which has an elevation of 148 ft (Topozone 2012).

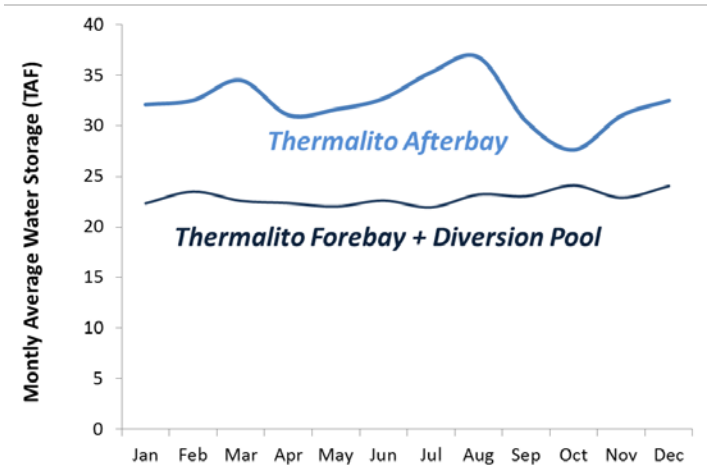

FIGURE 4-4 Historical Monthly Average Water Storage Levels for Thermalito Reservoirs

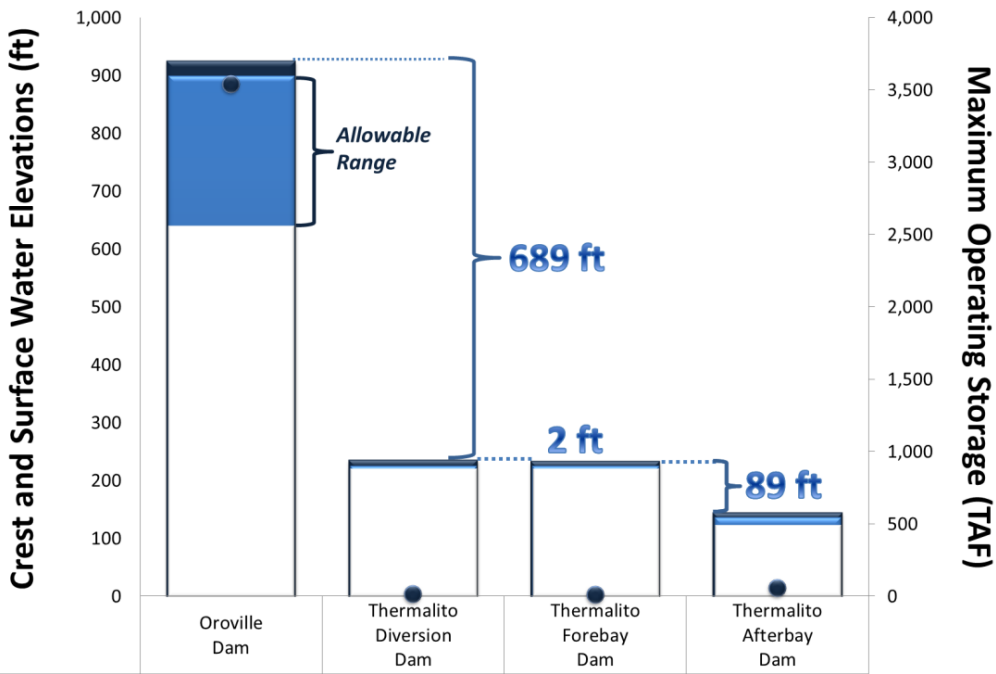

	Oroville Dam	Thermalito Diversion Dam	Thermalito Forebay Dam	Thermalito Afterbay Dam
Crest Elevation (ft)	922	233	231	142
Pool Elevation @ mos (ft)	900	225	225	136.5
Normal Minimum Elevation (ft)	640	221	221	124
Max Operating Storage (mos) (TAF)	3,537.6	13.4	11.8	57.0

FIGURE 4-5 Elevation Characteristics of the Oroville-Thermalito Complex Reservoirs

Although Lake Oroville dominates the system in terms of its size, operations at the Lake Oroville Dam must be sensitive to operational restrictions at the lower reservoirs and comply with Feather River flow requirements. This situation highlights the fact that system operations are tightly coupled; the actions taken at one facility impact one or more of the other parts of the system. For example, the amount of water released from Lake Oroville impacts water flows throughout the rest of the complex. Therefore, operations at the Hyatt and Thermalito PGPs have historically been conducted in tandem.

These system interdependencies occur not only over space, but also over time. For example, when water is released from a reservoir, the water is not available for release at some other point in time. A water release also lowers the reservoir elevation, thereby reducing the conversion of water releases into electricity by the power plant at the dam. Depending on the situation, this loss may persist over a long period of time.

Not only Lake Oroville dominate the system in terms of water, but the Hyatt PGP produced the majority of the system's power production and accounts for most of its capacity. The average annual generation shown in Figure 4-6 is based on data collected and maintained by the U.S. Energy Information Administration (EIA) (Forms EIA-423 and EIA-906) for the 1985 through 2005 time period.

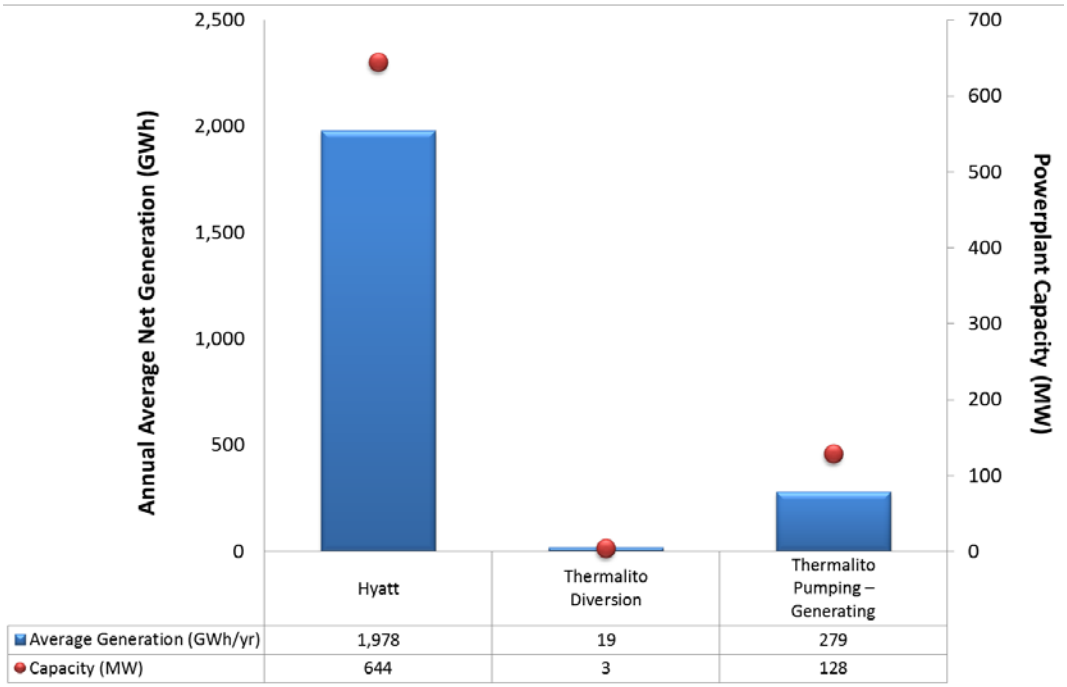

FIGURE 4-6 Oroville-Thermalito Complex Powerplant Capacities and Historical Generation Levels

Electricity produced by hydropower plants, including those in the Oroville-Thermalito Complex, is primarily a function of head (which is dictated primarily by reservoir elevation) and turbine flow rate. Therefore, water operations in the Oroville-Thermalito Complex drive hydropower production.

It should be noted, however, that water is not the only operational driver for electricity production. The demand for electricity, coupled with market forces (i.e., prices) for energy and ancillary services, significantly influence temporal electricity production patterns. Among many factors, the physical limitations and performance characteristics of power equipment, along with power dispatch goals and guidelines, further complicate system dynamics.

4.3.2 Colorado River Storage Project Demonstration Site

The second set of demonstration sites to which the WUOT was applied includes the Colorado River Storage Project with a focus on the Aspinall Cascade, a portion of the Salt Lake City Area Integrated Projects (SLCA/IP), which consists of 11 hydropower plants with a combined installed capacity of approximately 1,820 MW. These resources help provide electricity to cities and towns, rural electric cooperatives, agricultural irrigation districts, and Federal and state agencies. Capacity and energy are sold on the wholesale market under long-term firm (LTF) contracts. The marketing of SLCA/IP is currently conducted under the auspices of the Western Area Power Administration (Western). The Colorado River Storage Project (CRSP) Management Center, headquartered in Salt Lake City, Utah, sells the hydropower capacity and energy at LTF rates. The CRSP comprises the largest share of SLCA/IP resources. When energy exceeds LTF contractual obligations or when operational regulations result in generation levels above load, energy is sold on the spot market in order to maximize the value of the hydropower resource. The benefits of SLCA/IP include serving the energy requirements of special project uses, such as pumping for irrigation, and fulfilling utility system requirements for spinning reserves and regulation services.

The CRSP was authorized by a special Congressional Act on April 11, 1956, to develop the water resources of the Upper Colorado River Basin and control a drainage basin of approximately 108,335 square miles. The project regulates the flow of the Colorado River such that water-use developments in the Upper Colorado River Basin can take place while minimum water deliveries to the Lower Basin are maintained as mandated by the Colorado River Compact. Benefits provided by the CRSP include flood control, irrigation, recreation, municipal and industrial water supply, fish and wildlife conservation, electric power capacity, and electricity generation.

A map of SLCA/IP power plants is shown in Figure 4-7. CRSP resources consist of the Glen Canyon Dam upstream of the Grand Canyon National Park; Flaming Gorge Dam on the Green River in Utah near the Wyoming border; Navajo Dam on the San Juan River in New Mexico near the Colorado border; and the Wayne N. Aspinall Dams (formerly Curecanti) on the Gunnison River in west central Colorado. The power plants associated with Aspinall are Blue Mesa, Morrow Point, and Crystal. Glen Canyon Dam accounts for about three-fourths of the

FIGURE 4-7 SLCA/IP Power Plants

CRSP's total nameplate capacity. Although these resources are marketed by Western, the Bureau of Reclamation (Reclamation) owns and operates CRSP hydropower resources.

The Collbran Project, located in west central Colorado about 35 mi northeast of Grand Junction, was authorized by Congress in July of 1952. It developed a major part of the water in Plateau Creek and its principal tributaries. Major project works include Vega Dam and Reservoir, two power plants, two major diversion dams, about 37 mi of canal, and about 18 mi of pipeline and penstock. East Fork Diversion Dam and Feeder Canal, along with the Bonham-Cottonwood Collection System, carry water to Bonham Reservoir, which supplies water to operate the Molina power plants.

The Rio Grande Project, which is 125 mi north of El Paso, Texas, began operation in 1916 after the Rio Grande Reclamation Project congressional act in 1905 established a much-

needed irrigation project on the Rio Grande River in south central New Mexico and west Texas. The only dam with a power plant at the Rio Grande Project is Elephant Butte Dam.

The Seedskadee Project, a participating project of the SLCA/IP, is in the Upper Green River Basin in southwestern Wyoming. It provides storage and regulation of the flows of the Green River for power generation, municipal and industrial use, fish and wildlife conservation, and recreation. The Fontenelle Dam is the only power plant associated with the Seedskadee Project.

The Dolores Project is located in the San Juan and Dolores River basins of the Upper Colorado River Basin in southwestern Colorado. It extends through portions of Montezuma and Dolores counties and uses water from the Dolores River for irrigation, municipal and industrial use, recreation, fish and wildlife conservation, and production of hydroelectric power. There are hydroelectric power plants at the McPhee Dam and the Towaoc Canal.

4.3.2.1 Operating Constraints and Operational Guidelines

The SLCA/IP projects power plant operations are constrained by both physical limitations and institutional agreements that are designed to help protect the environment. Operations must also comply with laws that allocate water among various entities within the Colorado River Basin. Monthly water releases in the Upper and Lower Colorado River Basin are set by Reclamation to be consistent with various operating rules and guidelines, acts, international water treaties, consumption use requirements, State agreements, and other regulatory guidelines collectively known as the “Law of the River.” In addition to power production, monthly release volumes are set considering other uses of the reservoirs, such as for flood control, river regulation, consumptive uses, water quality control, recreation, and fish and wildlife enhancement and to address other environmental factors. One requirement is that a minimum annual average of 8.23 million acre-feet (MAF) of water must be released from Glen Canyon Dam.

4.3.2.2 Glen Canyon Dam

Glen Canyon Dam was built by Reclamation between 1956 and 1964. It is a 710-ft-high concrete arch structure with a crest length of 1,560 ft. The total capacity of its reservoir, Lake Powell, is 27.0 MAF, with an active capacity of approximately 20.9 MAF. The plant has eight turbines that are arranged in pairs such that two generators share one step-up transformer. The combined capacity of these eight generators is approximately 1,320 MW at unity power factor.

The Glen Canyon Dam is restricted by operating criteria to temper water release variability. On October 9, 1996, Bruce Babbitt, then-Secretary of the U.S. Department of the Interior, signed the Record of Decision (ROD) on operating criteria for the Glen Canyon Dam. Flow restrictions under the ROD, along with operational limits in effect prior to June 1, 1991, are shown in Table 4-3. The ROD criteria require water release rates to be 8,000 cfs or greater

TABLE 4-3 Operating Constraints Prior to 1991 and under the ROD (Reclamation 1996)

Operational Constraint	Historic Flows (Pre-1991)	ROD Flows (Post 1997)
Minimum release (cfs)	3,000 summer 1,000 rest of year	8,000 - 7 am - 7 pm 5,000 at night
Maximum release (cfs)	31,500	25,000
Daily fluctuations (cfs/24 hrs)	28,500 summer 30,500 rest of year	5,000; 6,000; or 8,000 depending on release volume
Ramp rate (cfs/hr)	Unrestricted	4,000 up 1,500 down

Source: Reclamation (1996).

between the hours of 7:00 a.m. and 7:00 p.m., and at least 5,000 cfs at night. The criteria also limit how quickly the release rate can be increased and decreased in consecutive hours. The hourly maximum increase (i.e., the up-ramp rate) is 4,000 cfs/hr, and the hourly maximum decrease (i.e., the down-ramp rate) is 1,500 cfs/hr. ROD operating criteria also restrict how much the releases can fluctuate during rolling 24-hr periods. This change constraint varies between 5,000 cfs and 8,000 cfs per day, depending on the monthly water release volume. Daily fluctuation is limited to 5,000 cfs in months when less than 600 TAF are released. The limit increases to 6,000 cfs when monthly release volumes are between 600 TAF and 800 TAF. When the monthly water release volume is 800 TAF or higher, the daily allowable fluctuation is 8,000 cfs.

The maximum flow rate is limited to 25,000 cfs under the ROD operating criteria. Maximum flow rate exceptions are allowed to avoid spills or flood releases during high runoff periods. Under very wet hydrological conditions, defined as when the average monthly release rate is greater than 25,000 cfs, the flow rate may be exceeded, but water must be released at a constant rate.

4.3.2.3 Flaming Gorge and Fontenelle Dams

Flaming Gorge is located on the Green River in northeastern Utah, about 32 mi downstream from the Utah/Wyoming border. Flaming Gorge Reservoir has a total capacity of approximately 3.8 MAF at a reservoir water elevation of 6,040 ft. The reservoir has an active capacity of 3,515,700 AF and a surface area of 42,020 acres.

The power plant has three generating units. Each unit has a capacity of 48 MW for a plant total of approximately 144 MW. However, because of turbine limitations, the operable capability of the power plant is approximately 141 MW. On average, the Flaming Gorge Dam Power Plant generates about 491.2 GWh of electricity annually.

In order to provide the warmer water needed to increase the growth rate of downstream trout, water releases are made through a selective withdrawal structure. The structure consists of a set of interlocking panels that can be manually raised to any height above the penstock intake to within 40 ft of the water surface. Around April first of each year, the upper gates are raised to an elevation about 40 ft below the surface of the reservoir. In order to maintain tailwater trout fisheries, the hourly average mandatory minimum flow rate from the reservoir is 800 cfs.

To protect endangered native fish in the Green River basin, the U.S. Fish and Wildlife Service released a biological opinion designed to protect the Colorado Squawfish and Razorback Sucker. The overall intent of the opinion is to structure releases from the Flaming Gorge reservoir so that they resemble natural hydrograph and water temperature conditions.

High spring release volumes are being structured to enhance river flows during spring spawning periods. Flow volumes include full power plant output (approximately 4,500 cfs or 141 MW) to full plant output plus bypass tubes and spillways combined. The high release period could be held for as long as 4 weeks. The actual volume and duration of release from Flaming Gorge is determined by the river volume desired on the Green River below the Jensen measurement gauge and inflow support from the Yampa River. The high spring release is patterned around the peak runoff period, which varies each year based on weather conditions and endangered fish spawning activities. High peak spring releases transfer water that historically was released during summer and winter peak months to spring months when both capacity and energy are not as valuable. This approach requires higher off-peak releases and contributes to higher project energy purchase costs.

After the high spring release is completed, flows on the Green River below the Jensen measurement gauge are held to an average of 1,600 cfs, if possible. The gauge is about 93 mi downstream of the reservoir, receiving water not only from Flaming Gorge but also from the uncontrolled Yampa River, which joins the Green River approximately 65 mi downstream of the dam. This means that releases from the Flaming Gorge power plant are held to a minimum of 800 cfs (25 MW) until a time when the combined average flows on the Green and Yampa rivers below Jensen fall below the desired level of 1,600 cfs average. This power plant restriction can last from 3 to 8 weeks, depending on snowpack conditions. Lost capacity at Flaming Gorge during this time is approximately 115 MW.

When the average flow on the Green River below the Jensen gauge falls below 1,600 cfs, releases from the Flaming Gorge power plant can be increased to a point where the Green River does not fluctuate between a limit of 1,400 and 1,800 cfs. The intent is for Flaming Gorge releases to moderate variations in Jensen Gauge flows to a stage change limit of 0.1 m/day.

Western uses a river flow attenuation model to model river flows at different generation levels to maximize power plant peaking capability within the river fluctuation limitations. CRSP schedulers coordinate with Reclamation to pattern Flaming Gorge generation to fit system requirements. Reclamation verifies that Western's proposed pattern meets operating requirements then notifies the Flaming Gorge Working Group of the pattern with an official directive for a flow change. Western does not make adjustments to the Flaming Gorge schedule without prior coordination with the Bureau of Reclamation. At times Western will carry operating reserves at Flaming Gorge when capacity is needed above the availability of Glen Canyon, Morrow Point, or Blue Mesa.

4.3.2.4 Gunnison Basin and the Aspinall Cascade

The Aspinall Cascade is located along a 40-mi section of the Gunnison River between the towns of Gunnison and Montrose, Colorado. As shown in Figure 4-8, from near the town of Gunnison water flows in the main stem of the Gunnison River flow through the Aspinall Cascade through the Black Canyon National Park. Further downstream, these waters link up to

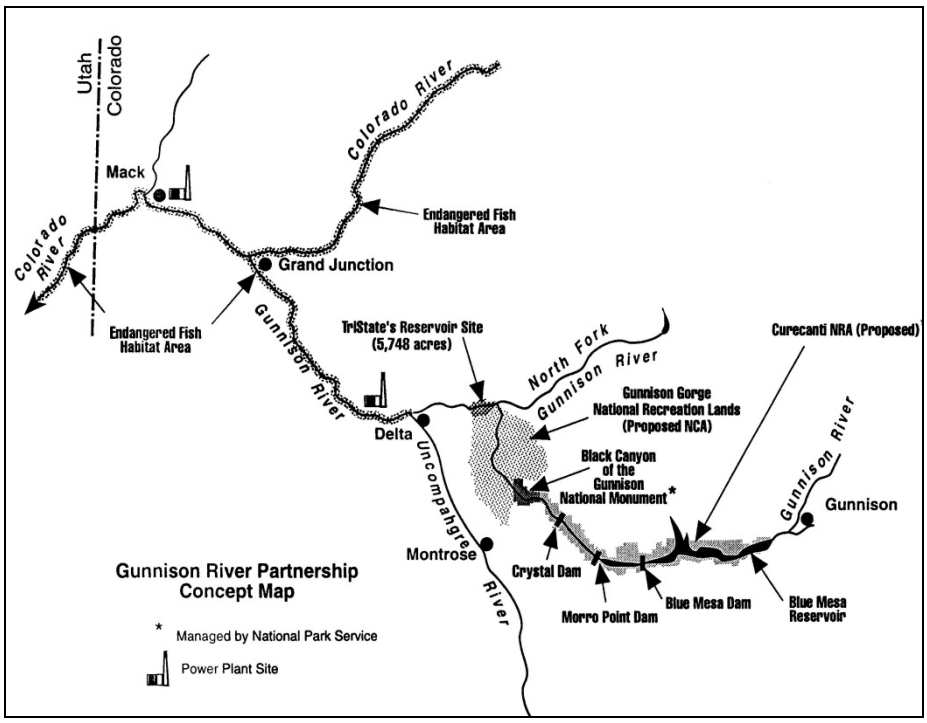

FIGURE 4-8 Map of the Aspinall Cascade and Surrounding Area

several tributaries including the Gunnison North Fork and the Uncompahgre River before linking up to the Colorado River at the city of Grand Junction, Colorado.

Water storage capacities of these three reservoirs in terms in TAF and power plant generating nameplate capacity in MW are shown in Figure 4-9.

The Aspinall Cascade is operated as a tightly coupled multipurpose system. Its primary purpose is to furnish the long-term regulatory storage needed to permit states in the upper basin to meet their flow obligation at Lees Ferry, Arizona, as defined in the Colorado River Compact. Operation of the Aspinall Cascade considers power generation, projected inflows to its reservoirs, flood control needs, existing water rights, minimum in-stream flows, target elevations for reservoirs, flow needs for endangered fish and other resources, recreation, hydropower needs, and other factors. Table 4-4 contains a summary of reservoir, dam, and power plant characteristics in the Aspinall Cascade.

The Blue Mesa Dam is on the Gunnison River about 30 mi downriver of Gunnison, Colorado. The dam is a zoned earth-fill embankment with a structural height of 390 ft and a crest length of 785 ft. It contains 3,093,000 yd^3 of materials. The maximum water discharge of the spillway is 34,000 cfs.

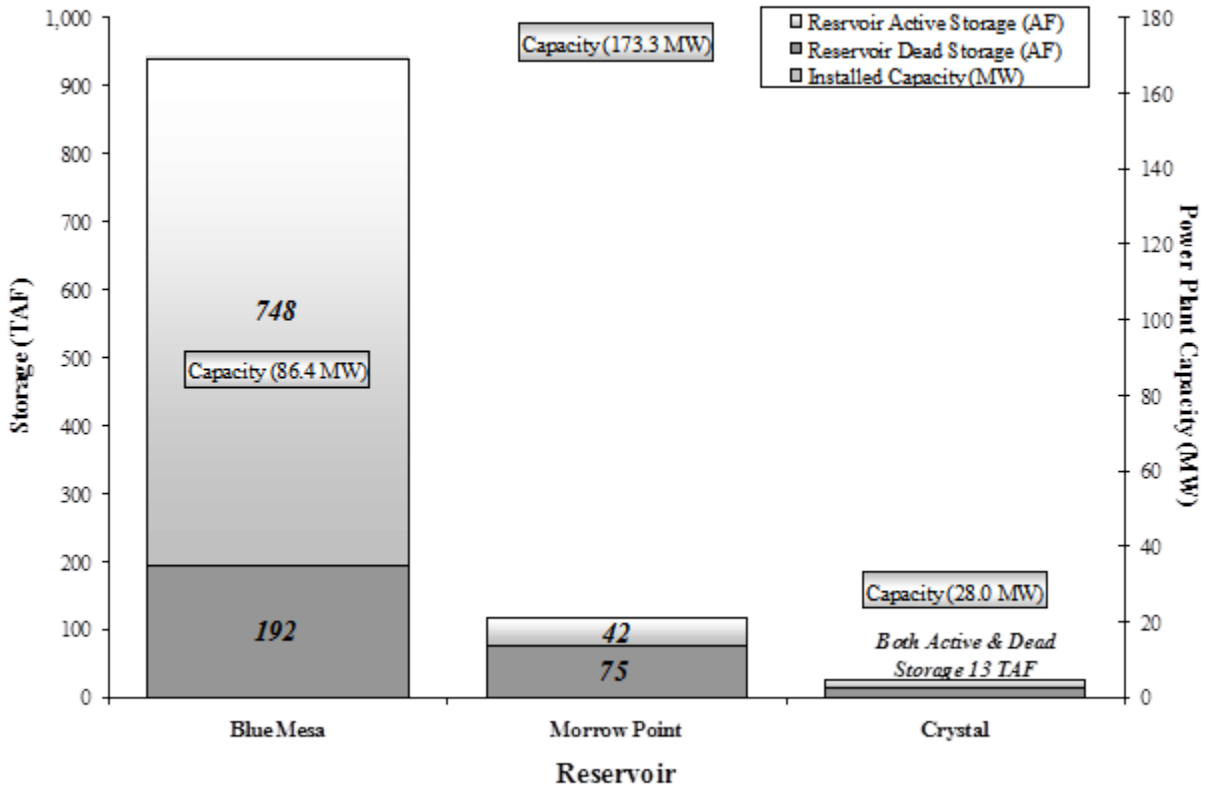

FIGURE 4-9 Aspinall Cascade Reservoir Storage Capacities and Power Plant Generating Capacities

TABLE 4-4 Characteristics of Aspinall Cascade Resources

Dam, Reservoir, and Power Plant Characteristics	Blue Mesa	Morrow Point	Crystal
Dam Type	Earthfill embankment	Double-curvature thin-arch	Double-curvature thin-arch
Primary Purpose	Water storage	Power production	Flow regulation
Dam Height (ft)	390	468	323
Spillway Crest Elevation (ft)	7,487	7,123	6,756
Crest Elevation (ft)	7,528	7,165	6,772
Active Reservoir Capacity (AF)	748,500	42,120	13,000
Surface Area (acres)	9,180	817	340
Power Plant In Service Year	1967	1970	1978
Installed Capacity (MW)	86.4	173.334	28.0
Number of Turbines	2	2	1
Production Mode	Peaking	Peaking	Base Load
Capacity Factor in FY 2001 (%)	22.54	17.98	62.14
Net Generation in FY 2001 (GWh)	205.322	271.157	151.270
Maximum Annual Generation 1992–2001 (GWh)	372	517	218
Minimum Annual Generation 1992–2001 (GWh)	205	271	151
Spinning Reserve	Yes	Yes	Yes[a]
Non-Spinning Reserve	Yes	Yes	Yes[a]
Replacement Reserve	Yes	Yes	Yes[a]
Regulation/Load Following	Yes	Yes	Yes[a]
Black Start			
Voltage Support			

[a] The Crystal Power Plant is physically capable of providing these ancillary services but institutional and environmental constraints preclude Crystal from operating in a mode such that these services can be sold on the market.

Blue Mesa's primary purpose is water storage. The reservoir has a total capacity of 940,800 AF. Live capacity is 829,500 AF, of which 748,400 AF is active capacity and is 81,100 AF is inactive capacity. Dead storage is 111,200 AF. At maximum water surface elevation, the reservoir occupies 9,180 acres.

Power generation from Blue Mesa using a single 30-MW generator was initiated in September 1967. Two months later, a second 30-MW unit was put into service. Both generators were uprated to 43.2 MW in 1988. Two 41,500-hp turbines drive the generators. The plant capacity factor in water year (WY) 2002 was 22.54%, generating 169-GWh of electricity. The power plant operates in a peaking mode with large hourly fluctuations in power production over the course of a day.

The Blue Mesa Dam and hydropower plant is at the top (i.e., highest elevation level) of the cascade, followed by Morrow Point and then Crystal. The Blue Mesa reservoir has the largest storage capacity, holding approximately 940.8 TAF of water. That is more than eight times larger than the Morrow Point reservoir and more than 36 times larger than the Crystal reservoir. The Aspinall Cascade is operated as a tightly coupled, multipurpose system. Its primary purpose is to furnish the long-term regulatory storage needed for states in the upper Colorado River Basin to meet the flow obligation at Lees Ferry, Arizona, as defined in the Colorado River Compact. Operation of the Aspinall Cascade takes into consideration power generation, projected reservoir inflows, flood control needs, existing water rights, minimum in-stream flows, reservoir target elevations, flow needs for endangered fish and other resources, recreation, and other factors.

The total power plant capacity of the cascade is 283.4 MW. At the top of the cascade, the Blue Mesa power plant has two generators, the total capacity of which is 86.4 MW. It operates in a peaking mode, with large hourly fluctuations in power production over the course of a day; from a starting point of zero, it has the potential to reach its maximum capacity in 1 hour.

The Morrow Point Power Plant has the largest generating capacity in the cascade. Its two units have a combined capacity of 165 MW. The plant is located 12 mi downstream from the Blue Mesa Dam.

Hydropower plant output levels can be ramped up or down from zero production levels to maximum capability in a matter of minutes without adverse effects on the power equipment. This attribute makes it well suited to providing the interconnected grid with various ancillary services, such as spinning and non-spinning reserves, regulation, and voltage support.

When a tour boat is operating in the Morrow Point Reservoir, there is a minimum reservoir elevation requirement of 7,151 ft. In addition, the Morrow Point Reservoir cannot be drawn down by more than 3 ft per rolling 24-hour period if the surface elevation is below an elevation of 7,144 ft. When the tour boat is not operating, the minimum reservoir elevation is 7,125 ft.

The Crystal Power Plant is located 6 mi downstream from the Morrow Point Dam and approximately 20 mi east of Montrose, Colorado. In addition to functioning as a power

generation unit, its operations stabilize the flow of water through Gunnison National Park. The dam is a double-curvature, thin-arch type that is 323 ft high, with a crest length of 635 ft. The plant currently has an installed capacity of approximately 32 MW from one unit driven by a 39,000-hp hydraulic turbine. Although the Crystal Power Plant has the physical capability to provide all types of ancillary services, a flat flow requirement precludes it from rapidly changing power output from one hour to the next.

Operations at the Crystal reservoir are season specific. During the wet season (March 1 to June 30), the operating criteria limit reservoir drawdown to no more than 4 ft in a 24-hour rolling period, 5 ft over a 48-hour rolling period, or 6 ft over a 72-hour rolling period, and so on. In addition, once an elevation of 6,748 ft is reached, the reservoir level may not be dropped by more than 0.5 ft per 24-hour rolling period. During the remainder of the year, there is a 10-ft-per-24-hr-period fluctuation limit, with a maximum 3-day drawdown of 15 ft. If the reservoir elevation is below 6,733 ft, there is a 5-ft-per-24-hr-period drawdown limit, with a maximum drawdown of 20 ft per 7-day period.

4.3.2.5 Collbran Project

The Collbran project has two power plants: Upper and Lower Molina. The Bonham Reservoir acts as a forebay for the Upper Molina Power Plant, which controls releases up to a maximum capacity of 50 cfs from the reservoir. The power plant consists of a single 8.6-MW generating unit. Releases from Upper Molina feed into the Lower Molina penstock, which extends 4.7 mi via the Molina Equalizing Reservoir to the Lower Molina Power Plant. The penstock consists of steel pipe that has a maximum capacity of 50 cfs. The Lower Molina Power Plant's capacity is 4.86 MW.

4.3.2.6 Rio Grande Project

Elephant Butte Dam can store approximately 2 MAF of water for irrigation and power generation. The concrete gravity dam is 301 ft high and 1,674 ft long. The power plant has a capacity of approximately 28 MW. Power production is restricted during non-irrigation months; that is, generation during the winter months is typically zero.

4.3.2.7 Small Projects

All other SLCA/IP hydropower plants, including those in the Seedskadee, Dolores, and Provo River Projects, have small power plants. These include Fontenelle with an installed capacity of 12 MW, Deer Creek with a capacity of just under 5 MW, McPhee with a capacity of about 1.3 MW, and Towaoc with a capacity of about 12 MW.

4.3.2.8 Operational Business Process

Western's Energy Management and Marketing Office (EMMO) in Montrose, Colorado, is responsible for forecasting and balancing firm loads and hydro generation for the SLCA/IP. Office staff schedule energy needs by planning resources for the next day or longer, including preparing seasonal and monthly outlooks. Every hour of every day, marketers perform a balancing act—buying and selling energy to ensure that customers have the energy they need when they need it. If marketers project long-term shortages, they may pursue a futures contract—buying a quantity of energy for up to 1 year in the future.

4.3.2.9 Long-Term Planning and Purchases

Reclamation provides an Annual Operating Plan (AOP) to Western for Colorado River hydro facilities. The AOP projects water releases for the upcoming 24 months based on historical hydrology and projected inflow to reservoirs, as well as on downstream demands. The AOP is updated monthly through the Bureau of Reclamation's 24-Month Study.

The CRSP Management Center (CRSP MC) in Salt Lake City, Utah, uses the 24-month study results and models hourly operations for a 1-year period with the Generation and Transmission Maximization (GTMax) software. Projected hourly generation and purchases are to be used for planning purposes. The model uses existing environmental restrictions and operating guidelines to determine the best hourly generation pattern. Customer loads and system losses are forecast to determine what shortages or surpluses are expected for the next 12 months. These results from the resource modeling are forwarded to the Montrose EMMO.

At the EMMO, the CRSP Project Manager reviews the resource model output, makes modifications as necessary, and determines how CRSP generation shortages would best be covered. Using forward price projections, the CRSP Project Manager determines whether to purchase on a seasonal or monthly basis. These decisions are made in consultation with the Montrose EMMO Manager and the CRSP MC Manager.

If it is determined that some purchases should be made, the CRSP Project Manager requests offers from suppliers for a specific amount of energy to be delivered at various CRSP transmission system delivery points. Once all of the offers have been received and analyzed, a contract is entered into using the existing Western Systems Power Pool (WSPP) agreement.

When the GTMax model shows that there will be additional hydropower generation available above the sustainable hydropower (SHP) level, the CRSP MC and Montrose EMMO will determine whether the amount is substantial enough to offer available hydropower (AHP) to the CRSP customers on a monthly basis. Small volumes of excess energy, or excess energy that is only available for time periods of less than a full month, are typically sold to the market on a day-ahead or real-time basis.

On a month-ahead basis, the CRSP Project Manager creates a Load and Resource (L&R) spreadsheet, which is used on a daily basis by the Preschedulers to compare hourly generation values against hourly load requirements. This comparison allows them to determine whether purchases or sales are needed. When developing the monthly L&R spreadsheet, the CRSP Project Manager uses the same customer load projection that was used in the GTMax model, patterns the hydro generation based on the GTMax output, and modifies information on unit outages if updated or on water releases if changed. At this point, the Project Manager determines whether to purchase any shortages remaining for the month or assign that determination to daily and real-time operations. Daily communication between the preschedule and real-time personnel provide information used to determine how best to cover any shortages.

4.3.2.10 Daily Scheduling

On a 2-day-ahead basis, the CRSP firm electric service customers submit their hourly schedules using the EMMO Web Scheduler System; the schedules are then downloaded into their respective e-tag templates in TIGER, an internal Energy Management System. From TIGER, the EMMO Prescheduler submits the e-tags through an associated Open Access Technology International's Energy Trading System (ETS) e-tagging software program. The L&R spreadsheet works in conjunction with TIGER so that hourly firm electric service loads and hourly prescheduled losses and interchange transactions can be calculated in the spreadsheet. Hydrogeneration and contractual purchases are also included in the L&R spreadsheet so that they can be netted against the load obligation to create an hourly "bottom line." The Prescheduler uses the L&R spreadsheet to adjust hourly generation values to best meet the actual hourly load requirements.

At the 1-day-ahead point, the Preschedulers make final adjustments to the generation patterns and purchase or sell so that the loads and resources are balanced close to zero. Preschedulers bring the bottom line within a manageable level (typically ±50 MW, depending on availability) so that the next day, the CRSP Real-time Desk can purchase or sell each hour to reach a net of zero. Prescheduling staff base their marketing decisions on the latest weather forecasts, available generation, water elevations in the reservoirs, availability of market energy on a real-time basis, and transmission system constraints.

4.3.2.11 Real-Time Scheduling

Once the Preschedulers have optimized resources to best meet hourly load obligations, responsibility shifts to the CRSP Real-time Desk. The CRSP Real-time Desk operator is responsible for balancing CRSP loads and resources by dispatching generation or purchasing or selling in the spot market. CRSP loads and resources may change on an hourly basis as a result of generation changes, transmission outages, unscheduled flow curtailments, Project Use load requirements, emergency responses, and/or CRSP firm electric service customers requesting additional energy through the Western Replacement Power (WRP) program. Hourly purchases, sales, load schedules, and generation resources are managed to balance commitments within the

Western Area Colorado Missouri and Western Desert Southwest (DSW) Customer Service Region in Lower Colorado (WALC) Balancing Authorities (BAs). Transmission system reservations and e-tag creation and profile changes are made as necessary. These activities are all performed each hour, ending on the half hour and culminating in a check-out process with the BAs' transmission operators.

In addition, the CRSP Real-time Desk operator monitors reservoir elevations, monitors operating contingency reserve requirements, finds alternative sources of energy during generation unit outages, and reroutes firm electric service schedules during transmission outages. While real-time operations are usually the end point for scheduling and balancing, all of the data from preschedule and real-time operations are checked and corrected by the Postschedule accounting staff.

A checkout process is carried out between Postschedule staff and all companies that bought or sold generation, transmission, or ancillary services with CRSP in the previous month to verify accounting records. The Postschedulers ensure that customer bills are accurate and timely. They also verify that incoming bills match CRSP records and see that bills are paid on time.

4.3.2.12 Management and Marketing Office Objectives

The primary objective of EMMO staff is to balance supply and demand, guided by the operational criteria and limitations previously described. However, within limitations, there are innumerable hourly release patterns and dispatch drivers that must comply with the operating criteria.

Prior to 1990, SLCA/IP power plant dispatch was primarily driven by market prices. This dispatch philosophy, coupled with a high level of operating flexibility at SLCA/IP projects, allowed Western to produce energy in a pattern that was often distinctly different from its firm loads. However, over the past several years, operational constraints imposed on SLCA/IP resources have become significantly more influential regarding the manner in which power is produced, sold, and purchased. Scheduling goals shifted from a model driven primarily by market prices to a new one driven by customer loads.

Within the boundaries of these operating constraints, SLCA/IP power resources are used to serve firm load. SLCA/IP generation resources are typically "stacked" on top of block purchases as a means of following firm customer load. The block purchases are typically made in constant strips that have a 16-hr duration in on-peak periods and an 8-hr duration when off-peak. Because of operational limitations, Western staff may need to either purchase or sell varying amounts of energy on an hourly basis. The volumes of these variable market purchases are relatively small under the vast majority of conditions.

Market sales can be significant when SLCA/IP resources exceed firm load. Under the load-following objective, excess hydropower generation is sold during hours with the highest

price while complying with operational limits. On-peak sales are limited by maximum SLCA/IP generation levels, which are constrained by hourly ramp rate limits and daily water level change constraints. However, significant excess power generation rarely occurs, because projected power production in excess of SHP is sold to SLCA/IP customers on a short-term basis as AHP. SHP is a fixed level of long-term capacity and energy available from SLCA/IP facilities during summer and winter seasons, which is based on an established risk level; this amount is the minimum commitment level for capacity that Western will provide to all SLCA/IP customers. AHP is the monthly capacity and energy that is actually available based on prevailing water release conditions; it is the amount that Western offers to its customers above and beyond their SHP levels.

The load-following scheduling objective facilitates a strong link between Western's contractual obligations and SLCA/IP operations, requiring dispatch among SLCA/IP power plants to be closely coordinated. This interdependency exists because loads and hydropower resources are balanced whenever feasible. Western is therefore able to indirectly affect SCLA/IP power plant operations and hourly reservoir releases via specifications in its contract amendments. Contract terms that indirectly affect power plant operations include SHP and AHP capacity and energy sales, as well as Minimum Schedule Requirement (MSR) specifications. The MSR is the minimum amount of energy that a customer must schedule from Western in each hour. The load-following dispatch philosophy minimizes scheduling problems and helps Western avoid noncompliant water releases.

In addition to load following, schedulers follow other practices that fall within mandated operational boundaries but are not strict requirements. One such practice involves reducing generation at Glen Canyon to approximately the same minimum level every day during low-price, off-peak hours. Also at Glen Canyon, Western avoids drastic changes to total water volume releases when they occur over successive days. Nearly identical volumes of water are released each weekday, while slightly less water may be released on Saturdays and Sundays. In addition, during the summer season, one cycle of increasing and decreasing Glen Canyon Power Plant output is recommended. This practice increases to a maximum of two cycles during other seasons of the year as dictated by the hourly load pattern.

Scheduling guidelines beyond those specified by strict criteria are practiced not only at Glen Canyon but also at other SLCA/IP power plants. For example, the Collbran Project's daily generation produced by Upper and Lower Molina power plants is scheduled at or near power plant maximum capability for continuous blocks of time, the lengths of which are based on the amount of water that is available for release during a 24-hr period. Western also has scheduling guidelines for daily water releases from Blue Mesa Reservoir. Water is released from Blue Mesa seven days a week to accommodate higher runoffs, except during November through February, when water is not released on Saturdays. The decision not to release water on Saturdays was made for economic reasons so that more water could be released during higher-priced hours during the week.

There are two voluntary informal ramping restrictions with which the SLCA complies in order to avoid conflicts with the U.S. Fish and Wildlife Service and trout fishery interests at the

Flaming Gorge Power Plant. One of these limitations includes an up and down ramp rate restriction of 800 cfs, which serves to limit change in generation to 25–30 MW per hour, depending on the head. In addition, generation fluctuations are limited to one peak per day (peaking capability) during summer months. In winter months a double peak pattern may be implemented depending on water release volumes. These two restrictions further restrict the load-following capability of Flaming Gorge power plant during periods in which there is some river system flexibility (i.e., November–April).

4.3.2.13 Ancillary Services

Ancillary services help maintain reliable system operations in accordance with good utility practice. Some of these services include spinning reserve, non-spinning reserve, replacement reserve, regulation/load following, black start, and voltage support. Their abilities to provide quick start times, fast ramping capabilities, and rapid corrective responses to changes in grid conditions mean that most SLCA/IP hydropower plants are an excellent resource for providing ancillary services. Hydropower plant output levels can be ramped up or down from zero production levels to maximum capability in a matter of minutes without adverse effects on the power equipment. This attribute makes the SLCA/IP assets well suited to providing the interconnected grid with various ancillary services, such as spinning and non-spinning reserves, regulation, and voltage support.

Typically, Glen Canyon Dam is used to serve all CRSP spinning reserve obligations and regulation services. This capacity is not available for serving firm load unless an emergency condition exists. CRSP will respond to a pool emergency request up to the CRSP obligation only. Any such request must be based on an actual emergency condition and not an economic condition. If more than the CRSP obligation is requested, the company making the request must be capable of documenting that energy is not available elsewhere on the system.

Regulation services (up and down) are viewed as a single service. Power production fluctuations at Glen Canyon Dam related to regulation (±40 MW) and/or responses to system emergencies are not included in water ramp-rate calculations given that the ramping criteria are specified on an average hourly ramp, not an “instantaneous” ramp.

Glen Canyon usually provides the entire CRSP pool reserve share for regulation services and spinning reserves. However, Aspinall units at Blue Mesa and Morrow Point are alternatives.

Motoring SLCA/IP units can be sold as spinning reserves up to the maximum output level of the unit based primarily on capacity and head. However, motoring units cannot supply regulation services. Units that are not operating (i.e., motoring or generating) but are ready for operation can be sold as non-spinning reserves.

4.3.2.14 Reclamation Unit Dispatch Guidelines

Although Western markets SLCA/IP capacity and energy, Reclamation owns the power plants and operates most of them. As described earlier, the Montrose EMMO balances supply and demand in part through the hourly scheduling of generation. Dispatch for all CRSP units is performed at the Reclamation Glen Canyon Field Office located at the Glen Canyon Dam in Page, Arizona. Therefore, except for Glen Canyon, all other facilities are scheduled remotely.

Reclamation dispatchers take generation orders from Montrose Preschedulers and real-time operations at the plant level. These orders typically go through the Western DSW BA. The WALC is registered with the North American Electric Reliability Corporation (NERC) as a BA, transmission operator, transmission owner, and planning authority. WALC is also a member of the Western Electricity Coordinating Council (WECC). Any changes in real-time operations, initiated either by Western or by Reclamation, are routed through the DSW Office.

To comply with NERC regulations, all online units are on Automatic Generation Control (AGC), whether or not they are already on regulation services or spinning reserves. Clearly, a unit cannot be on AGC when it is not operating. When operating in "rough zones," the power system stabilizer (PSS) of a generator does not perform up to specifications and cannot provide the NERC operating requirement of automatic voltage regulation (AVR). Therefore, generators are not dispatched to operate in these rough zones. For CRSP units, rough zones are at low output levels. This level is used to define minimum unit-specific generation levels.

Water releases from a reservoir are estimated by two methods, as follows: (1) water releases through turbines, estimated by lookup tables that estimate turbine flow as a function of pool elevation and power output; and (2) accusonic flow meter systems. The systems use the multiple parallel-path acoustic transit-time method to measure flow velocities at discrete elevations in the measurement section. Real-time flow rates are determined by the integration of the flow velocity profile, which achieves accuracies of up to 0.5% of the flow rate. For marketing purposes, flows are based on lookup tables.

Operator computer screens connected to the supervisory control and data acquisition (SCADA) system display the reservoir elevation level, maximum elevation level, full reservoir elevation, and minimum elevation. Also displayed are tailrace elevation, power releases, and bypass releases.

Reclamation determines the number of units put into service by simply taking the Western generation schedule at a plant and dividing it by the capacity of a unit at the plant. The resultant number is rounded up. All operating units at a plant are operated at identical output levels. In situations where one unit at the plant has a higher operating level than others, units will be maxed out at different levels.

If a unit is offline for 4 hours or less, it is put into condensing mode (referred to as motoring). This mode of operation consumes energy, the level of which is unit-specific. At Glen Canyon, the typical motoring of a unit consumes 4 to 5 MW.

If there is a line outage that affects a power plant, maximum generation schedules are reduced accordingly. The same is true of a full or partial unit outage. Some units are placed on "Unit Dropping" status, of which there are three levels. Status 1 units are dropped before status 2 units. Status 3 level units are dropped last. Not all plants have dropping units. Dropping is carried out automatically by DSW (WALC BA). When a unit is taken offline, the disconnect switch to the unit ground is closed, and the disconnect switch to the transformer is open. The opposite occurs when the unit is online.

In real-time operations, ramping for load following is performed during a 20-minute time span from 10 minutes before the hour to 10 minutes after the hour.

The preceding section has described the complexities that need to be considered for the scheduling and dispatch of SCLA/IP resources. In this section, we describe how the CHEERS model may be used to improve day-ahead and real-time scheduling by the Western EMMO and unit-level dispatch by the Reclamation Glen Canyon Field Office. Most of the discussion in the balance of this paper will focus on the CRSP and the Seedskadee Project. These two projects account for the great majority of SLCA/IP capacity and energy. A schematic diagram of CRSP and Seedskadee as represented in the CHEERS modeling framework is shown in Figure 2-3. In the actual model software, the user would create this network in the CHEERS GUI by dragging and dropping nodes, connecting these nodes with links, and providing essential information about each component. The user would also define the temporal granularity and simulation time period. In this case, the demonstration site will initially have a 1-hr time step and optimize operation over a 1-week time period; that is, a primary forecast period of 168 time steps. Historical information about the system will also be input into the model to ensure that future operations will not violate limitations, such as the multiple-day reservoir elevation change limits described in previous sections.

4.4 TOOLSET DEMONSTRATION RESULTS AND DISCUSSION

4.4.1 Aspinal Cascade

4.4.1.1 Hydroforecasting

The Aspinall Unit (210 MW) of the Colorado River Storage Project is located at the South Fork of the Upper Gunnison River Basin. It consists of a series of dams: Blue Mesa, Morrow Point, and Crystal. The flow is then largely diverted into the Gunnison Tunnel downstream of Crystal Dam. The area drained into the Aspinall Unit is about 10,000 km^2. Precipitation is relatively constant throughout the year, whereas temperature displays a strong seasonal cycle with temperatures below freezing from October to April (McCabe and Hay 1995). Around 70% of the flow from the Gunnison River is from snowmelt (Regonda et al. 2006a, 2006b; McCabe and Hay 1995). April 1 snowpack can account for about 70% of the variability in annual runoff, indicating the utility of long lead flow forecasts (McCabe and Hay 1995). The

Aspinall Unit is operated as a tightly coupled multi-purpose system and where the main purposes of operations are to provide seasonal water supply and short- to medium-range flood control. A secondary use at shorter time scales is hydropower. Blue Mesa operations are scheduled jointly with Crystal and Morrow Point, because Crystal has minimal storage capacity and needs to handle contributing flow for flood control and hydropower generation. Therefore, flow forecasts at both short-range and seasonal time scales are useful for optimizing Aspinall Unit operations.

The following sections discuss the application of EHFS to the Gunnison River basin. The EHFS provides medium-range and seasonal ensemble forecasts to WUOT seasonal hydrosystem analysis, environmental performance, and day-ahead scheduling and real-time operation components at Blue Mesa Dam, Morrow Point Dam, tributary inflow between Blue Mesa and Morrow Point, Crystal Dam, tributary inflow between Morrow Point and Crystal Dam, the North Fork of the Gunnison just above the confluence with the Gunnison, the Uncompahgre above the confluence with the Gunnison, Delta, and Grand Junction (Table 4-5).

TABLE 4-5 Forecast Point Lookup Manually Established to Provide a Relationship between the Specified Nodes of the Toolset System, the EHFS Routing Network at a One-Eighth-Degree Spatial Resolution, and On-the-Ground Latitude/Longitude of Existing Gauge Stations Whenever Possible

Forecast Point	Latitude	Longitude	Feature Description	EHFS Latitude	EHFS Longitude
BLMSA	38.453370	–107.334052	Blue Mesa Dam	38.4375	–107.3125
MRWPT	38.451700	–107.537900	Morrow Point Dam	38.4375	–107.4375
MRWPS			(side flow between BLMSA and MRWPT)		
CRYST	38.510402	–107.623774	Crystal Dam	38.4375	–107.5625
CRYSS			(side flow between MRWPT and CRYST)		
NFORK	38.782700	–107.837000	Just above the confluence of North Fork Gunnison and Gunnison	38.8125	–107.6875
UNCOM	38.756800	–108.090200	Uncompahgre River above the confluence with the Gunnison	38.6875	–108.1875
DELTA	38.753	–108.078	Delta (below confluence of the Gunnison and Uncompahgre)	38.8125	–108.313
GRNJC	38.983	–108.45	Whitewater (happens to be in preexisting Grand Junction GRNJC grid cell)	38.9375	–108.438

Setup of the EHFS to the Gunnison Basin. The calibration of the EHFS hydrology model was performed at Blue Mesa (Figure 4-10), which is the most upstream reservoir of the Aspinall Unit and the largest reservoir in the unit. Monthly naturalized and daily observed impounded inflow into the Blue Mesa reservoir are available for calibration. The regulation of the daily flow includes regulation from the upstream Taylor Park reservoir, transfers across the Continental Divide, diversion for 45,000 ha of irrigated land, and domestic and industrial uses. The calibration of the surface water hydrology model was performed in two steps: (1) a monthly calibration of the hydrology model ensures that we capture the annual and seasonal water balance of the observed monthly naturalized flow; (2) a daily calibration of the routing model is performed to capture daily variability of the observed daily regulated flow, but not the impounded annual water balance.

The hydrologic model was calibrated at a monthly time step at one-eighth-degree spatial resolution using an automatic calibration of soil parameters with respect to monthly naturalized flow at Blue Mesa (provided by the U.S. Bureau of Reclamation). Standard routing model parameters were used for this monthly VIC calibration. The monthly calibration was performed using the Multi-Objective Complex Evolution of the University of Arizona (MOCOM-UA) method (Yapo et al. 1998), as applied in Voisin et al. (2011) to identify optimal soil parameters.

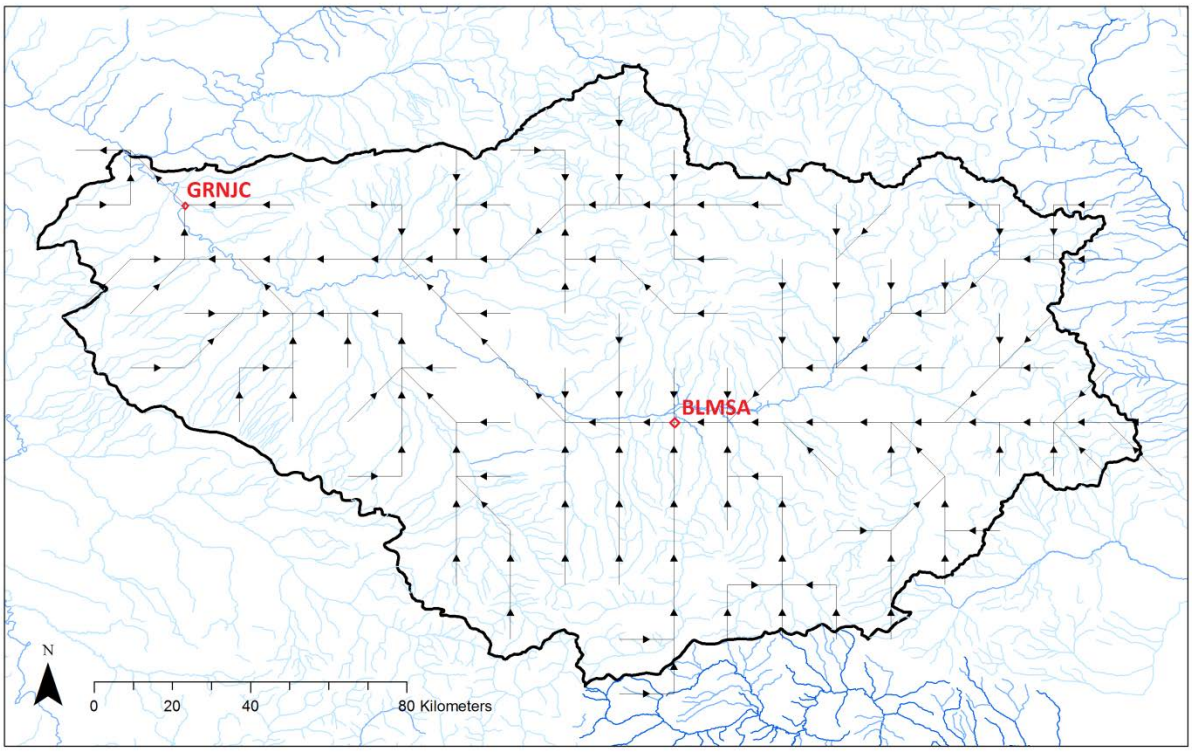

FIGURE 4-10 Routing Network for the Gunnison River Basin (Blue Mesa and Grand Junction forecast points are shown.)

Multi-objective automatic calibration saves time and allows fitting of the simulated hydrograph to different characteristics of the observed hydrograph. Objective measures to be optimized for this monthly calibration included Nash-Sutcliffe Efficiency (NSE), which combines correlation, bias, and variability; NSE of the logarithmic (base 10) monthly flows in order to focus on monthly peaks; and annual water balance bias. In addition, given the large interannual and decadal variability over the Gunnison Basin, the entire retrospective period (1980–2010) was used for the calibration of the routing model in order to overlap as many variations as possible.

Daily calibration of diffusion and velocity parameters of the routing model was performed at Blue Mesa over the same period with respect to observed daily regulated flow. The daily calibration was used to better capture the daily variability, but not necessarily the long-term water balance, which is affected by the impoundment in the basin. In order to give less weight to the annual water balance, the NSE metric was decomposed into three metrics. The objective functions for the daily calibration were the daily correlation, the daily bias and the variability, and the annual water balance. Despite the lower weight given to the annual water balance, the water balance errors between the monthly calibration of VIC calibrated to naturalized flow dominated the source of errors in the automatic calibration process, and calibration parameters did not converge to a stable solution. The default routing model parameters were kept for the setup. This affected the skill of the very short-term flow forecast but not the seasonal forecasts. The post-processor, which is integral part of the EHFS, handles the biases with a very short lead time. Data assimilation is expected to decrease the bias, not only for short-term and for a couple of days ahead, but also for a horizon that corresponds to the time of concentration.

Table 4-6 presents measurements of the flow simulation performance for the calibration period (1980–2000) and the verification period (2000-2010) for the following components: correlation, NSE, raw and relative RMSEs, and mean squared errors (MSEs), which are standard measures.

Figure 4-11 shows 1980–2000 monthly flow climatology at Blue Mesa, comparing observed naturalized flow, observed regulated flow and the EHFS calibrated simulation. The middle panel shows the mean monthly flow for the validation period, 2001 – 2010. This allows a visual check of the monthly performance.

TABLE 4-6 Monthly Streamflow Simulation Performance in the Gunnison River Basin at Blue Mesa

Parameter	Calibration (1980–2000)	Validation (2001–2010)
Annual Relative Bias	0.08	0.08
Monthly correlation	0.93	0.96
Monthly NSE	0.82	0.81
Monthly NSE (log flow)	0.81	0.77
RMSE/Obs Mean	0.47	0.44
MSE/Obs Mean	0.18	0.19
RMSE (cms)	19.2	13.5

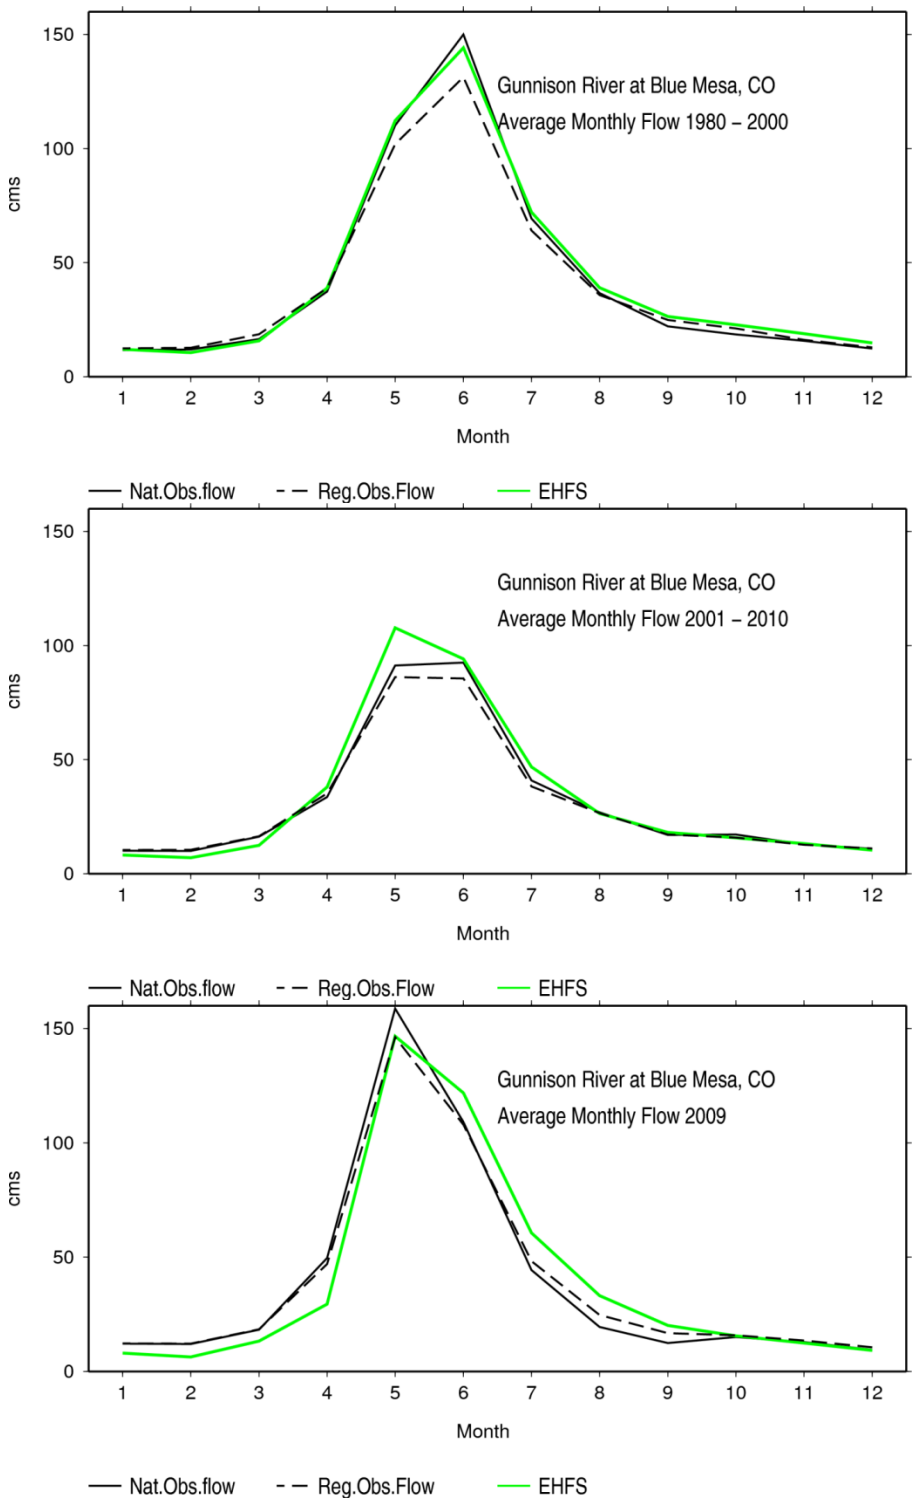

FIGURE 4-11 Top: 1980–2000 Calibrated Monthly Flow Climatology at Blue Mesa, with Observed Naturalized Flow, Observed Regulated Flow, and EHFS Simulation; Middle: Mean Monthly Flow for 2001–2010 Validation Period; Bottom: Mean Monthly Flow for the 2009 Demonstration Period

Data Assimilation. Within the EHFS system, the hydrology model is forced with observed meteorology prior to issuing the forecast. This process, termed the "nowcast," is used to estimate the basin snow and soil moisture state at the time of the forecast. The EHFS nowcast for 2009 is compared to observed naturalized streamflow in Figure 4-12. Although the simulated flow captures the general trend of observed flow, there are noticeable differences as well. These differences reflect errors in the simulated soil moisture state that will directly influence the accuracy of model forecasts (Figure 4-13). The accuracy of the nowcast can be improved by using observed stream gauge data.

Because streamflow represents the integrated hydrologic response of a basin, streamflow observations contain indirect information about the moisture state of the basin. These observations can potentially be used to correct modeled soil moisture storage values and provide better initial states for forecast generation. Thus, assimilation of observed streamflow to update the modeled soil moisture state has the potential to improve forecast skill, especially at short forecast lead times.

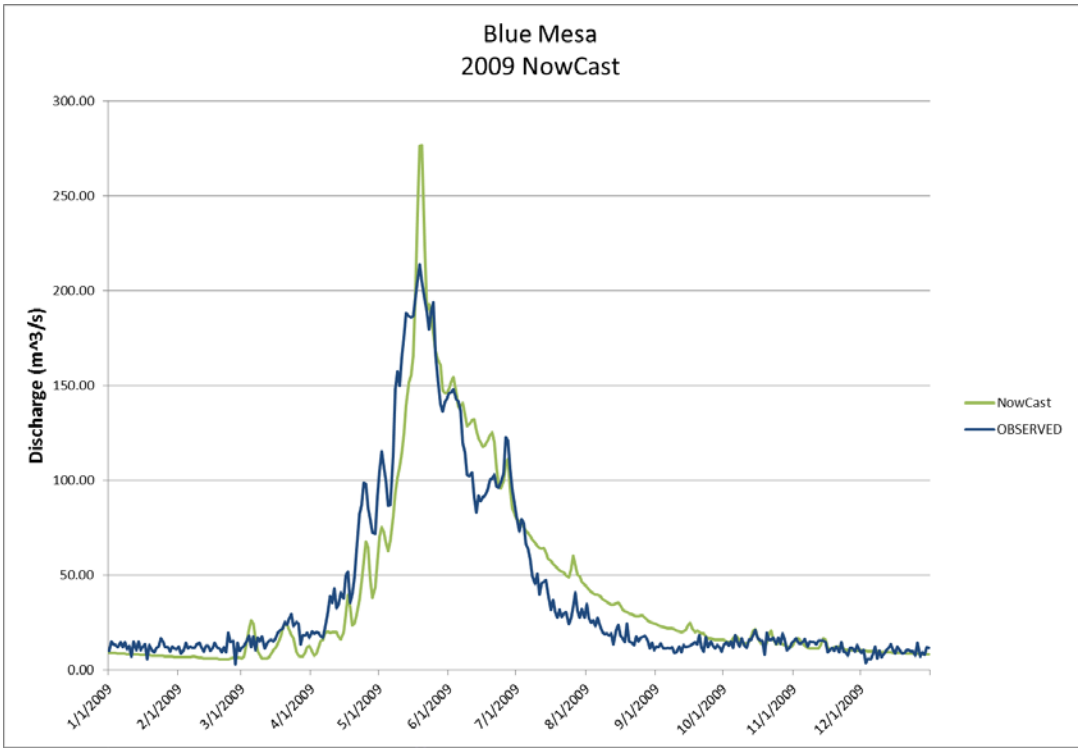

FIGURE 4-12 The EHFS Nowcast and Observed Flow at Blue Mesa for 2009

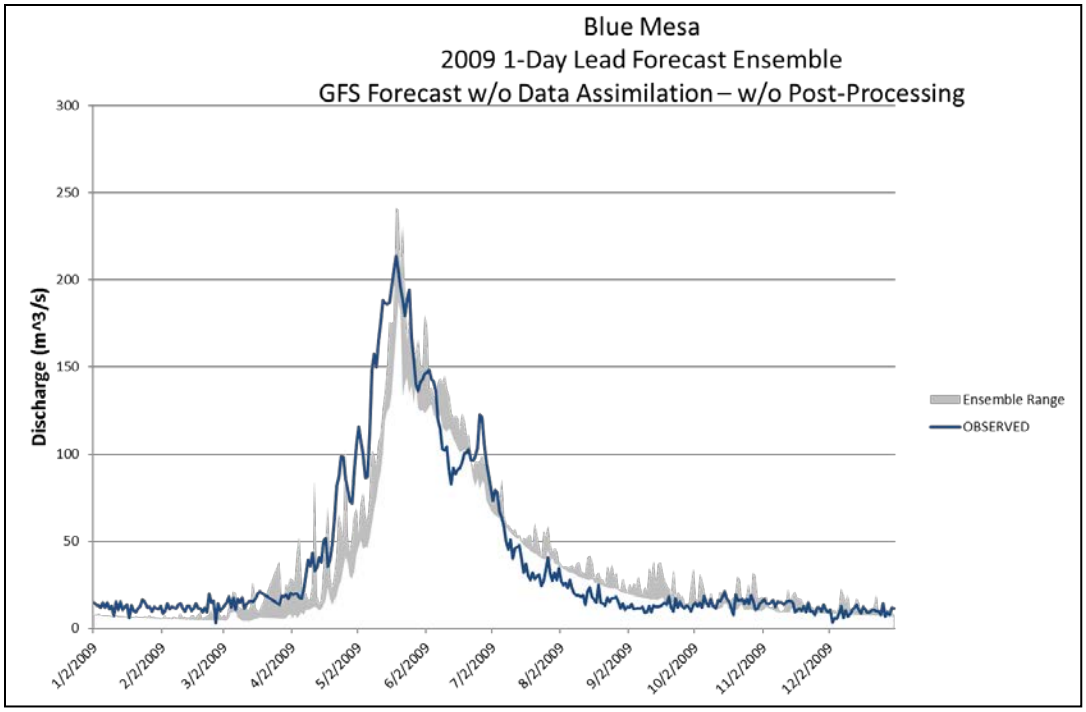

FIGURE 4-13 Observed Inflow at Blue Mesa for 2009 and EHFS 1-day Lead Ensemble Forecast without Streamflow Assimilation or Post-processing

A streamflow assimilation procedure was implemented that uses near–real-time observed streamflow to update the model initial moisture state (see Appendix M for details). In this first implementation, soil moisture values in the deepest model soil layer, the layer responsible for the generation of baseflow, are directly replaced with values estimated from the observed streamflow. Soil moisture values are updated for all areas upstream of a given streamflow gauge in such a way that the spatial patterns of wetter and drier cells are preserved.

Within VIC, baseflow is modeled as drainage from a linear reservoir when the moisture storage is less than a threshold amount of storage. Above this threshold, baseflow generation is modeled as a nonlinear reservoir. In addition to this baseflow generation mechanism, the model also has a fast response runoff mechanism, which generates flow based on a parameterized saturation curve for the top soil layer. It is this latter curve that gives the model its name, but flow is generated through this mechanism mostly during high-flow events.

The assimilation procedure consists of two parts. The first part establishes a relationship between modeled streamflow and the spatial mean moisture storage in the deepest model soil layer (the third layer in our setup for the Gunnison and Feather Rivers). This relationship is developed based on a retrospective model simulation. Figure 4-14 shows an example of this relationship for the area upstream of the gauge at Blue Mesa in the Gunnison River Basin. A single third-order polynomial was used to model the relationship, with the same curve used for all seasons. The values used to establish the curve are smoothed using an *n*-day forward-looking window to diminish the impact of timing errors on the assimilation process. Note that the update

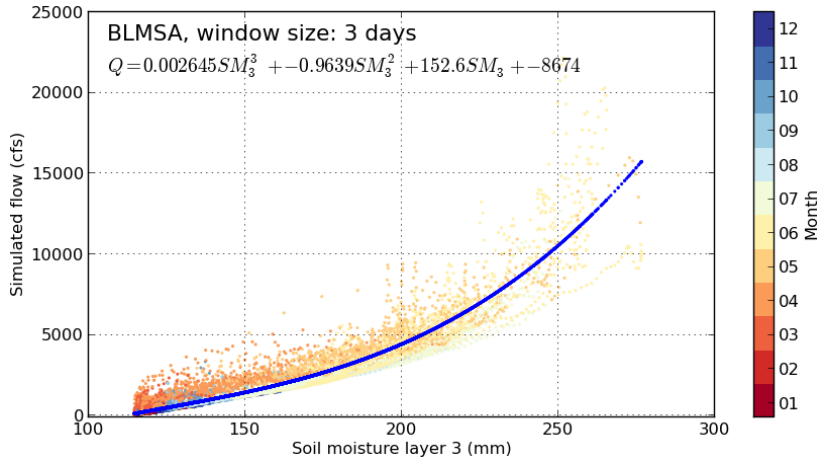

FIGURE 4-14 Relationship between Simulated Streamflow and Spatially Averaged Model Soil Moisture in Lowest Layer for the Area Upstream of Blue Mesa (BLMSA) in the Gunnison River Basin

is performed a few days in advance of the forecast date, because it will take some time for the updated soil moisture to impact the flow at the gauge.

Figure 4-15 shows an example of the soil moisture update for three subbasins in the Gunnison River Basin. Early results from the application of the assimilation procedure show that the assimilation tends to improve the correlation between the forecasted and observed flows and reduce the RMSE in the forecasts, but that it somewhat increases the bias. Reduction of the bias at short forecast lead times is the focus of current work to post-process the forecasted streamflow sequences.

Medium-Range Forecasts. The medium-range forecast provides distributed hourly meteorological forecasts and daily flow forecasts out to 13 days for 15 traces and an average trace for specified forecast points (see Table 4-5). As described previously, the GFS medium-range weather forecasts were downscaled to the grid of the hydrological model and calibrated to remove systematic biases and improve the ensemble forecast probabilistic skill. The downscaled data were used as forcings in the hydrology model, which was previously spun-up using near-real-time observed meteorology.

Medium-range flow forecasts were generated for each day in 2009, and the 1-, 2-, and 3-day-ahead flow values at Blue Mesa were extracted for each day's forecast and these values compared to the regulated flow values. The comparison of these flow ensembles to the daily regulated flows are shown in Figures 4-16 through 4-19. One-day lead streamflow forecasts for 2009 at Blue Mesa with data assimilation are shown in Figure 4-16. Comparison of Figures 4-13 and 4-16 demonstrates the value of streamflow assimilation to update EHFS basin soil moisture, producing a significant improvement in forecast accuracy.

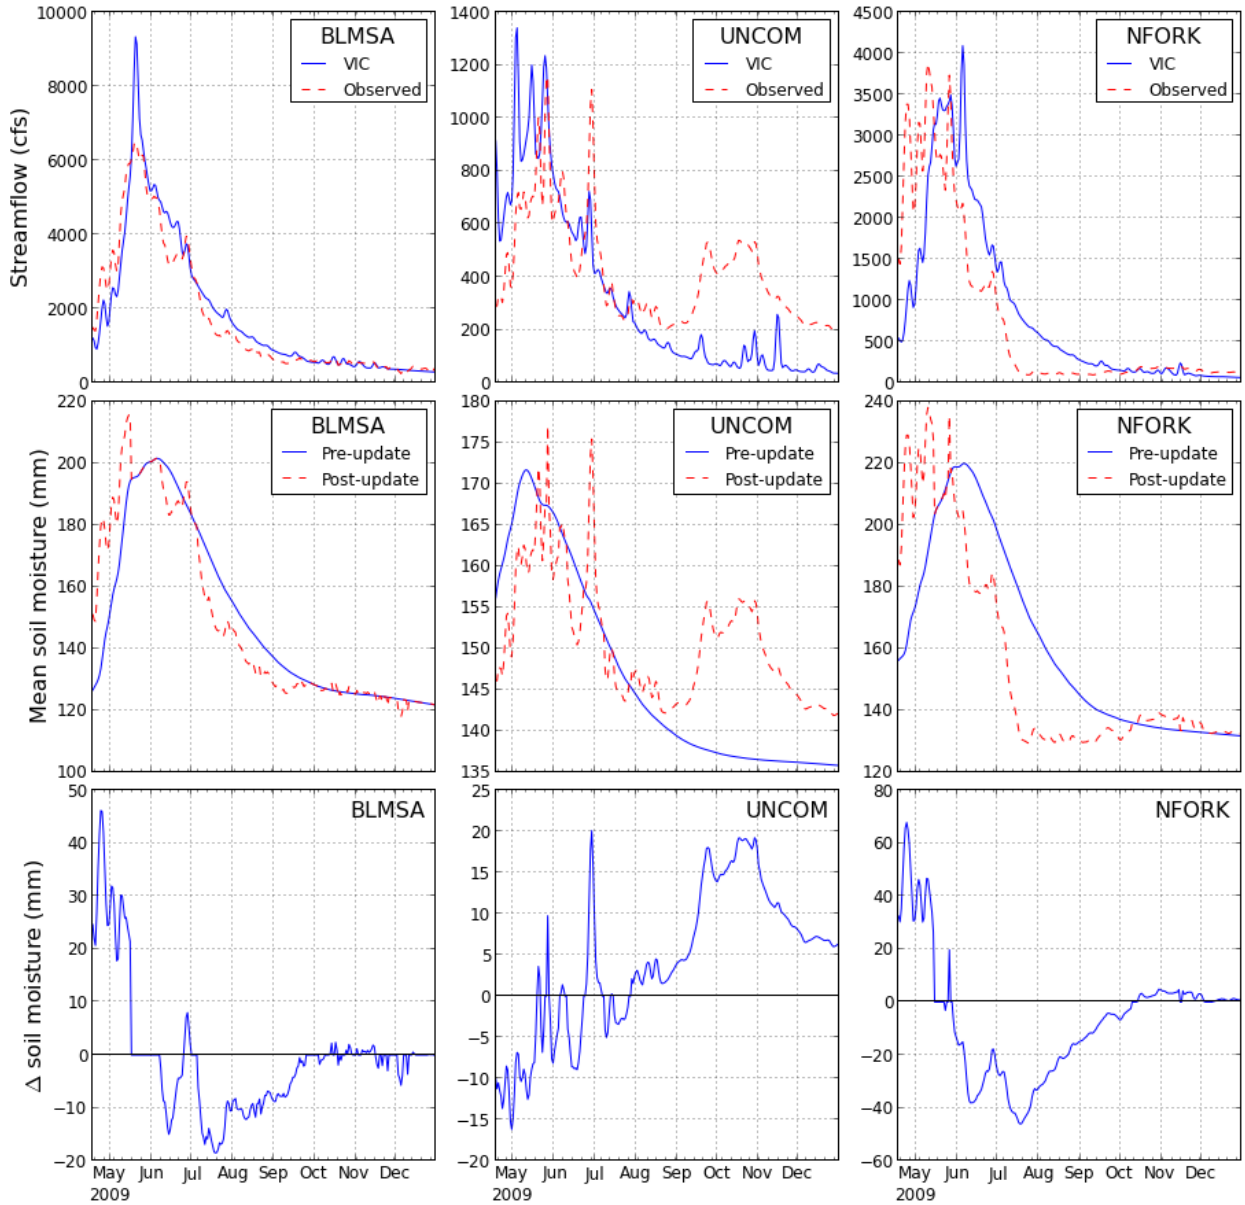

FIGURE 4-15 Effect of Streamflow Assimilation on the Model's Soil Moisture State for Three Subbasins in the Gunnison River Basin (BLMSA = Blue Mesa; UNCOM = Uncompaghre; NFORK = North Fork of the Gunnison River. Top: observed and simulated streamflow for the period from mid-April through the end of 2009. Middle: soil moisture in the third layer pre- and post-update. Bottom: change in soil moisture in that layer as a result of the updates. No update is performed when the update would increase the deviation of the simulated flow from the observed. Note that these updates were performed in a non-sequential manner to test the performance of the update algorithm.)

The EHFS forecast with streamflow assimilation (Figure 4-16) still displays a relatively small bias in the forecast ensemble average compared to observations (–2.85 m^3/s). We employed a simple post-processing step that used near-real-time observed streamflow prior to the forecast to reduce this bias (Figure 4-17). The post-processing step reduced the 1-day lead forecast bias from –2.85 to –0.02 m^3/s and the forecast mean absolute error (MAE) from 6.39 to 5.53 m^3/s.

EHFS 2- and 3-day lead forecasts are presented in Figures 4-18 and 4-19. The EHFS ensemble average forecast bias after post-processing remains small, –0.02 m^3/s for both forecasts. The MAE increases from 6.85 to 7.67 m^3/s as the lead time is increased from 2 to 3 days.

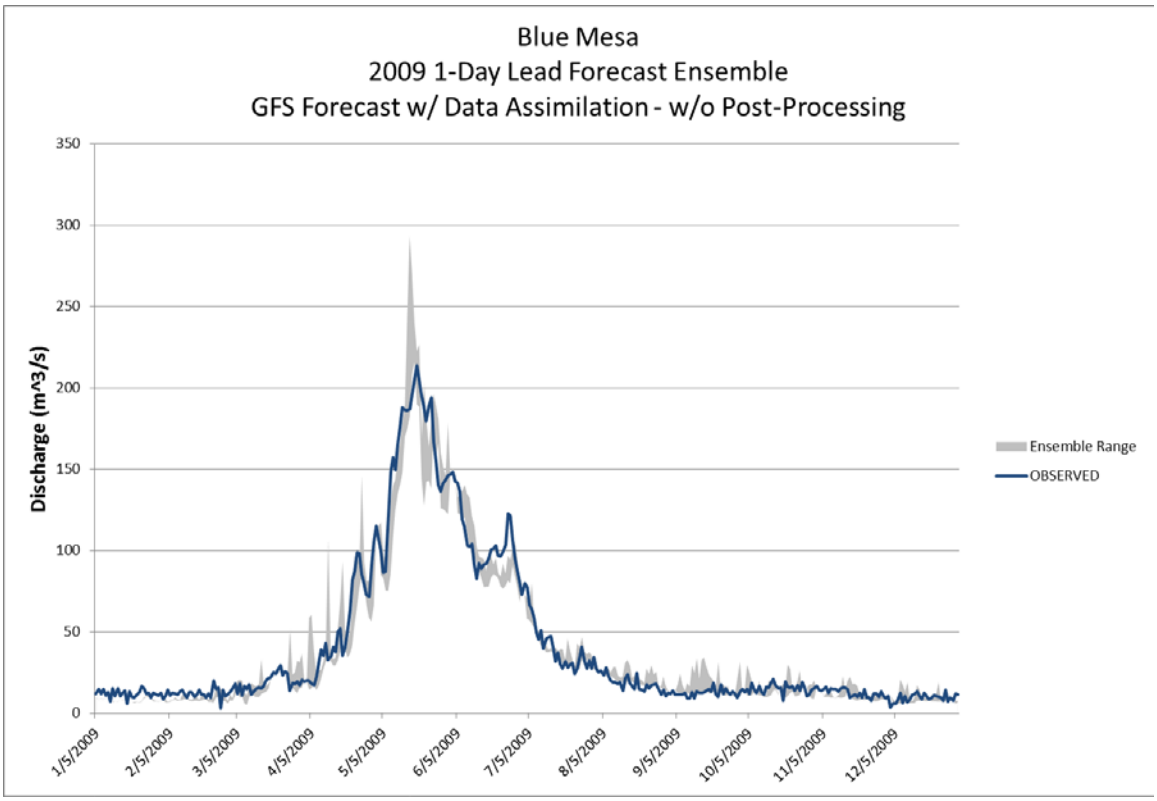

FIGURE 4-16 Observed Flow at Blue Mesa for 2009 and EHFS 1-day Lead Ensemble Forecast with Streamflow Assimilation, without Post-processing (EHFS model bias = –2.85 m^3/s; mean absolute error = 6.39 m^3/s; average ensemble spread = 9.35 m^3/s)

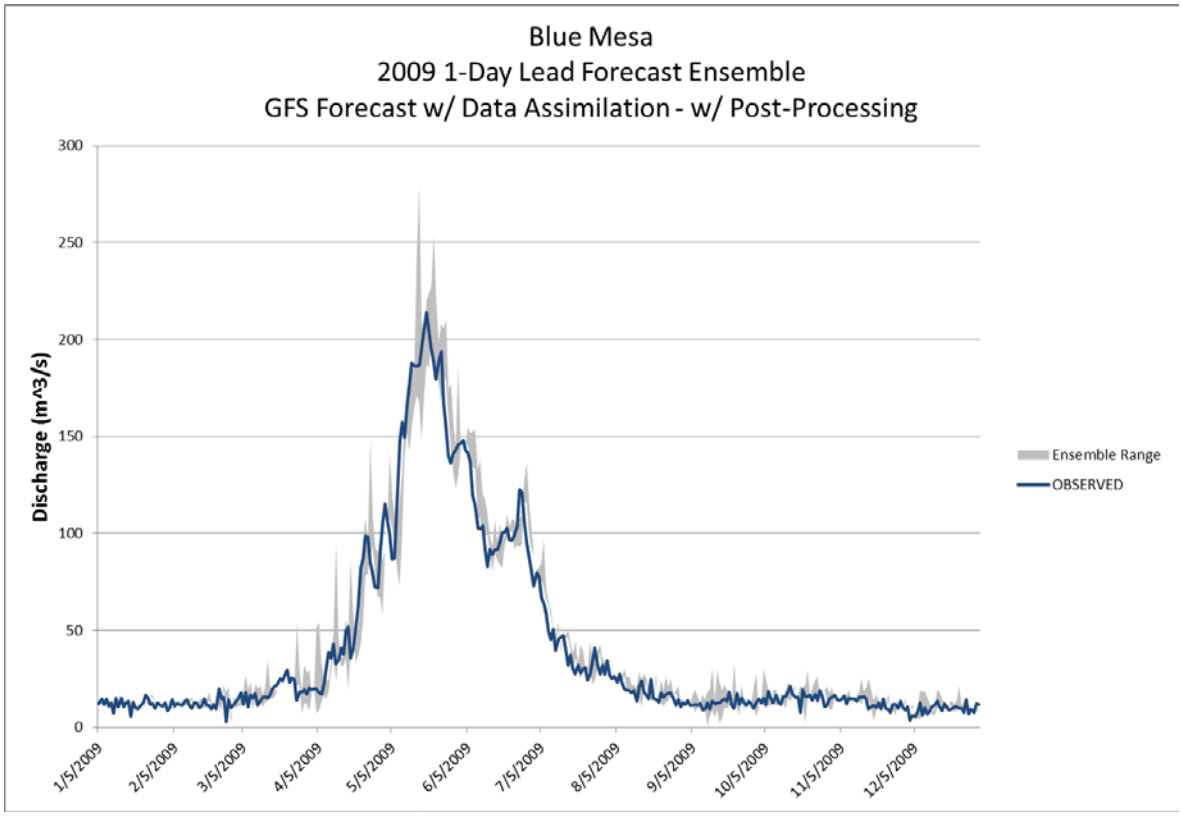

FIGURE 4-17 Observed Flow at Blue Mesa for 2009 and EHFS 1-day Lead Ensemble Forecast with Data Assimilation and Post-processing (EHFS model bias = –0.02 m^3/s; mean absolute error = 5.53 m^3/s; average ensemble spread = 9.35 m^3/s)

Seasonal Forecasts. The seasonal forecasts are generated weekly and include an ensemble of daily data (49 traces + average trace) out to 1 year from the forecast date. Seasonal monthly ensemble flow forecast results are presented for Blue Mesa and Grand Junction, representing the uppermost forecast location with relatively high flows and the most downstream location with the largest flows (Figures 4-20 and 4-21), respectively. Each figure presents 12-month seasonal ensemble forecasts issued April 1, 2009, with forecast flows to March 30, 2010. The black line shows actual observed flow (perfect forecast), the blue line shows EHFS results based on observed meteorology (our best forecast), and the box and whisker plots show EHFS results based on actual ensemble meteorological forecasts. The boxes show the 25%, 50%, and 75% non-exceedance monthly flows, while the whiskers represent the minimum and maximum ensemble members.

Comparison of the blue and black lines indicates the ability of the hydrologic model to represent seasonal hydrologic conditions. Best forecast results are excellent for Blue Mesa and generally good for Grand Junction. The NOAA National Weather Service (NOAA-NWS) provides operational short-range hydrologic model-based deterministic forecasts. NOAA-NWS also provides probabilistic volumetric seasonal forecasts using a blend of model-based ensemble

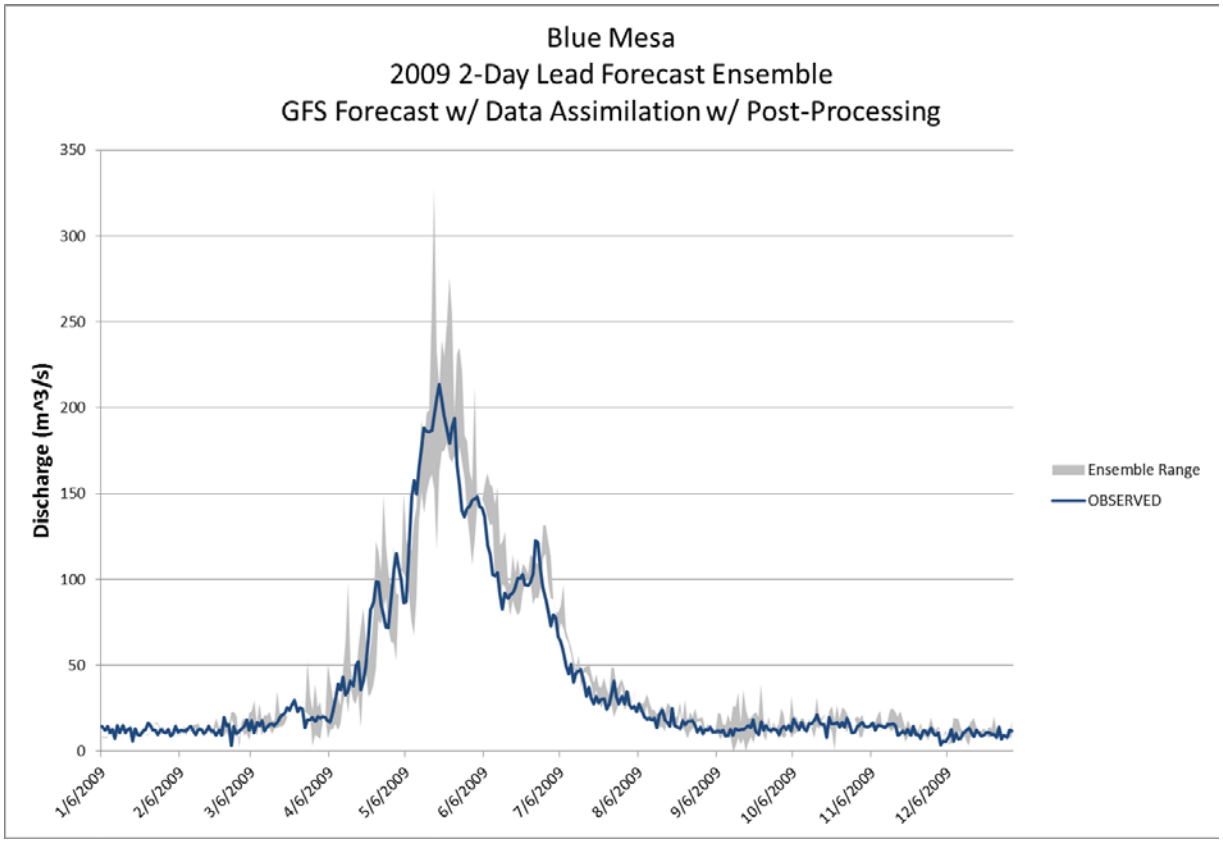

FIGURE 4-18 Observed Flow at Blue Mesa for 2009 and EHFS 2-day Lead Ensemble Forecast with Data Assimilation and Post-processing (EHFS model bias = –0.02 m^3/s; mean absolute error = 6.85 m^3/s; average ensemble spread = 13.74 m^3/s)

forecasts and regression- based methods. The probabilistic forecasts provide different types of information (Figure 4-22): each ensemble member has an equal probability of occurrence. As such, (1) the ensemble mean forecast represents the convergence between the traces and (2) reliable probabilistic forecasts can be drawn from the ensemble. The ensemble mean is evaluated with respect to observed flow; the maximum and minimum values, which define the range of possible outcomes, are expected to bracket the observed flow; the range quantifies the uncertainty around the deterministic forecast. Because we assume that each ensemble member is equally probable, chances that a certain volume will or will not be exceeded can be derived. This is the probabilistic information given by the distribution of the ensemble members. The April 1, 2009, EHFS probabilistic forecast for Blue Mesa shows good agreement with a comparable forecast issued by the NOAA-NWS (Figure 4-22). Although the EHFS forecast means are below those of the NWS, the uncertainty (range) and probabilistic information (quantiles) agree very well.

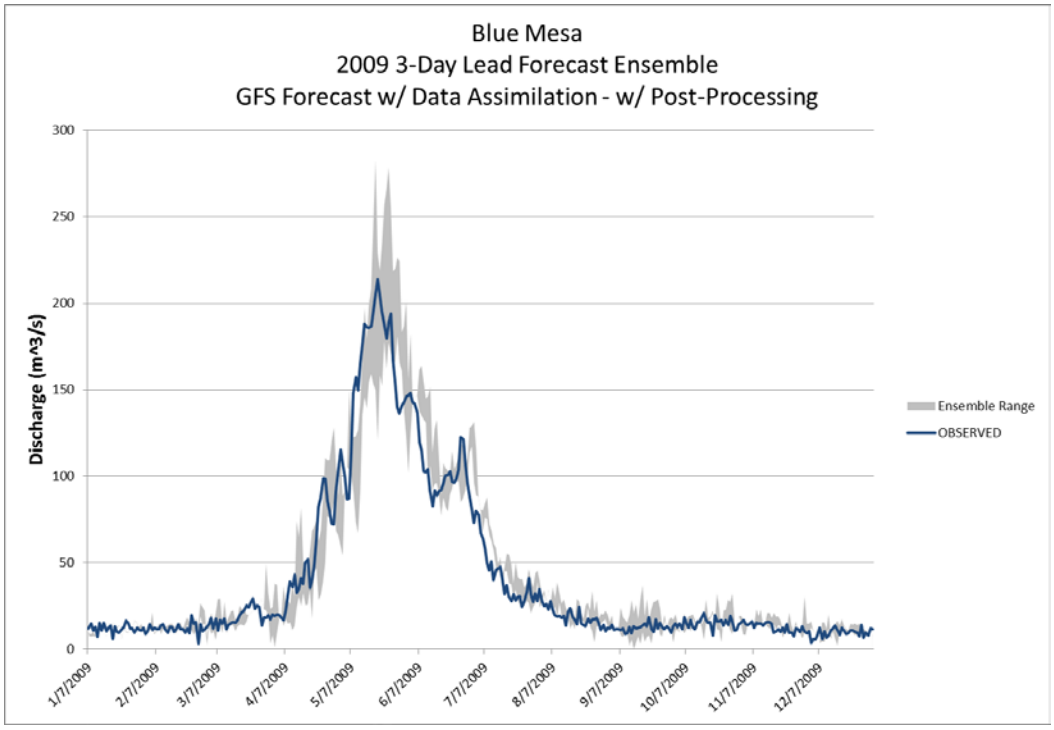

FIGURE 4-19 Observed Flow at Blue Mesa for 2009 and EHFS 3-day Lead Ensemble Forecast with Data Assimilation, with Post-processing (EHFS model bias = –0.02 m^3/s; mean absolute error = 7.67 m^3/s; average ensemble spread = 14.2 m^3/s)

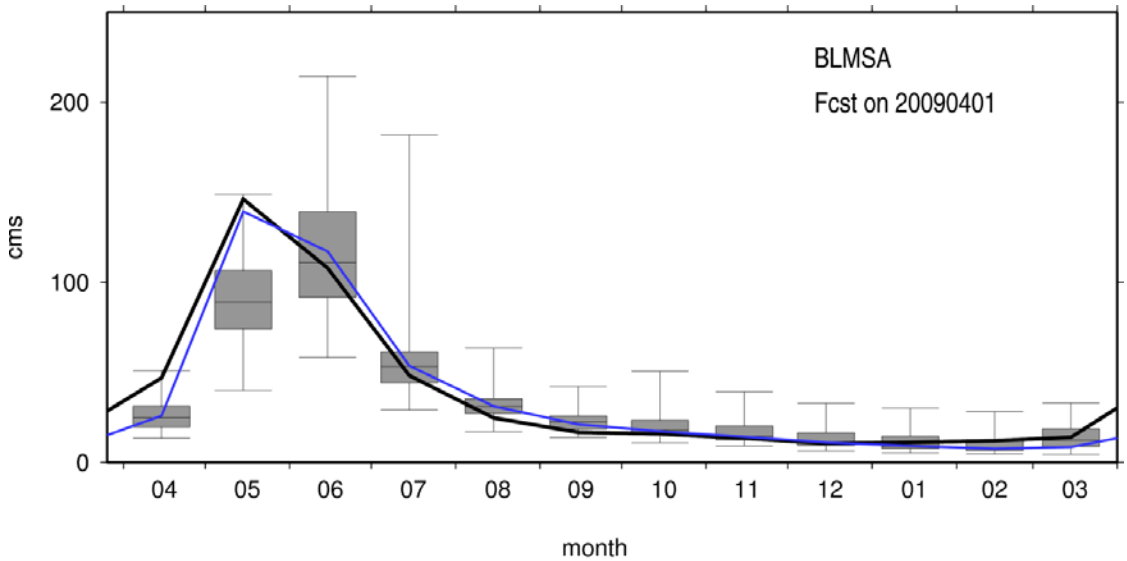

FIGURE 4-20 12-month Seasonal Ensemble Flow Forecast for Blue Mesa Issued April 1, 2009, with Forecast Flows to March 30, 2010 (The black line shows observed naturalized flow, the blue line shows EHFS results based on observed meteorology, and the box and whisker plots show EHFS results based on ensemble meteorological forecasts. The boxes show the 25%, 50%, and 75% exceedance flows, while the whiskers represent the minimum and maximum ensemble members.)

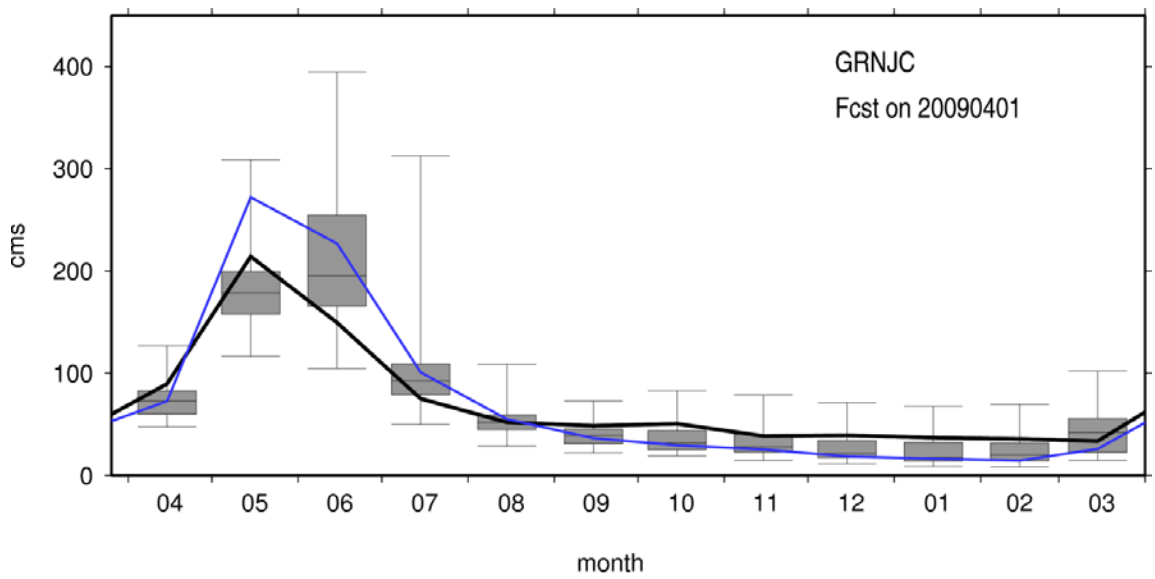

FIGURE 4-21 12-month Seasonal Ensemble Flow Forecast for Grand Junction Issued April 1, 2009, with Forecast Flows to March 30, 2010 (The black line shows observed naturalized flow, the blue line shows EHFS results based on observed meteorology, and the box and whisker plots show EHFS results based on ensemble meteorological forecasts. The boxes show the 25%, 50%, and 75% exceedence flows, while the whiskers represent the minimum and maximum ensemble members.)

The NOAA-NWS forecasts rely on extensive site-specific forecaster experience to process the weather forecasts based on retrospective simulation performances and manually adjusting the model state and selected parameters based on observed streamflow. In addition, the NOAA-NWS operational flow forecasts are currently issued using a lumped hydrology model, limiting the opportunity to assimilate spatially distributed observations, such as satellite-based snow-covered extent, and are not broadly available for the nation's reservoirs. Figure 4-22 suggests that the EHFS may have the capacity to provide a more general and broadly applicable approach that produces very similar forecast results without the reliance on the experienced personnel for each geographic region.

The NWS developed and set up a Community Hydrologic Prediction System (CHPS), a platform from which multiple hydrology models, sources of weather forecasts, data assimilation approaches, and post-processing approaches could be run. It would allow interevaluation of the products, promote diversity of approaches, and accommodate the fact that some approaches work better in some basins and applications than others. Most importantly, the CHPS platform provides a standard format that would allow reservoir operators to choose the optimum products for their specific application and location. To gain a better understanding of CHPS, the forecast team visited the NWS Northwest River Forecast Center in Portland, Oregon, in January 2011. In addition, two meetings with the NOAA-NWS Office of Hydrologic Development (OHD) took place to collaborate on the data assimilation and forecast approaches.

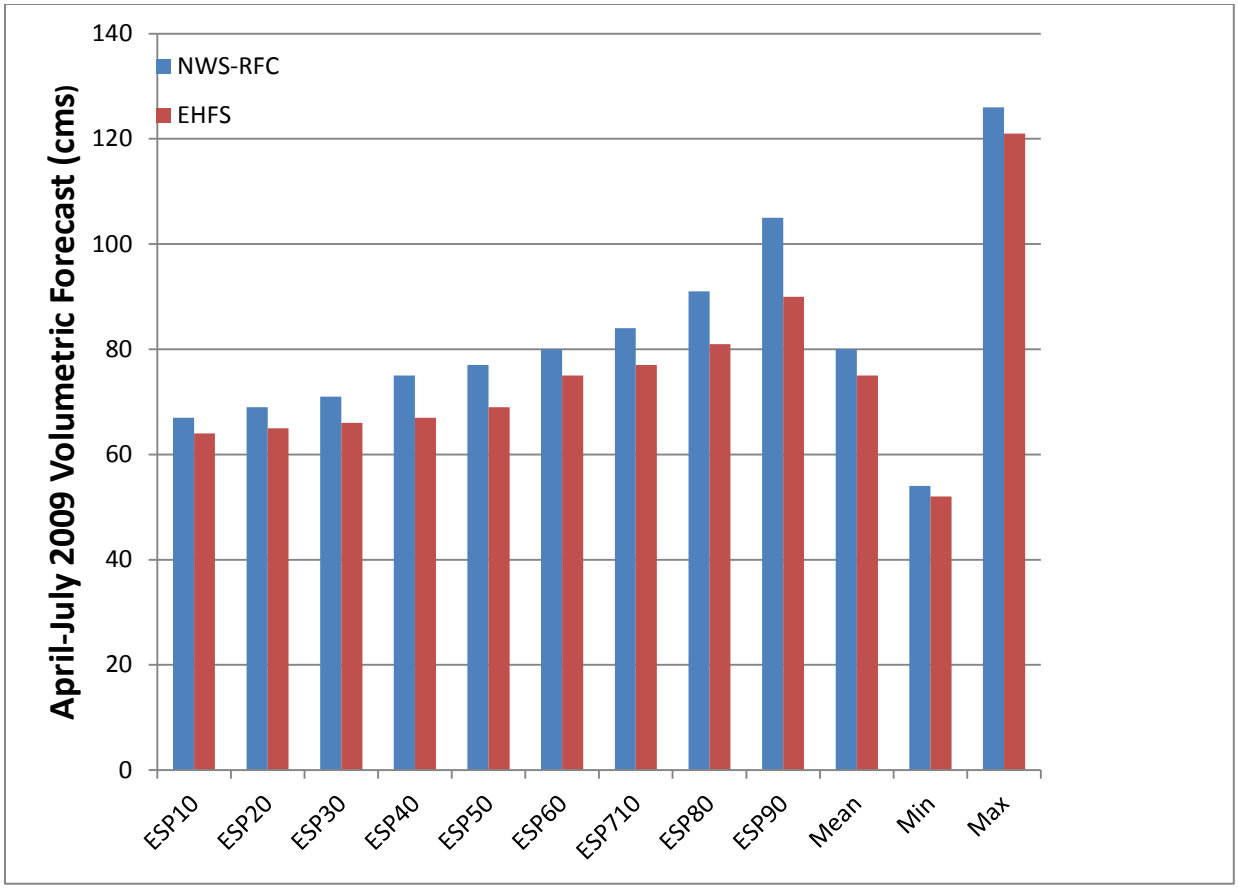

FIGURE 4-22 Evaluation of the April 1, 2009, EHFS Volumetric Probabilistic Forecast with Respect to the Equivalent Forecast Issued by the NWS Colorado Basin River Forecast Center (Retrospective seasonal forecast provided by Dr. Andy Wood, NWS. The ESP*xx* shows the seasonal volume corresponding to 10, 20, … 90 % non-exceedance probabilities. The ensemble forecast mean, minimum and maximum are also shown to provide the entire range of the uncertainty.)

Products. The database was populated with streamflow and meteorological forcing data. For each streamflow forecast day (daily April 1, 2009, through September 30, 2009), an assemblage of the ensemble average of the medium-range forecast and the seasonal forecast issued before the medium-range forecast date was entered into the database. Each forecast day had 344 days with forecast flows reported in CFS. The first 13 days are produced from the medium-range forecast, and the rest of the forecast period are produced from the daily data of the seasonal forecast.

The medium-range ensemble meteorological forecasts are derived from the pre-processed GFS medium-range weather forecasts. They include one-eighth-degree gridded, 13-day, 15-member, hourly meteorological data, issued daily for the period from 1990 to 2010. They also include data parameters such as precipitation, air temperature, vapor pressure, atmospheric density, wind speed, specific humidity, incoming shortwave radiation, and incoming longwave radiation.

The seasonal forcing group includes a 49-member ensemble hourly meteorological forcing data at one-eighth-degree spatial grid resolution for a horizon of 366 days. The dataset includes data parameters such as precipitation, air temperature, vapor pressure, atmospheric density, wind speed, specific humidity, incoming shortwave radiation, and incoming longwave radiation. The dataset is static because each trace is a resampling from each of the 49 or 50 years (leave 1 year out when overlapping) in the 1960–2010 period.

4.4.1.2 HydroSCOPE

For the Aspinall Cascade, a retrospective analysis was performed to simulate how the WUOT would actually be implemented over a 6-month period that ran from April 1, 2009, to September 30, 2009. For HydroSCOPE, a full DAKOTA optimization is run for each day of the simulation. While the multi-objective genetic algorithm (MOGA) employed by DAKOTA requires a high number of simulations to exhaust its search space (generally >>1000, depending on the problem), we found that an adequate balance between simulation run time and accuracy was obtained with just 500 simulations. Each day of the simulation consists of seven steps, as follows:

1. Obtain the 6-month hydrologic forecast from the EHFS.

2. Run the DAKOTA optimization to adjust releases to maximize revenue and the IRF score.

3. Select the optimal release schedule from the DAKOTA results and extract the predicted reservoir releases.

4. Perform a deterministic simulation that uses the predicted reservoir releases, historical inflows from data for the first day, and forecasted inflows for the balance of the simulation.

5. Write the results from Step 4 to the common database for use by the other tools.

6. Extract the system conditions at the start of the second day of the simulation for use as the initial condition of the next day's simulation.

7. Return to Step 1 and repeat for each day running from April 1, 2009, through September 30, 2009.

Before running the optimization, an operational ruleset was created that allows HydroSCOPE to realistically determine the baseline release schedule based on the inflow forecast from the EHFS. For each simulation in the optimization, the baseline releases are adjusted as described in Section 3.3.3.10 to maximize both the total revenue and the IRF score. To calculate the IRF score after each simulation, the flow rate and stage elevation for the two

environmental evaluation points were sent to the IRF tool, which returned a single integrated score for that simulation. The calculation of the IRF score is described in detail in Section 4.4.1.3. The balance of this section describes the application of HydroSCOPE to the Aspinall Cascade, the details for implementing the seven steps listed above, and the results from the analysis.

Model Domain. Figure 4-23 shows a conceptualization of the Aspinall Cascade model as used by HydroSCOPE. The system consists of the three reservoirs (Blue Mesa, Morrow Point, and Crystal), and four river reaches. The forecasts from the EHFS tool are used to supply the five entry points into the model, the Gunnison River above Blue Mesa, an unnamed side inflow

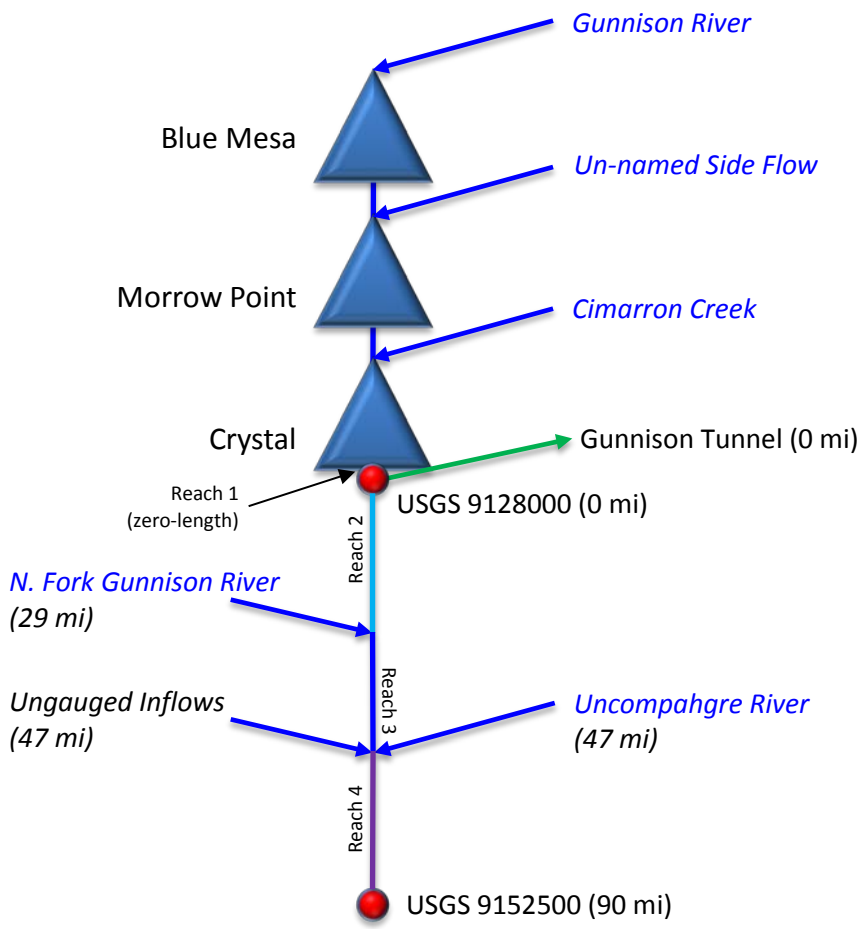

FIGURE 4-23 Conceptualization of the HydroSCOPE Model for the Aspinall Cascade Demonstration Site

above Morrow Point, Cimarron Creek above Crystal, the North Fork of the Gunnision at river mile 29 below Crystal, and the Uncompahgre River at river mile 47 below Crystal. The lower boundary of the model is at the WhiteWater Gauge (USGS gauge 09152500), just south of Grand Junction at the confluence of the Gunnison and Colorado Rivers. A withdrawal from the model occurs just below Crystal Reservoir.

The three reservoirs are modeled as one-dimensional vertical systems with a maximum of 60 layers for Blue Mesa and Morrow Point and 42 layers for Crystal. Maximum and minimum allowed layer thicknesses are 3.0 and 1.5 m for Blue Mesa, 4.2 and 2.0 m for Morrow Point, and 3.0 and 1.5 m for Crystal. Recall that if a layer exceeds its maximum thickness it is split in half and that if it falls below its minimum thickness, it is merged with the layer above it.

The four river reaches consists of one zero-length reach, and three defined-length reaches. The reaches are numbered 1 through 4 from top to bottom. Within HydroSCOPE, a zero-length reach is used to provide a mechanism for water to be added or removed from the river system where the additions and losses are in close proximity to one another. For the Aspinall Cascade, Reach 1 is a zero-length reach that is defined to receive the discharge from Crystal, as well as to release water to the Gunnison Tunnel for agriculture needs in the Uncompahgre basin. The inflow point from Crystal and the outflow point to the tunnel are within a few hundred meters of each other and modeling changes along such a short length is not needed for this type of analysis. The outflow from a zero-length reach to the next reach is the instantaneous inflow (i.e., the release from Crystal) minus the instantaneous losses (i.e., the withdrawal to the tunnel). Similarly, for a zero-length reach, the outflow temperature is equal to the inflow temperature. For the retrospective analysis, the tunnel withdrawal rates were obtained directly from the Bureau of Reclamation Website (www.usbr.gov/uc/crsp/GetSiteInfo) and were used as-is, except where adjustments may have been needed to prevent a violation of downstream minimum flow restrictions. This condition could only occur if the DAKOTA optimization created a scenario where the releases from Crystal were greatly reduced.

An unforecasted inflow point (indicated as “ungauged” in Figure 4-23) is added to the last river reach as a closure term to the water balance and is calculated as the difference between the historical and simulated river flows at the WhiteWater Gauge, where the historical flows are the actual readings at the WhiteWater Gauge and the simulated river flows are based on a “perfect forecast” at the five other entry points and the actual releases from the reservoirs. The ungauged inflows represent the accumulated inflow from numerous minor side inflows, and possibly groundwater discharges that are not captured at the resolution of the EHFS tool. During an optimization run, the ungauged inflows are scaled based on the ratio of the combined forecasted inflow volume to the historical inflow volume from the North Fork and the Uncompahgre inflow points for the timespan of the simulation.

Operational Logic. In order to run the optimization, a set of operational rules were created that allow the model to operate the reservoirs in a representative fashion. The operational rules provide target criteria from which HydroSCOPE calculates the baseline release schedule; it is the baseline release schedule that is adjusted during the optimization process. As discussed in

more detail below, penalty functions on the total revenue are applied as part of the optimization process if the reservoirs are operated too far from their operational targets.

For Blue Mesa, the goal is to operate to the storage curve that was developed with the State of Colorado water allocation model, StateMod (StateMod 2012). A plot of the Blue Mesa elevation curve and the actual elevations for 2009 is shown in Figure 4-24. The calculation for the baseline release involves several steps, the first of which is to calculate the necessary net release (inflow minus outflow) to pull the reservoir elevation towards the elevation curve. This is done by converting the difference between the actual elevation and the target elevation to a volume difference and then dividing that by a user-defined number of days to meet the target (21 days in this case). The days-to-meet-target parameter is used to attenuate the operations and prevent operating regimes from oscillating wildly. Note that the net release is the difference between the total inflows, which is known from the forecasts, and the yet-to-be-calculated outflows.

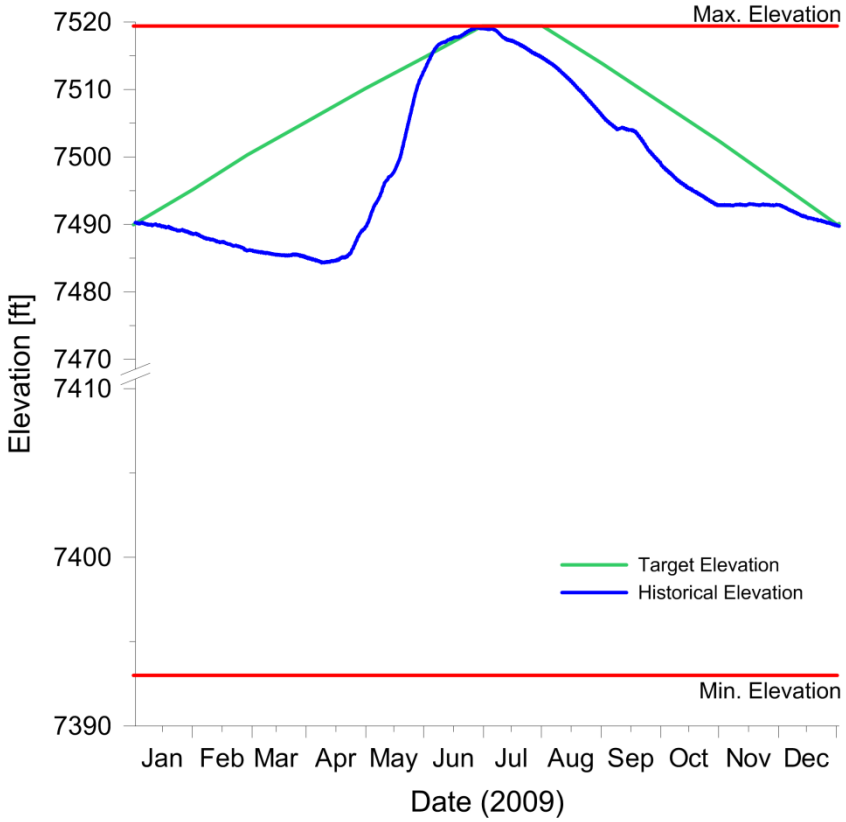

FIGURE 4-24 Plot of the Blue Mesa Target Elevations and the Actual Elevations for 2009

The second step limits the release to a stipulated downstream flow target or to the reservoir inflow minus the sum of the net release from Step 1, the downstream flow target, and the evaporation. The downstream flow target is based on the minimum flow restriction of 300 cfs, the amount of water needed to meet the water demand at the Gunnison Tunnel, and the required deliveries to meet the National Park Service water rights for flows within Black Canyon (Auble et al. 2009). The release from Step 2 is the maximum (or minimum) possible release that complies with the flow targets and the desire to approach and maintain the elevation curve.

The third step limits the release calculated in Step 2 to the maximum physical capacity of the reservoir outlets, while the fourth step makes adjustments to the outflows to insure that maximum and minimum volume restrictions are not violated. The entire process is shown below as equation (4-4):

Step 1
$$Q_1 = \frac{V_n - V_E}{D_T} \tag{4-1}$$

Step 2
$$Q_2 = \max\begin{cases} Q_t \\ Q_i - (Q_1 + Q_t + Q_e) \end{cases} \tag{4-2}$$

Step 3
$$Q_3 = \min\begin{cases} Q_2 \\ Q_{\max} \end{cases} \tag{4-3}$$

Step 4
$$Q_4 = Q_3 \pm \frac{\Delta V}{\Delta t} \tag{4-4}$$

where Q_1 (L^3/t) is the net release rate needed to reach the elevation curve; V_n (L^3) is the current volume of the reservoir; V_E (L^3) is the target volume; D_T (t) is a user-defined parameter that describes the days to meet the elevation target; Q_t (L^3/t) is the target release rate; Q_i (L^3/t) is the sum of the reservoir inflows; Q_e (L^3/t) is the evaporation rate; Q_{max} (L^3/t) is the maximum release possible at the current elevation; ΔV (L^3) is the difference between the minimum or maximum volume and the current volume (equal to zero when the current volume is greater than the minimum and less than the maximum); and Δt (t) is the simulation timestep. Q_4 (L^3/t) is the baseline release for the current timestep.

Morrow Point is operated to the mid-elevation between its maximum and minimum operating elevations while attempting to keep its releases equal to its inflows (i.e., stable elevations). Like Blue Mesa, the calculation involves four steps; the second step is the only step that differs from Blue Mesa. The second step for Morrow Point is based on the desired release, which is defined to be equal to the inflows. An intermediate release is calculated as the desired release minus the net release from Step 1, minus evaporation. The Step 2 release is calculated as the maximum of the downstream flow target and the intermediate release if the current elevation is lower than the target, or the maximum of the downstream flow target, the intermediate release, and the desired release if the elevation is above the target. This preserves the assumption that, when possible, the downstream requirements for the Gunnison will be met by releasing water

from Blue Mesa and passing it through Morrow Point and Crystal. The logic for Morrow Point Step 2 is shown as equation (4-5):

Step 2

$$Q_2 = \begin{matrix} \text{for } E > E_t\text{: max}\begin{cases} Q_1 \\ Q_{des} \\ Q_t \end{cases} \\ \text{for } E < E_t\text{: max}\begin{cases} Q_1 \\ Q_t \end{cases} \end{matrix} \tag{4-5}$$

where E (L) is the current elevation, E_t (L) is the target elevation, Q_{des} (L^3/t) is the desired release, and all other terms are defined above.

Crystal reservoir is operated as a re-regulating reservoir and is used to attenuate the releases from the upstream reservoirs. Within HydroSCOPE, Crystal is operated similar to Morrow Point in that it is always trying to operate to the midpoint of its operating range. However, because Crystal is re-regulating, it uses the forecasted inflows over the next 3 days and bases the net outflow needed to reach its target on an exponential function that reduces the days to target parameter as a function of the elevation difference from the target. In addition, Crystal is only allowed to change its release rate once a day (releases from the other reservoirs can change at each timestep) and its releases are limited by the elevation change limitations shown in Table 4-7. Finally, a hard lower limit of 300 cfs is set to insure that the minimum downstream requirements are always met. Like the other reservoirs, the calculations involve a four-step process, with Steps 3 and 4 the same as the other reservoirs and Steps 1 and 2 shown in equations (4-6) and (4-7):

Step 1

$$Q_1 = Q_{4mp} + \bar{Q}_{i3} \tag{4-6}$$

Step 2

$$Q_2 = \begin{matrix} \text{for } E > E_t\text{: max}\begin{cases} Q_1 \\ Q_{4mp} + Q_i + Q_{\text{add}} \end{cases} \\ \text{for } E < E_t\text{: max}\begin{cases} Q_1 \\ Q_{4mp} + Q_i - Q_{\text{add}} \end{cases} \end{matrix} \tag{4-7}$$

where Q_{4mp} (L^3/t) is the release from Morrow Point, Q_{i3} (L^3/t) is the 3-day average of the forecasted inflows, Q_i (L^3/t) is the inflow at the current timestep, and Q_{add} (L^3/t) is defined as

$$Q_{\text{add}} = \beta \exp\left(\frac{\Delta V}{\Delta V_{\max}} \alpha\right) \tag{4-8}$$

where ΔV (L^3) is the absolute value between the current volume and the target volume, ΔV_{max} (L^3) is the absolute value between the target volume and either the maximum or minimum

TABLE 4-7 Maximum and Minimum Operating Elevations for the Aspinall Cascade (Crystal Reservoir has additional restrictions that limit the rate of change of the water elevation.)

	Elevation or Elevation Change (ft)								
					Crystal				
	Blue Mesa		Morrow Point						
Month	Min	Max	Min	Max	Min	Max	Max Daily Change	Elev for 0.5 ft/day Change	3-Day Elevation Change
Jan	7393	7519.4	7143	7160	6739	6760	10	6733	15
Feb	7393	7519.4	7143	7160	6739	6760	10	6733	15
Mar	7393	7519.4	7143	7160	6739	6760	4	6748	6
Apr	7393	7519.4	7143	7160	6739	6760	4	6748	6
May	7393	7519.4	7143	7160	6739	6760	4	6748	6
Jun	7393	7519.4	7151	7160	6739	6760	4	6748	6
Jul	7393	7519.4	7151	7160	6739	6760	10	6733	15
Aug	7393	7519.4	7151	7160	6739	6760	10	6733	15
Sep	7393	7519.4	7151	7160	6739	6760	10	6733	15
Oct	7393	7519.4	7143	7160	6739	6760	10	6733	15
Nov	7393	7519.4	7143	7160	6739	6760	10	6733	15
Dec	7393	7519.4	7143	7160	6739	6760	10	6733	15

volume (dependent on the state of the reservoir), and α and β are coefficients. Q_{add} is added to the release for cases when the elevation is above the target and subtracted for cases when it is below the target.

Plots for each of the reservoirs operating range, the no-penalty zone, and default and optimized release schedules over time are shown in the results discussion below.

Power Generation. Power generation calculations are based on the available history of each reservoirs actual performance from the Western Area Power Administrations (WAPA) Colorado River Storage Project (CRSP) Website (http://www.wapa.gov/CRSP/opsmaintcrsp/scada.htm). For each reservoir, a polynomial surface fit is made that describes the power generation in megawatts as a function of the water surface elevation and combined flow through the power generating outlets. The fit for Crystal reservoir is shown in Figure 4-25. The fits for Blue Mesa and Morrow Point are similar.

The power is calculated at each timestep and then converted to an equivalent energy by multiplying by the timestep. Revenue is calculated by multiplying the energy produced at each timestep by monthly average electricity prices that are split between peak and off-peak hours. Historical prices from the WAPA Website (http://www.wapa.gov/CRSP/opsmaintcrsp) from September 2008 through September 2012 were used to calculate the monthly average peak and off-peak prices. Since the retrospective analysis assumes that we are forecasting the future and

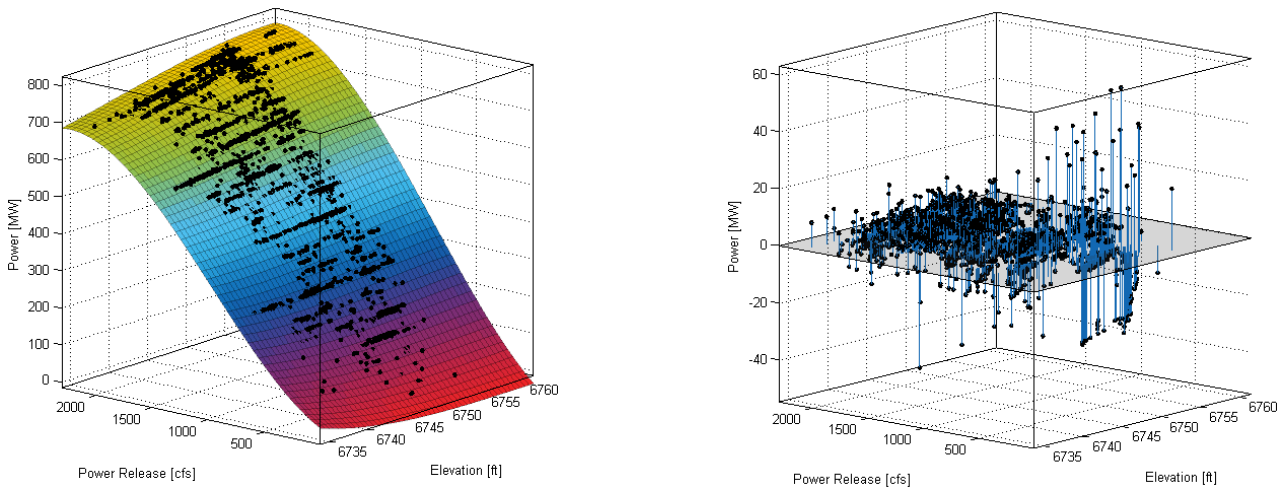

FIGURE 4-25 Plot of the Power Generation Curve for Crystal Reservoir (In all cases, the power is calculated as a function of the water surface elevation and the total outflow through the power generating outlets. The plot on the right shows the residuals between the actual data and the curve fit. With the exception of cases with high elevations and low flows, the surface fit is within 5% of the actual performance. Plots for the other reservoirs show similar traits.)

thus will not know the actual price during any given month, a surrogate in the form of the monthly averages was used (we did exercise some latitude in this assumption by including dates beyond our analysis time period in the analysis). In reality, the approach used to calculate the prices is not important as long as the approach is a representative reflection of the dynamics between reservoir releases, power generation, and revenue. For comparison purposes, a plot of the average monthly prices used in the model and the historical average monthly price for the April 2009 through March 2010 timespan of the retrospective analysis is shown in Figure 4-26.

Penalty Functions. If the optimization is to vary the operations realistically, it must be kept "honest" with respect to the operational rules. This is accomplished through a set of penalty functions that reduces the calculated revenue at each time step based on how far away each reservoir is operating from its respective target. Recognizing that different hydrologic conditions can make it impossible to reach the targets, an operational range is defined where no penalty is incurred. In all cases, the no-penalty operational range is defined to be the halfway point between the maximum allowed elevation as defined in Table 4-7 and the target on the upper end, and the halfway point between the minimum allowed elevation and the target on the lower end. As operations stray beyond these boundaries, a penalty in the form of a multiplier against the calculated revenue is applied. The penalty is an exponentially increasing function of the relative distance between the current elevation and the appropriate boundary of the no-penalty zone. The function is set to be modest early on, imposing a 20% reduction if the violation is 50% beyond

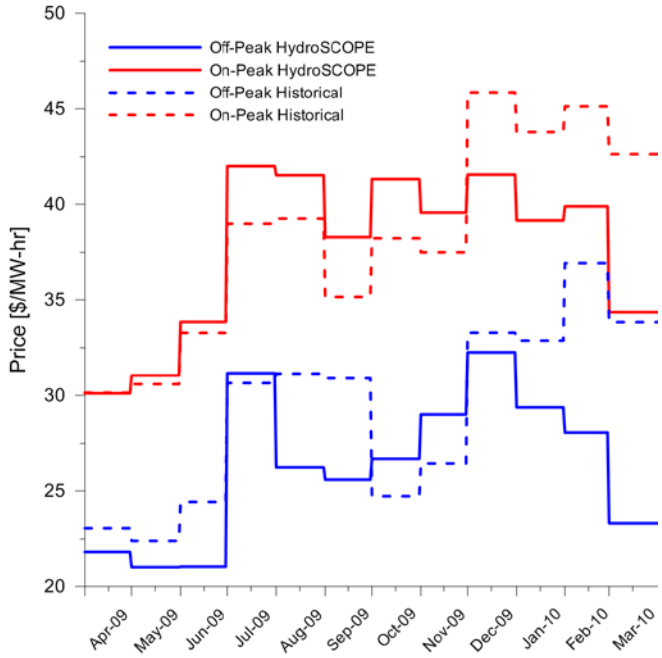

FIGURE 4-26 Historical Monthly Electricity Prices for April 2009 through March 2010 and the Average Monthly Prices Used in the Retrospective Analysis

the boundary, and a 50% reduction if the violation is 75% beyond the boundary (Figure 4-27). The total revenue for a simulation is calculated using equation (4-9) as:

$$R = \sum_{i=1}^{nt} \sum_{k=1}^{nres} r_{ik} p_{ik} \tag{4-9}$$

where *R* (dollars) is the total revenue for the simulation, *r* (dollars) is the revenue at time *i* for reservoir *k*, *p* (unitless) is the penalty function at time *i* for reservoir *k*, *nt* (unitless) is the number of timesteps in the simulations, and *nres* (unitless) is the number of power-producing reservoirs in the model.

Retrospective Analysis—Optimization Results. As discussed above, each day of the retrospective analysis is optimized using 500 simulations in which the baseline release schedule is adjusted to simultaneously maximize the total revenue, *R*, and the IRF score. DAKOTA outputs two files when using the MOGA option, one that lists the input multipliers and objective function values (i.e., R, and the IRF score) for all 500 simulations, and another that lists the Pareto-optimal solutions (i.e., the Pareto front). The second file is a subset of the first file. When HydroSCOPE is used in practice, it is intended that a user would manually pick the optimal solution from the Pareto front simulations based on user experience and what the user knows about the relative tradeoffs between the objectives. However, to automate the analysis process

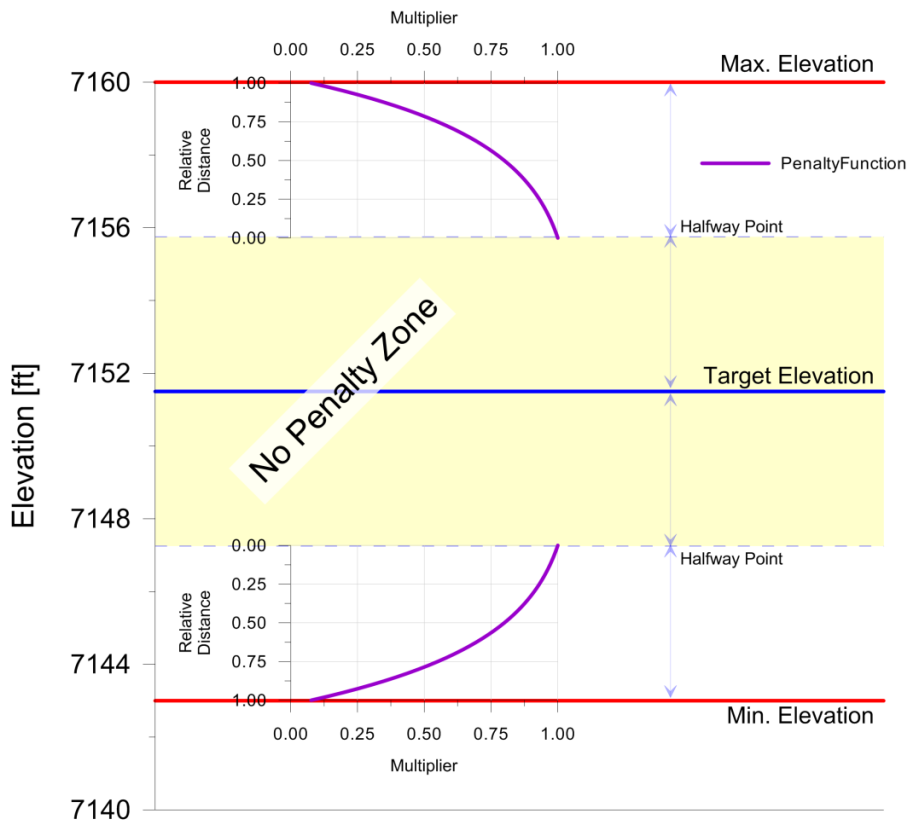

FIGURE 4-27 Schematic for Implementing the Penalty Function (The no-penalty zone is defined as the zone bounded by the halfway points between the maximum and minimum operating elevations and the target elevation. If the water surface elevation exceeds the halfway point, an exponential multiplier is applied to the calculated revenue based on the elevations relative distance from the halfway point to the maximum or minimum operating elevation. The elevations for the plot are based on the operating targets for Morrow Point from October through May.)

for the retrospective analysis, a single optimal solution was mathematically defined to be the solution with the lowest common ranking of the two objectives, where a ranking of 1 is the highest and 500 the lowest. Figure 4-28 shows plots of *R* versus the IRF score at the beginning of each month as well as at the end of the simulation. The red diamonds are the results of all 500 simulations, the blue points are the Pareto optimal solutions chosen by DAKOTA, the yellow point is the optimal solution chosen with the lowest common ranking method, and the green point is the baseline condition before optimization. Note that in all cases, DAKOTA finds release schedules that produce improvements in at least one of the objectives. With the exception of the May 1 example, the improvements came without negatively impacting the other objective. The May 1 example had the largest improvement in the IRF score (5.3%), but this came with a slight decline in total revenue (–1.9%). In this case, a user might chose to select another point on the May 1 Pareto front that has a smaller increase in IRF score and either a smaller decrease or an increase in total revenue. The values and percent differences are listed in Table 4-8.

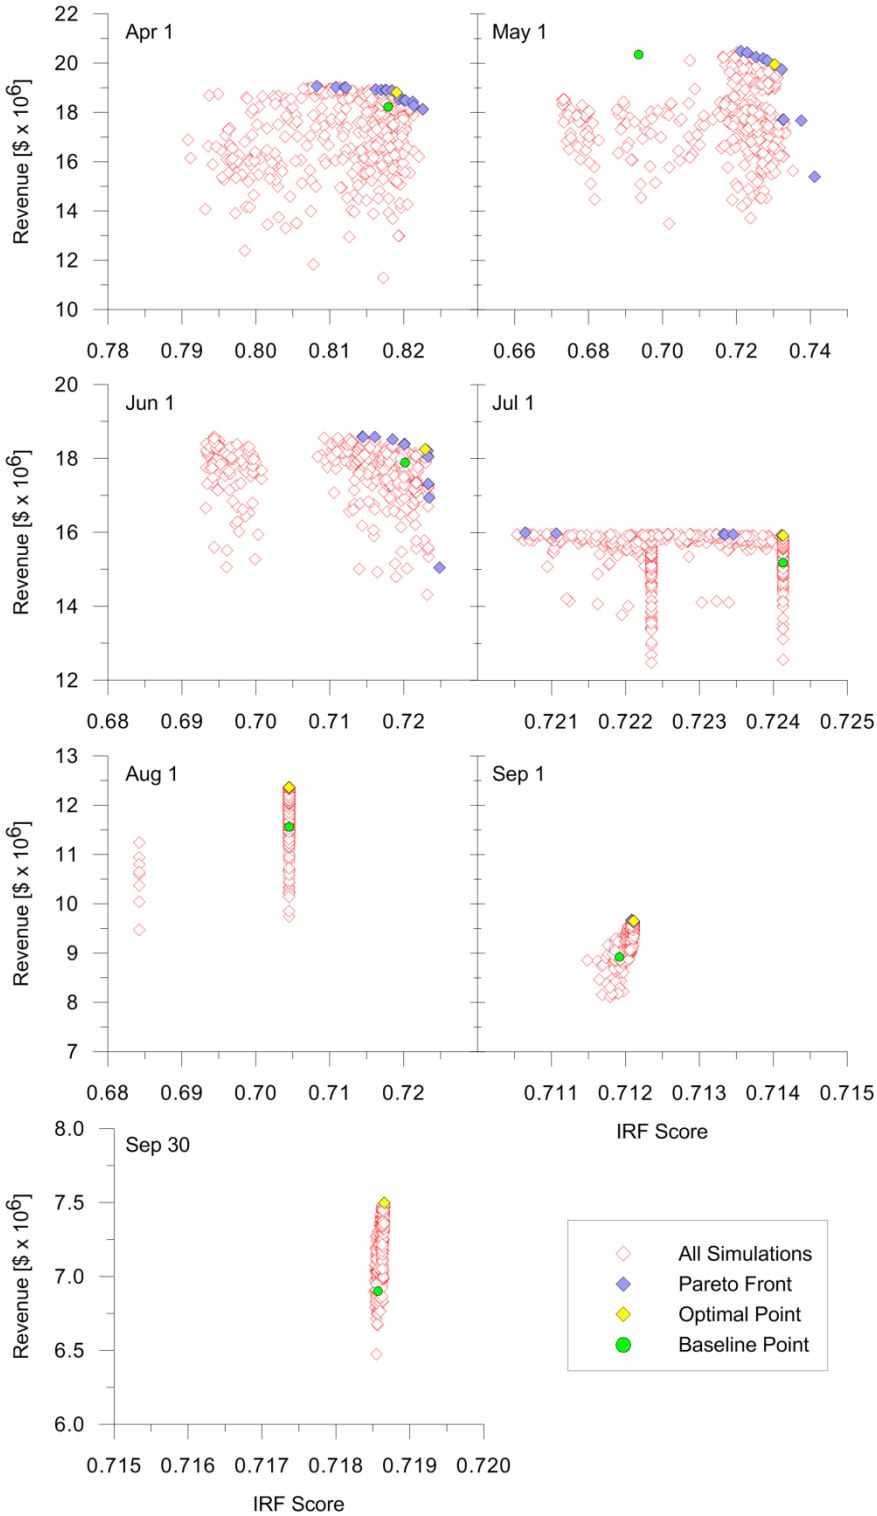

FIGURE 4-28 Plots of the Tradeoff between the Total Revenue and the IRF Score Produced by DAKOTA during an Optimization Run (The yellow point is the simulation that was chosen as optimal, and the green point is the baseline values before optimization. The blue points represent the Pareto front.)

TABLE 4-8 Baseline and Optimized Values and Percent Change for the IRF Score and Total Revenue[a]

	IRF Score			Total Revenue		
Date	Baseline	Optimal	% Change	Baseline	Optimal	% Change
April 1	0.818	0.819	0.138	$18,225,052.91	$18,813,307.62	3.228
May 1	0.694	0.730	5.292	$20,345,565.66	$19,953,981.34	–1.925
June 1	0.720	0.723	0.385	$17,886,503.27	$18,246,587.98	2.013
July 1	0.724	0.724	0.000	$15,183,538.73	$15,916,939.29	4.830
August 1	0.705	0.705	0.000	$11,561,481.52	$12,361,876.91	6.923
September 1	0.712	0.712	0.027	$8,918,083.69	$9,652,586.49	8.236
September 30	0.719	0.719	0.012	$6,901,302.69	$7,496,994.13	8.632
Average	0.727	0.733	0.803	$14,145,932.64	$14,634,610.54	3.455

[a] Note that the IRF score and total revenue are for the 6-month period following and including the listed date. The average is for the selected dates only.

Apparent in Figure 4-28 is the narrowing range in IRF score over the course of the analysis, especially during August and September. This is due to a combination of the lower flows in the rivers in late summer, fall, and early winter and the multiplier approach used by DAKOTA; smaller flows translate into a narrower range of variation that DAKOTA is able to explore. Groupings of the IRF score occur in the June, July, and August plots and are due to flow regimes that fall on either side on one or more IRF objective thresholds where the objective(s) either succeed or fail, producing a jump in the overall IRF score. Figures 4-29 and 4-30 show the range of flow variation produced by DAKOTA during the optimization process at the two environmental assessment points: the Gunnison Gauge and the Whitewater Gauge. The red line in each of the plots represents the optimal simulation and can be seen to range across the entire spectrum of flows. Note how the optimal solution can coincide with both the maximum and minimum of the flow range during a single optimization depending on the time of year. In addition, the distribution of flows sampled by DAKOTA at any one point in time during the simulation changes over time. An example of this is shown in the April 1 plot for the Whitewater Gauge (Figure 4-30) for the month of September, where the distribution of flows are skewed toward the minimum value.

Retrospective Analysis—Daily Results. To save disk space and reduce the simulation time, none of the model outputs are saved during the 500 simulation optimization step. Thus, for each day of the analysis, two post-simulations are run: Run B and Run C. Run B runs the model using the forecasted inflows and the chosen optimized releases. Run B results are written to a series of text files for post-processing and archiving. The Run C simulation uses the historical inflows (i.e., the perfect forecast) for the first day of the simulation, the forecasted inflows for the balance of the simulation, and the optimal releases for the entire simulation. Run C sets the initial conditions for the next day's simulation by assuming that at the start of the next day, the system

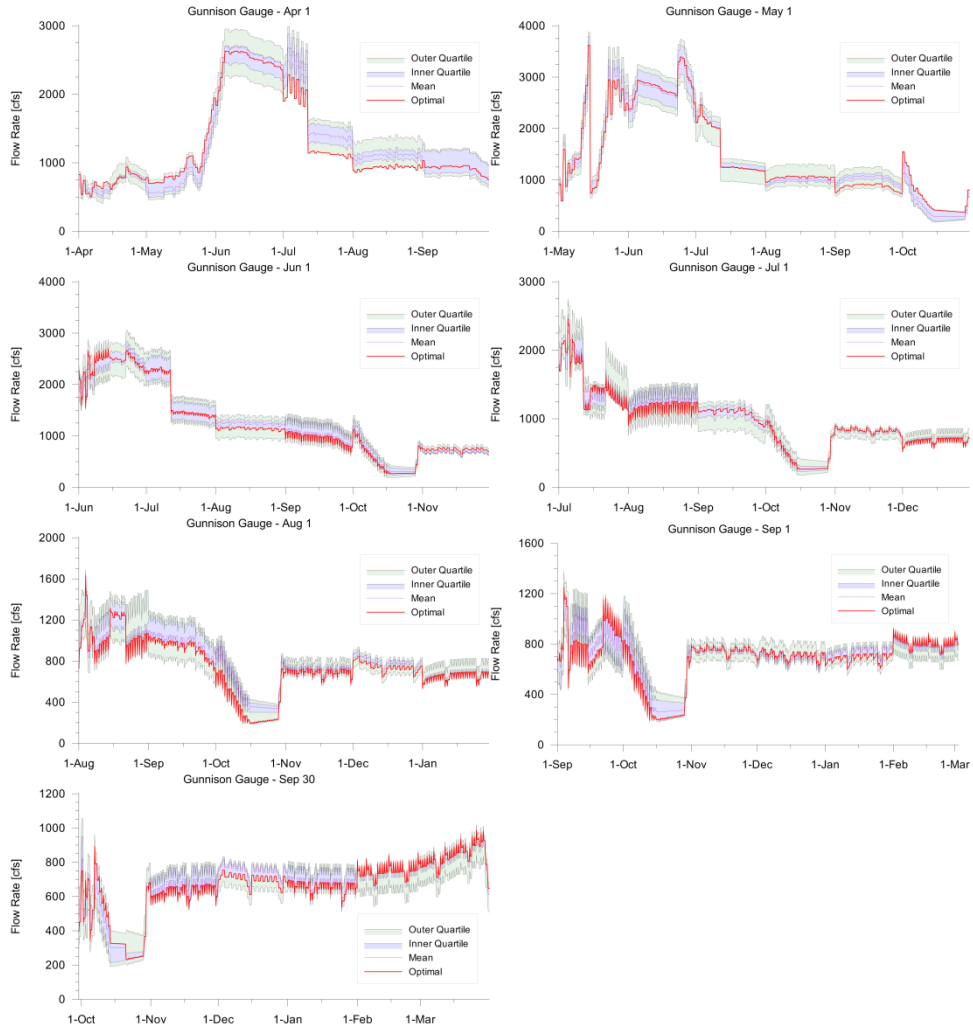

FIGURE 4-29 Plots Showing the Range of Variability of Flows Sampled by DAKOTA during the Optimization at the Gunnison Gauge (USGS #9128000) (The red line in each plot represents the flow at the gauge produced by the chosen, optimal solution.)

has experienced the actual inflows as opposed to the forecasted inflows. Run C results are saved to the database for use by the other downstream models.

Once the retrospective analysis is complete, a synthetic history is created for the entire six month period by concatenating the first days' results of each Run C simulation. The synthetic history contains the reservoir elevations and discharge rates, as well as the hydrographs at the two environmental assessment points that would have resulted if the optimized releases had been used under historical inflow conditions. It is the synthetic history of the daily release volumes that the day-ahead model shapes on an hourly basis to further optimize the total revenue.

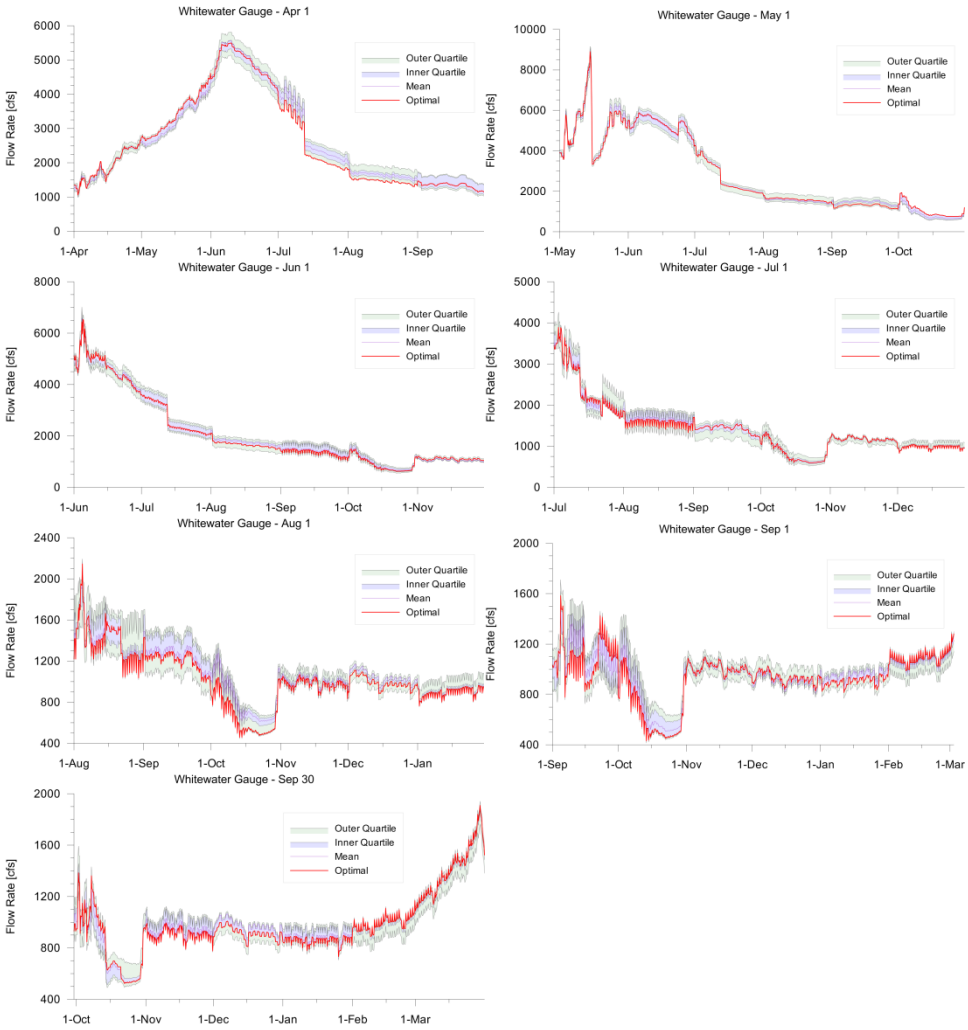

FIGURE 4-30 Plots Showing the Range of Variability of Flows Sampled by DAKOTA during the Optimization at the Whitewater Gauge (USGS #9152500) (The red line in each plot represents the flow at the gauge produced by the chosen, optimal solution.)

Figures 4-31 through 4-33 show the optimized and historical releases for Blue Mesa, Morrow Point, and Crystal reservoirs, respectively. The most striking difference between the optimized and historical releases is a shift in the high-flow peak that occurs in the historical operations on May 13 but is delayed until May 24 in the optimized case. The optimized peak also has a lower magnitude and a larger timespan than the historical case. According to the Bureau of Reclamation minutes for their April 2009 operation meeting (http://www.usbr.gov/uc/wcao/water/rsvrs/mtgs/pdfs/archives/am2009_04.pdf), the peak is based on the ability to meet the Black Canyon reserved water rights and the April 15 inflow forecast, and was timed to coincide with the Cimarron Creek peak. In 2009, the peak began on May 7 when releases from Crystal were increased daily until peaking at 7500 cfs on May 13. At this point, the releases were ramped down on a daily basis, starting on the morning of May 15, until leveling off at 2900 cfs

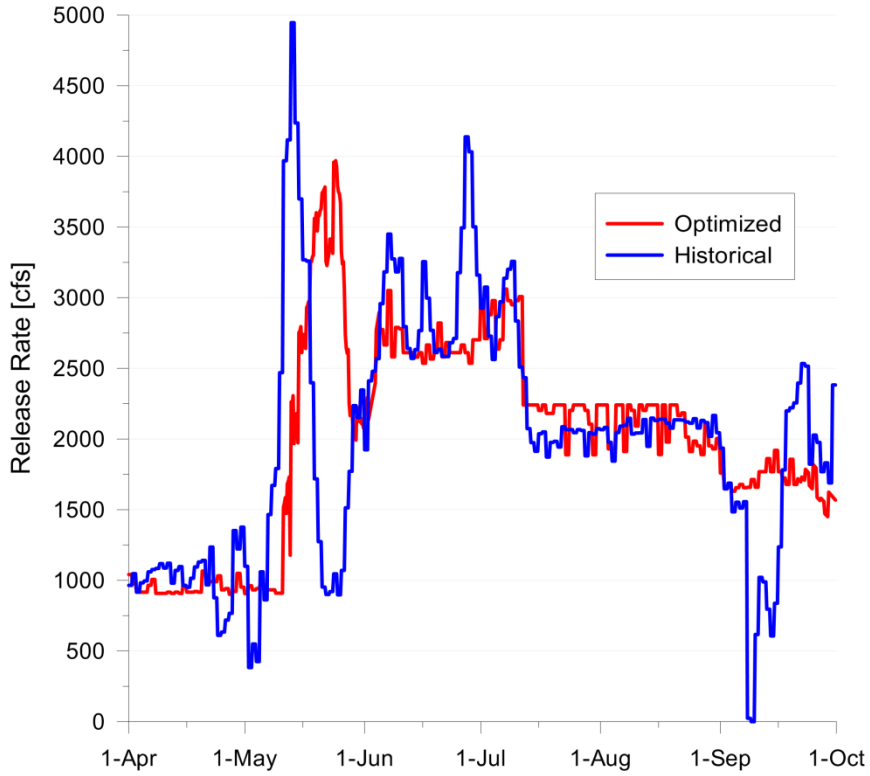

FIGURE 4-31 Optimized and Historical Releases for Blue Mesa Reservoir

on May 23. The peak for the optimal solution results from both the operational logic within the seasonal model and the optimization process whereby feedback from the IRF tool helped DAKOTA maximize the total revenue and IRF score by shifting and attenuating the springtime peak. The optimal and historical hydrographs at the two environmental assessment points as shown in Figures 4-34 and 4-35 show how the peak propagates downstream; these are the hydrographs used by the environmental tool to calculate the final IRF score. The ability to account for complex operational rules through the optimization process is a design element of the seasonal tool that is meant to increase DAKOTA's ability to explore solutions that would otherwise be missed. It should be noted that for this exercise, additional constraints on the springtime peak such as recreational opportunities and other competing factors were not considered.

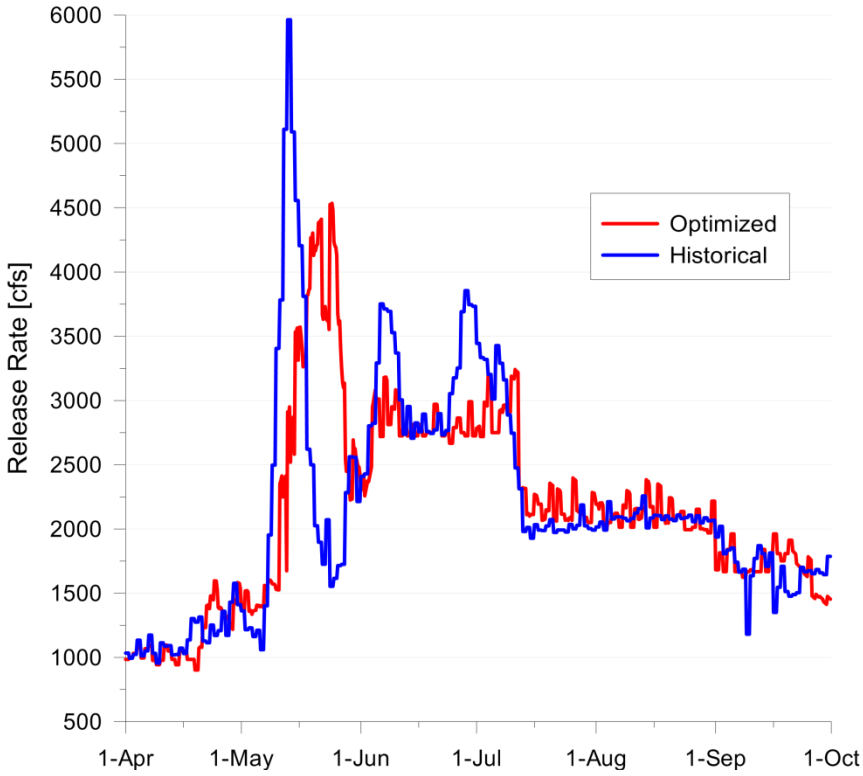

FIGURE 4-32 Optimized and Historical Releases for Morrow Point Reservoir

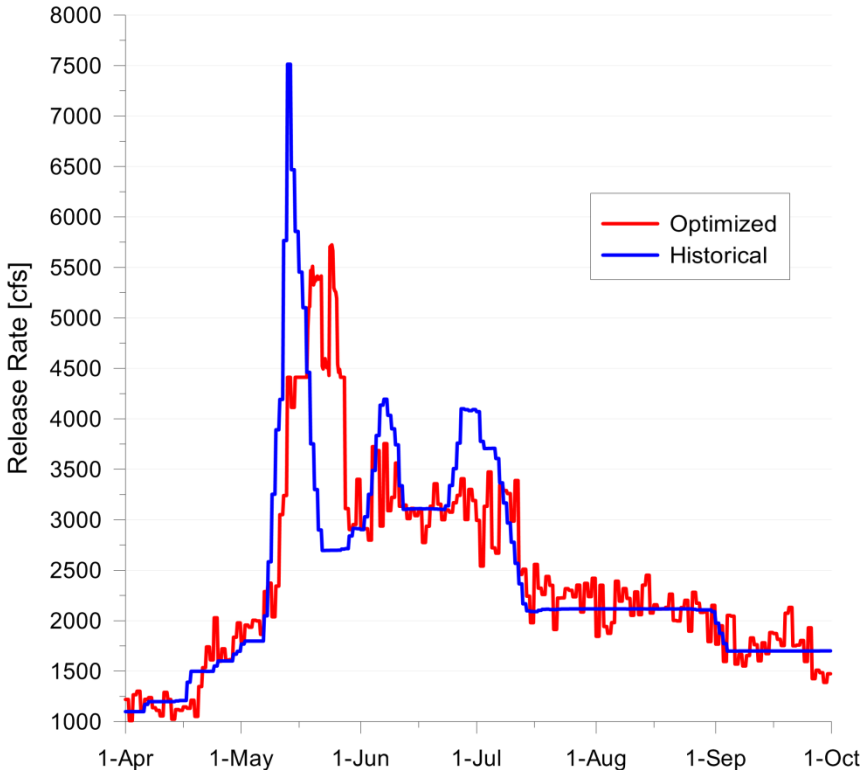

FIGURE 4-33 Optimized and Historical Releases for Crystal Reservoir

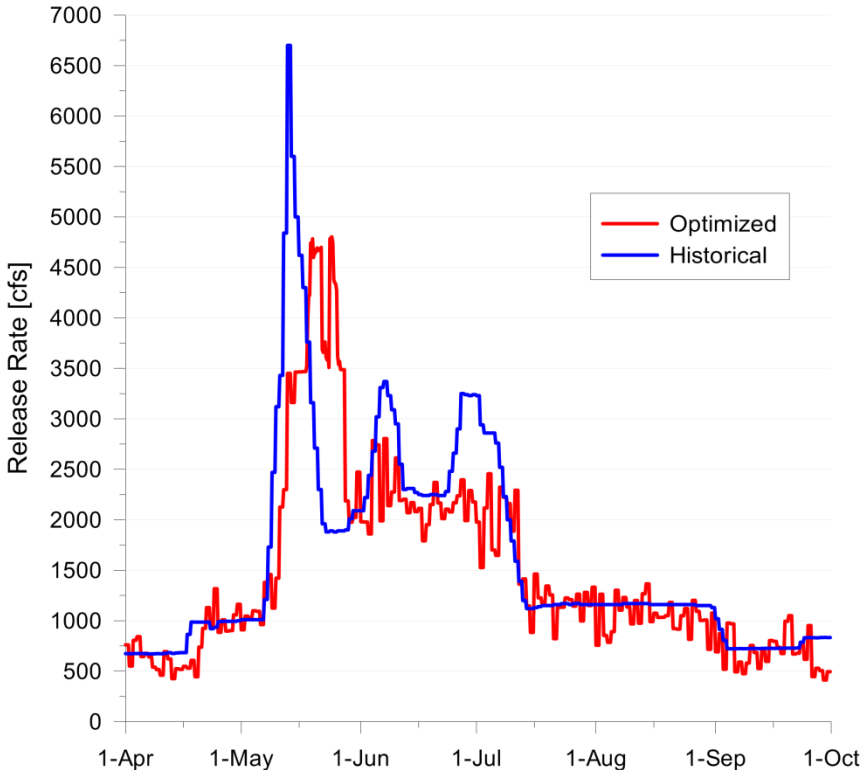

FIGURE 4-34 Optimized and Historical Releases at the Gunnison Gauge

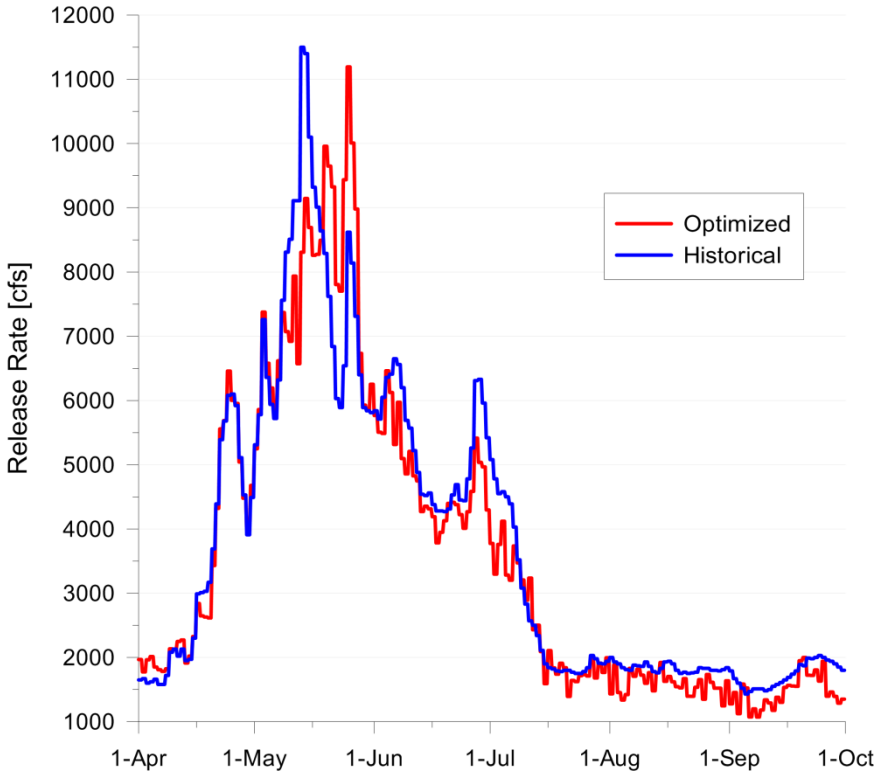

FIGURE 4-35 Optimized and Historical Releases at the Whitewater Gauge

The differences in the daily release volumes between the optimized and historical operations for the three reservoirs are shown in Figures 4-36 through 4-38. Because the daily optimized volumes from the seasonal model are products of a 6-month optimization process, they contain information about the projected system performance that is consistent with physical and regulatory requirements and accounts for the latest forecasted inflows. This information is passed on to the day-ahead planning model that uses those volumes to calculate an optimized hourly release schedule to maximize revenue production over the next 24 hours. This "telescopic" optimization approach means that seasonal objectives are inherently included in the daily optimization process, which helps keep the daily operations more in line with longer term objectives in light of the forecasted inflows. This approach is unique to the WUOT and is the primary reason behind its integrated design.

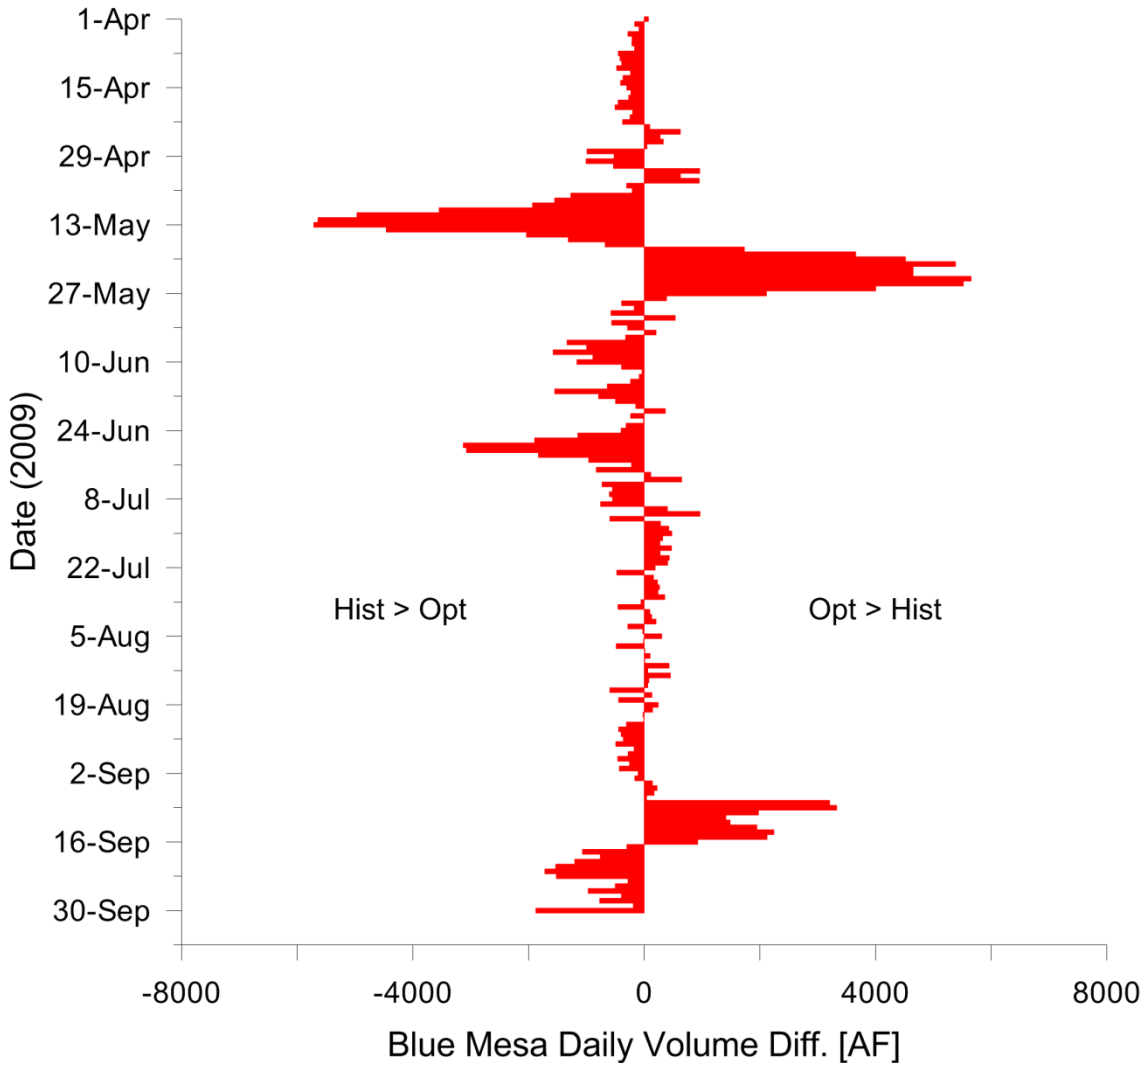

FIGURE 4-36 Differences in the Daily Volumes between the Optimized and Historical Operations for Blue Mesa Reservoir

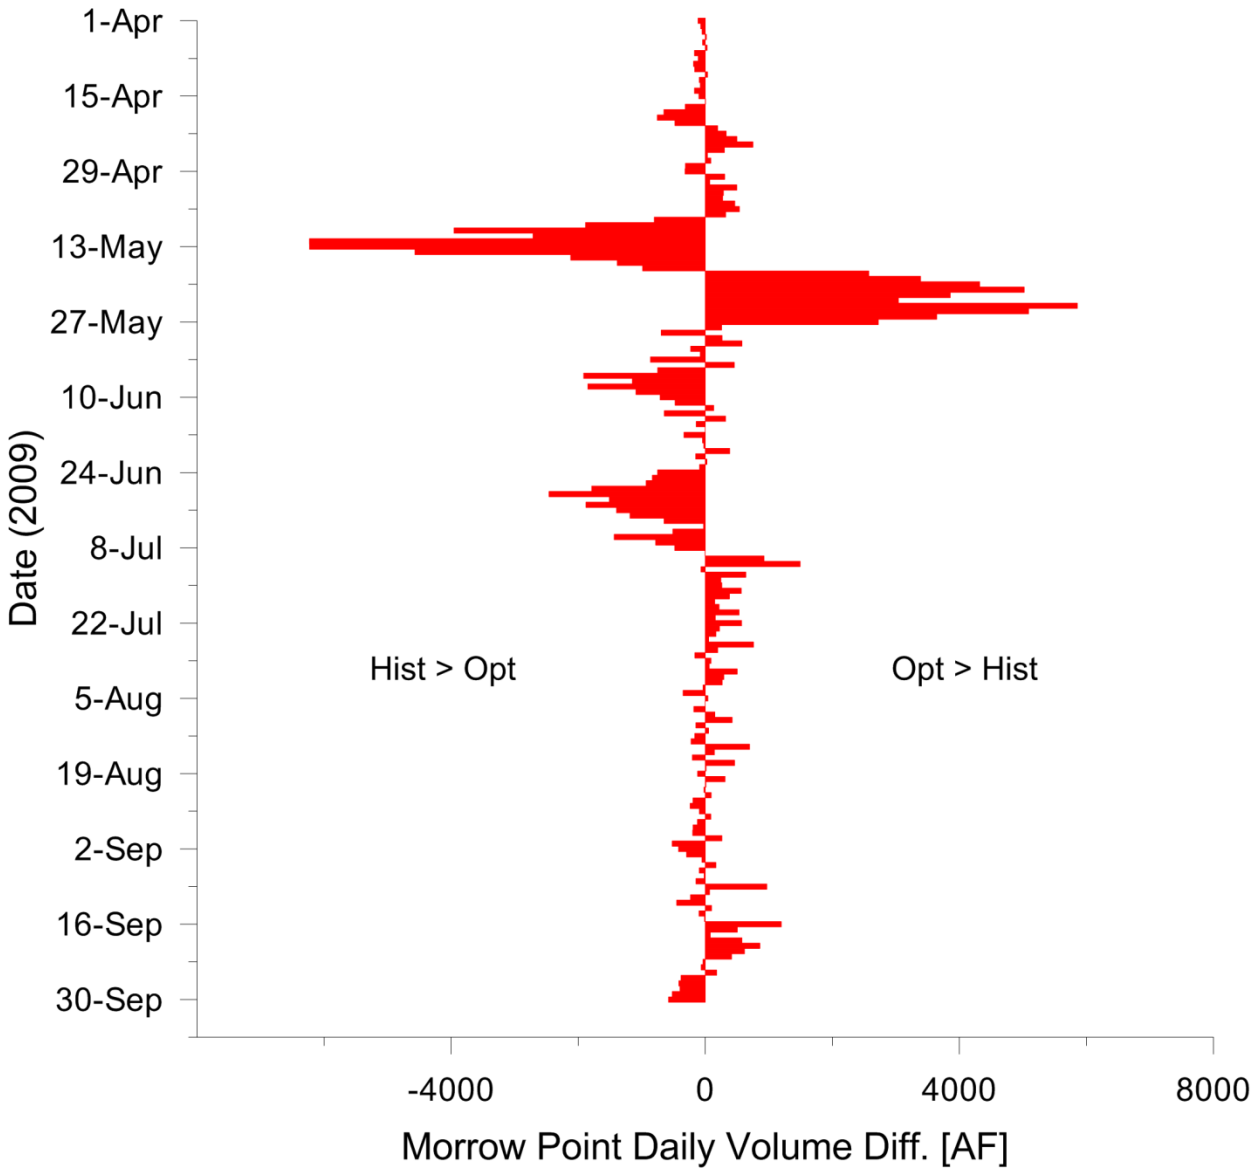

FIGURE 4-37 Differences in the Daily Volumes between the Optimized and Historical Operations for Morrow Point Reservoir

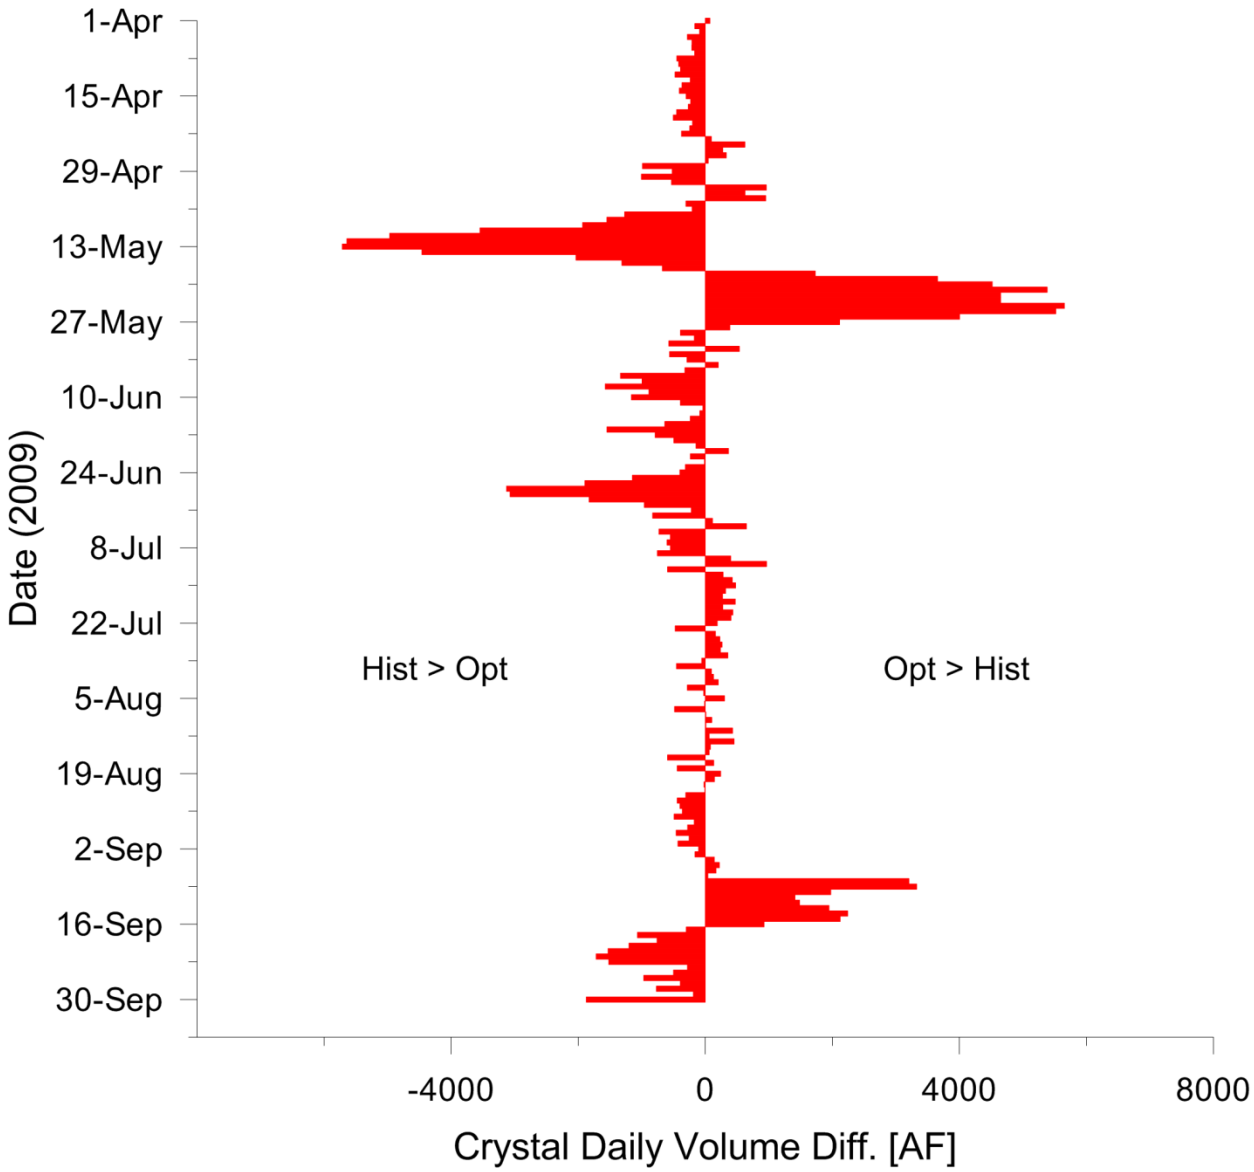

FIGURE 4-38 Differences in the Daily Volumes between the Optimized and Historical Operations for Crystal Reservoir

4.4.1.3 Environmental Performance

Description of Environmental Objectives for the Gunnison River. Objectives were defined for two locations on the Gunnison River:

1. At Whitewater, Colorado (meant to represent the reach between Roubideau Creek and the Redlands diversion dam, including the Escalante State Wildlife Area [ESWA]): Flow condition information for this location is based on readings from the USGS gauging station located near Grand Junction (USGS Gauge Station No. 09152500).

2. Below Crystal Dam (meant to represent the reach between the Gunnison Tunnel and confluence of the Gunnison River with the North Fork of the Gunnison River, including Black Canyon and Gunnison Gorge [BCGG]): Flow condition information for this location is based on readings from the USGS gauging station located below the Gunnison Tunnel (USGS Gauge Station No. 09128000).

We are aware of additional downstream flow restrictions that are in place for environmental reasons on the Colorado River at the Colorado-Utah state line, downstream of the confluence with the Gunnison River. However, those are not defined here because Location 1 (Redlands Diversion Dam) represents the downstream boundary of our network model and is located upstream of the confluence of the Gunnison River with the Colorado River. Other operational constraints, such as the prevention of flooding at Delta, Colorado, and Grand Junction, Colorado, are not included as part of the environmental objectives. For this exercise, all the objectives defined were considered to be equally important (i.e., relative weight of 1). Objectives related to flow requirements for endangered fish species at ESWA were adapted primarily from McAda (2003), and flow requirements for the two trout species at BCGG were adapted from Nehring and Miller (1987). Each of the environmental objectives is summarized in the following sections.

Environmental Objectives for Location 1—Escalante State Wildlife Area.

Objective 1.1: Substrate Maintenance for Endangered Fish Spawning
Relevant Period: May 1–June 30

Condition: The purpose of this objective is to provide appropriate spawning conditions for endangered fish, specifically within the interstitial spaces of gravel, where the eggs incubate and larvae spend their initial days after hatching. The magnitude was based on the known site-specific flow-stage relationship (McAda 2003). Mean hourly discharge should be at least 8,070 cfs for at least 1 full day within the relevant period. Maintaining this flow (1/2 bankfull) should provide adequate water velocities for achieving initial motion of sediments in order to remove fines from interstitial spaces of spawning substrate.

As defined, this objective is desired to be met in 7 out of 10 years over a long-term average (exceptions were made for dry and critically dry periods), but after 3 consecutive years of failure to address it, the score for this objective would go to zero.

Objective 1.2: Entrainment of Larval Fish into Floodplain Areas
Relevant Period: June 1–September 15

Condition: This objective pertains to improving the recruitment success of endangered fish species by providing suitable habitat conditions for their early life stages. The idea is to maximize floodplain habitat availability for fish larval development. The relationship between habitat suitability and the minimum daily discharge at the Whitewater gauge was adapted from the final environmental impact statement for operation of the Aspinall Unit (Reclamation 2012a) and is presented in Figure 4-39. Suitability values were represented as points based on a rating curve that identified the floodplain inundation extent corresponding to discharge levels in this segment of the river. Points would be accumulated every day as the suitability of minimum daily discharge values throughout the relevant period, based on the relationship depicted in Figure 4-39.

A minimum of three points should be accrued within the relevant period; however, 12 points would fully satisfy the objective. This objective is desired to be met (at least at its minimum level of three points) in 7 out of 10 years on a long-term average, but after 3 consecutive years of failure to address it, its score would go to zero.

Objective 1.3: Juvenile Fish Habitat Quality as Defined by Amount of Stage Change
Relevant Period: August 1–March 31

Condition: This objective pertains to improving the recruitment success of endangered fish species by providing habitat conditions suitable for their early life stages. The idea is to provide stable flows by limiting within-day fluctuations in stage, which provides suitable in-stream conditions for young-of-year during the base flow period, and would also minimize stranding of juveniles. The relationship between within-day changes in stage at the Whitewater gauge and habitat suitability (represented as points on maximum within-day stage change) is presented in Figure 4-40. Points would be accumulated every day based on daily changes in stage values at the Whitewater gauge throughout the relevant period, based on the relationship depicted in Figure 4-40.

As defined, a minimum of 235 points should be accrued within the relevant period. Assuming that one point could be achieved every day (i.e., within-day change in stage is restricted to less than 0.33 ft), this objective would be accomplished by maintaining the required condition for 235 days of the theoretical maximum of 243 days that are available within the mentioned period. This objective is desired to be met in 7 out of 10 years on a long-term average, but after 3 consecutive years of failure to address it, its score would go to zero.

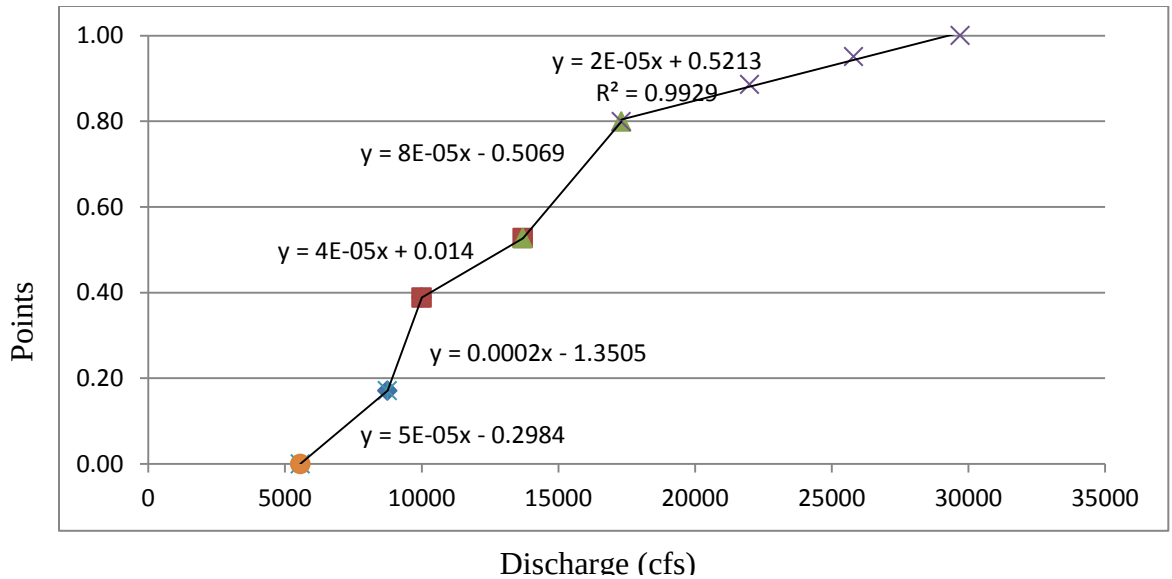

FIGURE 4-39 Relationship between Minimum Daily Discharge at the Whitewater Gauge and Suitability (as defined by points indicating portion of floodplain inundation) for Fish Larval Entrainment into Floodplain Areas in the Lower Gunnison River

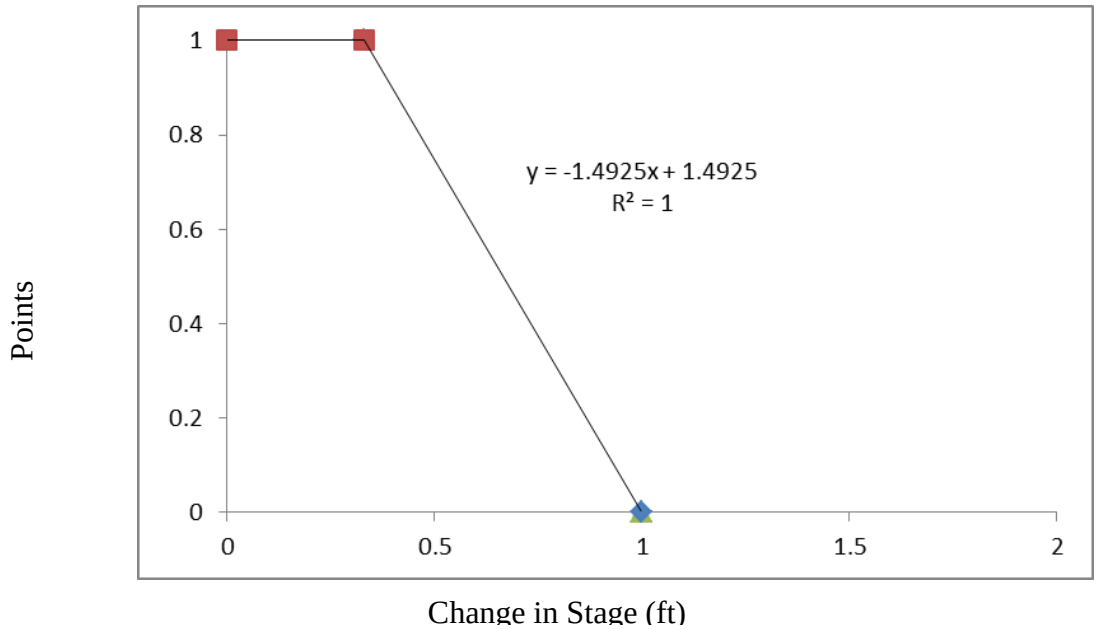

FIGURE 4-40 Relationship between Within-Day Change in Stage at Whitewater and Suitability (as defined by points indicating level of within-day change in stage) for Juvenile Fish in the Lower Gunnison River

Objective 1.4: Fish Habitat as Defined by Amount of Available Backwater during High Flow Season
Relevant Period: April 15–August 15

Condition: This objective pertains to improving the recruitment success of endangered fish species by providing backwater habitat along channel margins. Backwater habitat provides in-stream refugia during high flow periods and enhances growth. The idea is to enhance fish growth during the spring and summer months by providing the maximum possible amount of backwater habitat available in this segment of the river. The relationship between habitat suitability and the mean hourly discharge at the Whitewater gauge was adapted from Reclamation (2012a) and is presented in Figure 4-41. Suitability values were indicated as points and based on a rating curve that identified the potential backwater area availability corresponding to discharge levels in this segment of the river. Points would be accumulated every hour according to the suitability of the mean hourly discharge values measured at the Whitewater gauge throughout the relevant period, as depicted in Figure 4-41.

A minimum of 100 points should be accrued within the relevant period; however, 200 points would fully satisfy the objective. This objective is desired to be met (at least at its minimum level of 100 points) in 7 out of 10 years on a long-term average, but after 3 consecutive years of failure to address it, its score would go to zero.

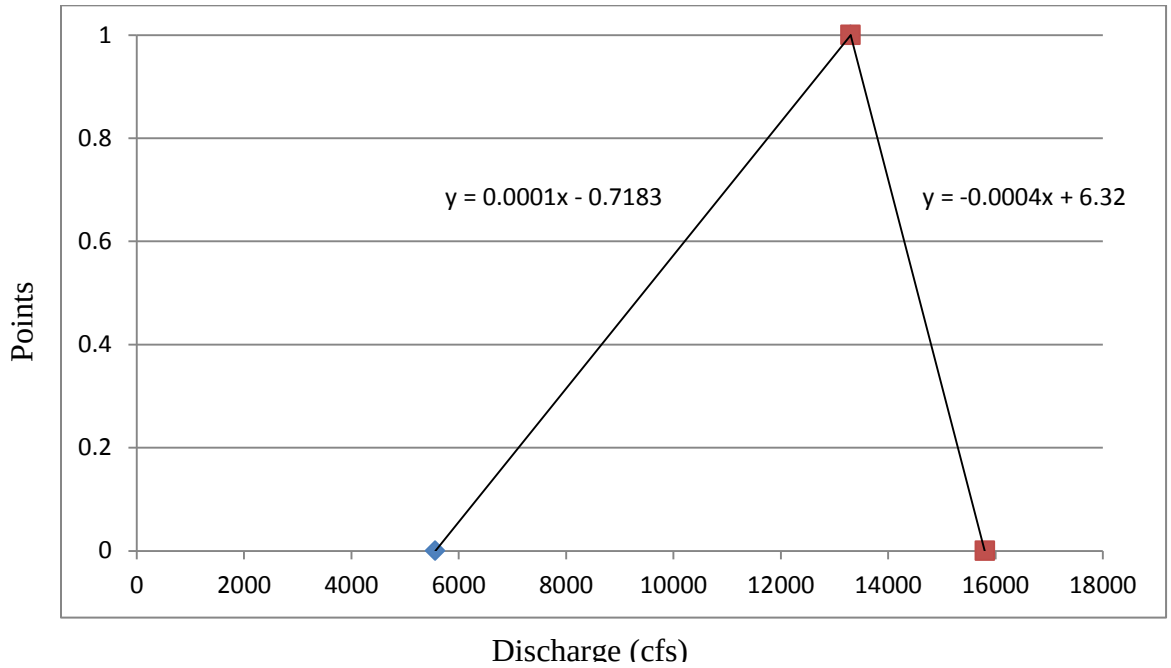

FIGURE 4-41 Relationship between Mean Hourly Discharge and Availability of Backwater Habitat for Endangered Fish in the Lower Gunnison River

Objective 1.5: Regeneration of Cottonwoods in Bottomland Habitat
Relevant Period: April 25–June 25

Condition: Minimum daily discharge should be at least 14,500 cfs as measured at the Whitewater gauge for at least 7 continuous days within the relevant period. Maintaining flows above the bank-full level would provide water on the floodplain for native hardwood (cottonwood/willow) germination and stand development, which in turn can provide nesting habitat for riparian birds and a future source of large, woody debris in the stream. Bank-full discharge also maintains channel form and prevents the encroachment of vegetation, including the non-native tamarisk.

As defined, this objective is desired to be met in 1 out of 10 years on a long-term average, but after 25 consecutive years of failure to address it, its score would go to zero.

Objective 1.6: Foraging Habitat for Riparian and Shoreline Birds
Relevant Period: March 15–July 31

Condition: An instantaneous discharge of at least 14,500 cfs at the Whitewater gauge for 48 instances (an instance is the smallest time step consisting of a 15-minute period) within the relevant period is required to achieve the minimum requirement; however, 96 instances would accomplish the objective to its maximum potential. Achieving the appropriate discharge levels (over bank-full) should provide water on the floodplain to submerge bottomland soil and organic matter and initiate development of a prey base (insects/worms) for riparian and shoreline birds (including the endangered flycatcher and cuckoo). Such pulse flows would also provide undergrowth for nesting and vegetation cover for chicks and fledglings.

As defined, this objective is desired to be met in 3 out of 10 years on a long-term average, but after 5 consecutive years of failure to address it, its score would go to zero.

Objective 1.7: Minimum Base Discharge (Yearly for 95% of Years)
Relevant Period: January 1–December 31

Condition: Minimum base-flow requirements were adopted from information provided in the Record of Decision (ROD) for Aspinall Unit Operations (Reclamation 2012b), and were based on minimum discharge for the period (the entire year). The objective requires that a minimum of 750 cfs (as measured at the Whitewater gauge) should be maintained throughout the year in 95 out of 100 years on a long-term average (relaxation made for extreme drought hydrological years); however, after 2 consecutive years of failure to address the objective, its score would go to zero. It should be noted that the frequency of extreme drought (critical dry) years is as specified by Reclamation (2012b) (forecasted April–July inflow into Blue Mesa Reservoir). This objective is meant to ensure an absolute year-round minimum discharge and to maintain ecological integrity of the river during all years, with an exception for critically dry (95% exceedance) years.

Objective 1.8: Minimum Base Discharge during High Flow Period
Relevant Period: June 1–August 31

Condition: Minimum base-flow requirements were adopted from Reclamation (2012b), based on minimum discharge for the high flow period. The objective requires that a minimum of 1,050 cfs be maintained at the Whitewater gauge from June 1 to August 31 in 95 out of 100 years on a long-term average (relaxation made for extreme drought hydrological years); however, after 2 consecutive years of failure to address the objective, its score would go to zero. Objective 1.8 imposes a more restrictive flow requirement than objective 1.7 during the high flow period of the year, and the two objectives in conjunction with one another create an additive effect of minimum flow restriction. The minimum discharge during the high flow season will allow the possibility of some level of reproduction by endangered fish in most years, except critically dry years.

Objective 1.9: Minimum Discharge during High Flow Period
Relevant Period: April 15–July 31

Condition: Minimum flow requirements were based on flow recommendations specifying minimum peak magnitude and duration for the high flow period. The relationship between instantaneous discharge at the Whitewater gauge and objective suitability is presented in Figure 4-42. As defined, an instantaneous discharge value greater than 2,600 cfs is required at the Whitewater gauge to accrue some points for this objective. A minimum of 5,000 points is required to achieve the minimum requirement, where a point could be accumulated for every instance (an instance is the smallest time step consisting of a 15-minute period) within the relevant period; however, 10,000 points (10,368 instances are available during the mentioned period) would accomplish the objective to its maximum potential. This objective is meant to provide a certain minimum level of flow for an adequate duration during the spring peak flow period to achieve geomorphic goals such as maintaining connections between the main channel and floodplains, reworking larger sediments on cobble bars, and reforming high-elevation sandbars. In this type of relationship, the duration has a degree of flexibility, since reduced duration can be compensated for with an increase in flow magnitude. As defined, this objective is desired to be met in 7 out of 10 years on a long-term average, but after 3 consecutive years of failure to address it, its score would go to zero.

Objective 1.10: Minimum Base Discharge (Spring)
Relevant Period: June 1–July 31

Condition: Minimum base-flow requirements were adopted from Reclamation (2012b), based on minimum discharge for the months of June and July. The objective requires that a minimum of 1,500 cfs be maintained throughout the period in 1 out of 2 years on a long-term average, but after 7 consecutive years of failure to address it, its score would go to zero. This level of flow is intended to ensure a minimum base discharge during the high flow season to provide suitable conditions for adult and sub-adult native endangered fish in at least half of the years.

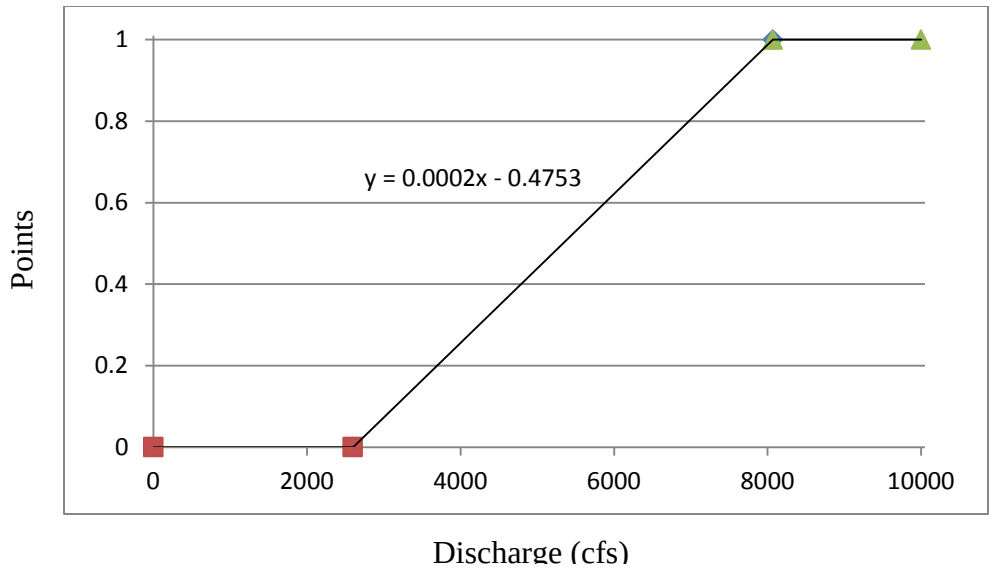

FIGURE 4-42 Relationship between Instantaneous Discharge and Suitability of Conditions as Indicated by Objective Points Associated with Minimum Peak Flow Requirements in the Lower Gunnison River

Objective 1.11: Minimum Base Discharge (Yearly for 70% of Years)
Relevant Period: January 1–December 31

Condition: Minimum base-flow requirements were adopted from Reclamation (2012b), based on minimum discharge for the entire year. The objective requires that a minimum of 1,050 cfs should be maintained throughout the year in 7 out of 10 years on a long-term average, but after 5 consecutive years of failure to address it, its score would go to zero. This level of flow ensures a minimum year-round base discharge for native endangered fish.

Objective 1.12: Minimum Peak Discharge (1/2 Bank-full)
Relevant Period: May 15–June 15

Condition: Minimum peak flow requirements were adopted from Reclamation (2012a), based on maximum daily instantaneous discharge values for the identified period during the spring high flow season. McAda (2003) recommends an instantaneous daily peak flow of at least 8,070 cfs for at least 10 days in 70% of the years in this stretch of the Gunnison River. As defined, an instantaneous peak discharge (based on daily maximum) of at least 8,070 cfs at the Whitewater gauge is needed for a minimum of 10 days within the defined period is required to accomplish this objective. It should be accomplished in 7 out of 10 years on a long-term average, but after 5 consecutive years of failure to address the objective, its score would go to zero. The peak flow for this objective is similar to that of objective 1.1; however, in addition to simply providing spawning conditions for the endangered fish, the prolonged duration will promote spawning and maintain suitable conditions for incubation of eggs and development of larvae within the interstices of the gravel substrate.

Objective 1.13: Minimum Peak Discharge (Bank-full)
Relevant Period: May 15–June 15

Condition: Minimum peak flow requirements were adopted from Reclamation (2012a), based on maximum daily instantaneous discharge values for the entire year. As defined, the objective requires an instantaneous peak discharge (based on daily maximum value) of 14,350 cfs for a minimum of 2 days within the relevant period. It should be accomplished in 1 out of 2 years on a long-term average, but after 4 consecutive years of failure to address it, its score would go to zero. This magnitude of flow will help maintain channel complexity and associated habitat heterogeneity. The definition for this objective was based on information available in Reclamation (2012a) and McAda (2003).

Environmental Objectives for Location 2—Black Canyon and Gunnison Gorge.

Objective 2.1: Prepare Trout Spawning Habitat (Remove Fines)
Relevant Period: May 20–June 10

Condition: Minimum daily discharge between 2,000 and 4,000 cfs for at least 1 day within the relevant period accomplishes the objective by flushing out fines from the interstitial spaces of gravel, where the trout eggs incubate and larvae spend their initial days after hatching. This objective is desired to be met in 1 out of 2 years on a long-term average, but after 10 consecutive years of failure to address it, its score would go to zero. A longer frequency limit was set for trout objectives, because they are a part of a managed recreational fishery subject to supplemental stocking.

Objective 2.2: Provide Rainbow Trout Spawning Habitat
Relevant Period: April 1–May 15

Condition: The relationship between habitat suitability and discharge is adapted from information provided in Reclamation (2012a) and depicted in Figure 4-43. Objective points would be accumulated every hour during the relevant period based on the suitability of the mean hourly discharge. A minimum of 700 points would be required within the relevant period to minimally accomplish the objective. Assuming that one point could be achieved every hour (i.e., mean hourly discharge is between 500 and 600 cfs), this objective will be minimally met by maintaining the required condition for 700 hours; however, maintaining the condition for 1,000 hours of the theoretical maximum of 1,080 hours that are available within the mentioned period accomplishes the objective to its full potential. This objective is expected to provide ideal spawning habitat for rainbow trout and is desired to be met (at least to its minimum level of 700 points) in 1 out of 2 years on a long-term average, but after 10 consecutive years of failure to address it, its score would go to zero.

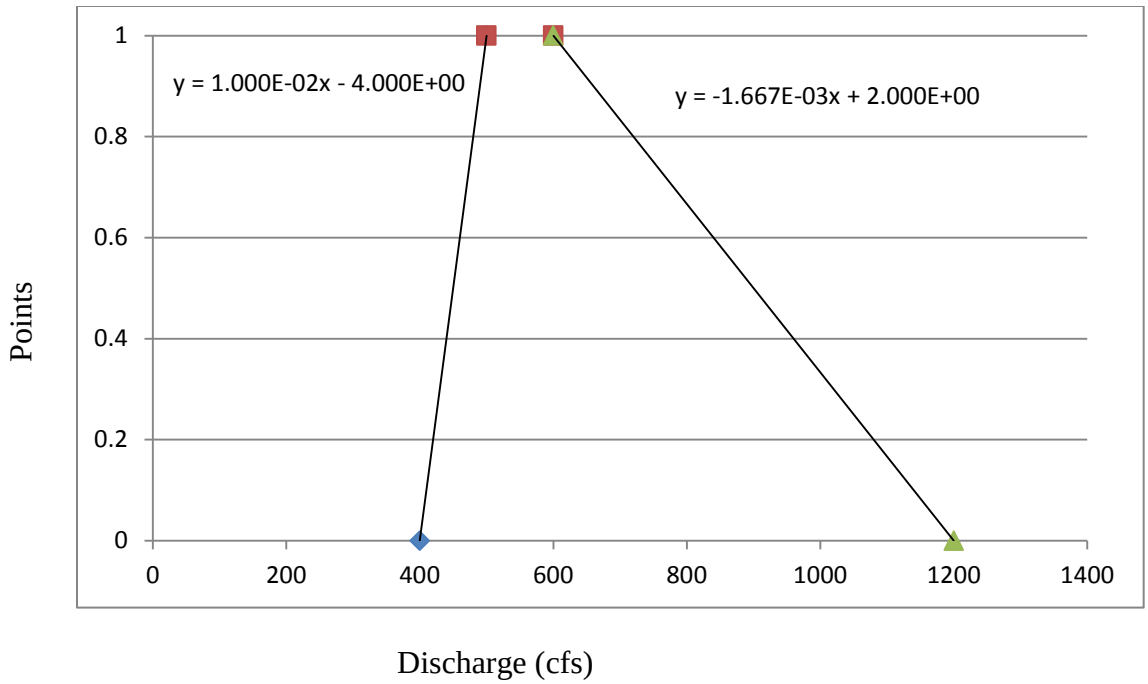

FIGURE 4-43 Relationship between Mean Hourly Discharge and Availability of Trout Spawning Habitat for Endangered Fish in the Black Canyon and Gunnison Gorge Reaches of the Gunnison River

Objective 2.3: Provide Brown Trout Spawning Habitat
Relevant Period: October 15–November 15

Condition: The relationship between habitat suitability and discharge for brown trout is the same as that for the rainbow trout, as depicted in Figure 4-43, although the relevant timing differs due to differences in spawning seasonality. Objective points would be accumulated every hour during the relevant period, based on mean hourly discharge. A minimum of 500 points would be required within the relevant period to accomplish the objective. Assuming that one point could be achieved every hour (i.e., mean hourly discharge is between 500 and 600 cfs), this objective will be minimally met by maintaining the required condition for 500 hours; however, maintaining the condition for 700 hours of the theoretical maximum of 768 hours that are available within the mentioned period would accomplish the objective to its full potential. This objective is anticipated to provide appropriate spawning habitat for brown trout and is desired to be met (at least to its minimum level of 500 points) in 1 out of 2 years on a long-term average, but after 10 consecutive years of failure to address it, its score would go to zero.

Objective 2.4: Provide Preferred Habitat Conditions for Adult Trout
Relevant Period: January 1–December 31

Condition: At least 10,000 instances of instantaneous (smallest possible time step, i.e., 15 minutes) discharge between 400 and 1,200 cfs within the relevant period (the theoretical maximum possible is 35,040 instances). Improvements beyond the minimum requirement could

be achieved by accruing points for up to 20,000 instances within the relevant period. This objective is desired to be met (at least to its minimum level of 10,000 instances) in 8 out of 10 years on a long-term average, but after 5 consecutive years of failure to address it, its score would go to zero.

Objective 2.5: Ramp-rate Considerations for Rainbow Trout Recruitment
Relevant Period: June 10–July 1

Condition: Maximum within-day change in discharge should be less than 250 cfs within the relevant period to accomplish this objective. A minimum of 17 days within the relevant period is required for this objective; however, 22 days of meeting the condition fully satisfies the objective. This objective is desired to be met (at least to its minimum level of 17 days) in 8 out of 10 years on a long-term average, but after 5 consecutive years of failure to address it, its score would go to zero.

Objective 2.6: Maximum Flow Considerations for Rainbow Trout Recruitment
Relevant Period: June 10–July 1

Condition: Maximum daily discharge has to be less than 3,500 cfs within the relevant period to accomplish this objective. A minimum of 20 days within the relevant period is required for this objective; however, 22 days of meeting the condition fully satisfies the objective. This objective is desired to be met (at least to its minimum level of 20 days) in 8 out of 10 years on a long-term average, but after 5 consecutive years of failure to address it, its score would go to zero.

Objective 2.7: Ramp-rate Considerations for Brown Trout Recruitment
Relevant Period: April 15–June 1

Condition: Maximum within-day change in discharge has to be less than 250 cfs within the relevant period to accomplish this objective. A minimum of 36 days within the relevant period is required for this objective; however, 48 days of meeting the condition fully satisfies the objective. This objective is desired to be met (at least to its minimum level of 36 days) in 8 out of 10 years on a long-term average, but after 5 consecutive years of failure to address it, its score would go to zero.

Objective 2.8: Maximum Flow Considerations for Brown Trout Recruitment
Relevant Period: April 15–June 1

Condition: Maximum daily discharge has to be less than 3,500 cfs within the relevant period to accomplish this objective. A minimum of 44 days within the relevant period is required for this objective; however, 48 days of meeting the condition fully satisfies the objective. This objective is desired to be met (at least to its minimum level of 44 days) in 8 out of 10 years on a long-term average, but after 5 consecutive years of failure to address it, its score would go to zero. The specified year-round base flow will provide the minimum in-stream habitat conditions for adult and growing juvenile trout.

Objective 2.9: Minimum Base Flow
Relevant Period: January 1–December 31

Condition: Minimum instantaneous flow has to be at least 300 cfs below the Gunnison tunnel gauge at any given time of the year. A slight relaxation of 3 days (i.e., 362 days out of 365 days in the year) has been made to accommodate for unexpected emergency and accidental maintenance issues at the dam. Minimum base flow requirements were adopted from Reclamation (2012b) for the Aspinall Unit Operations, based on minimum discharge for the period (the entire year). As defined, this objective should be met in 95 out of 100 years on a long-term average, but after 2 consecutive years of failure to address it, its score would go to zero. The specified year-round base flow will provide the minimum in-stream habitat conditions for the adult and growing juvenile trout.

IRF Scores for the Gunnison River Using the IRF-Historical Tool in Standalone Mode. IRF scores were calculated using the IRF historic mode and the objective scores for the two Gunnison River locations (described in the previous sections) were calculated (Figure 4-44) using the annual hydrographs reported from 1993 to 2008 at the gauging stations for each location (Figure 4-45). The overall annual historic IRF scores for the Gunnison River from 1993 to 2008 are also shown in Figure 4-44. Annual historic scores for each individual objective are presented in Figures 4-46 and 4-47 for locations 1 and 2, respectively.

Annual overall environmental performance score (IRF score) varied from 0.38 to 0.52 during the 16-year time period (Figure 4-44). The annual IRF score showed an increase in trend from 1994 to 1997, but has generally been stable since 1997. The increase in IRF score was primarily due to alternate high flow years between 1993 and 1997 (Figure 4-45), when several of the poorly performing objectives were accomplished (Figure 4-46). A clear difference in environmental performance score was noticed between the two sites. Site scores for the ESWA site were approximately half those of the BCGG site for most of the years within the 16-year period (Figure 4-44). These results suggest that the conditions were more favorable for the trout that predominantly occur in the upper BCGG reach than for the endangered native fish known to occur in the lower ESWA site. Field surveys from this period also indicate such a trend at these sites. Overall trout biomass, which is an indicator of the population abundance, has been at an all-time best; the endangered fish species barely show any improvement in their standing stocks. This is also evident from the objective scores for the two sites. For example, scores for objectives 1.2 (entrainment of the larval endangered fish) and 1.9 (peak spring-flow magnitude) at the ESWA site were approximately 25% of their desired levels for the last 8 of the 16 years (Figure 4-46). In contrast, at the BCGG site, scores for six of the nine objectives were at least 75% of their desired levels for the last 7 of the 16 years (Figure 4-47).

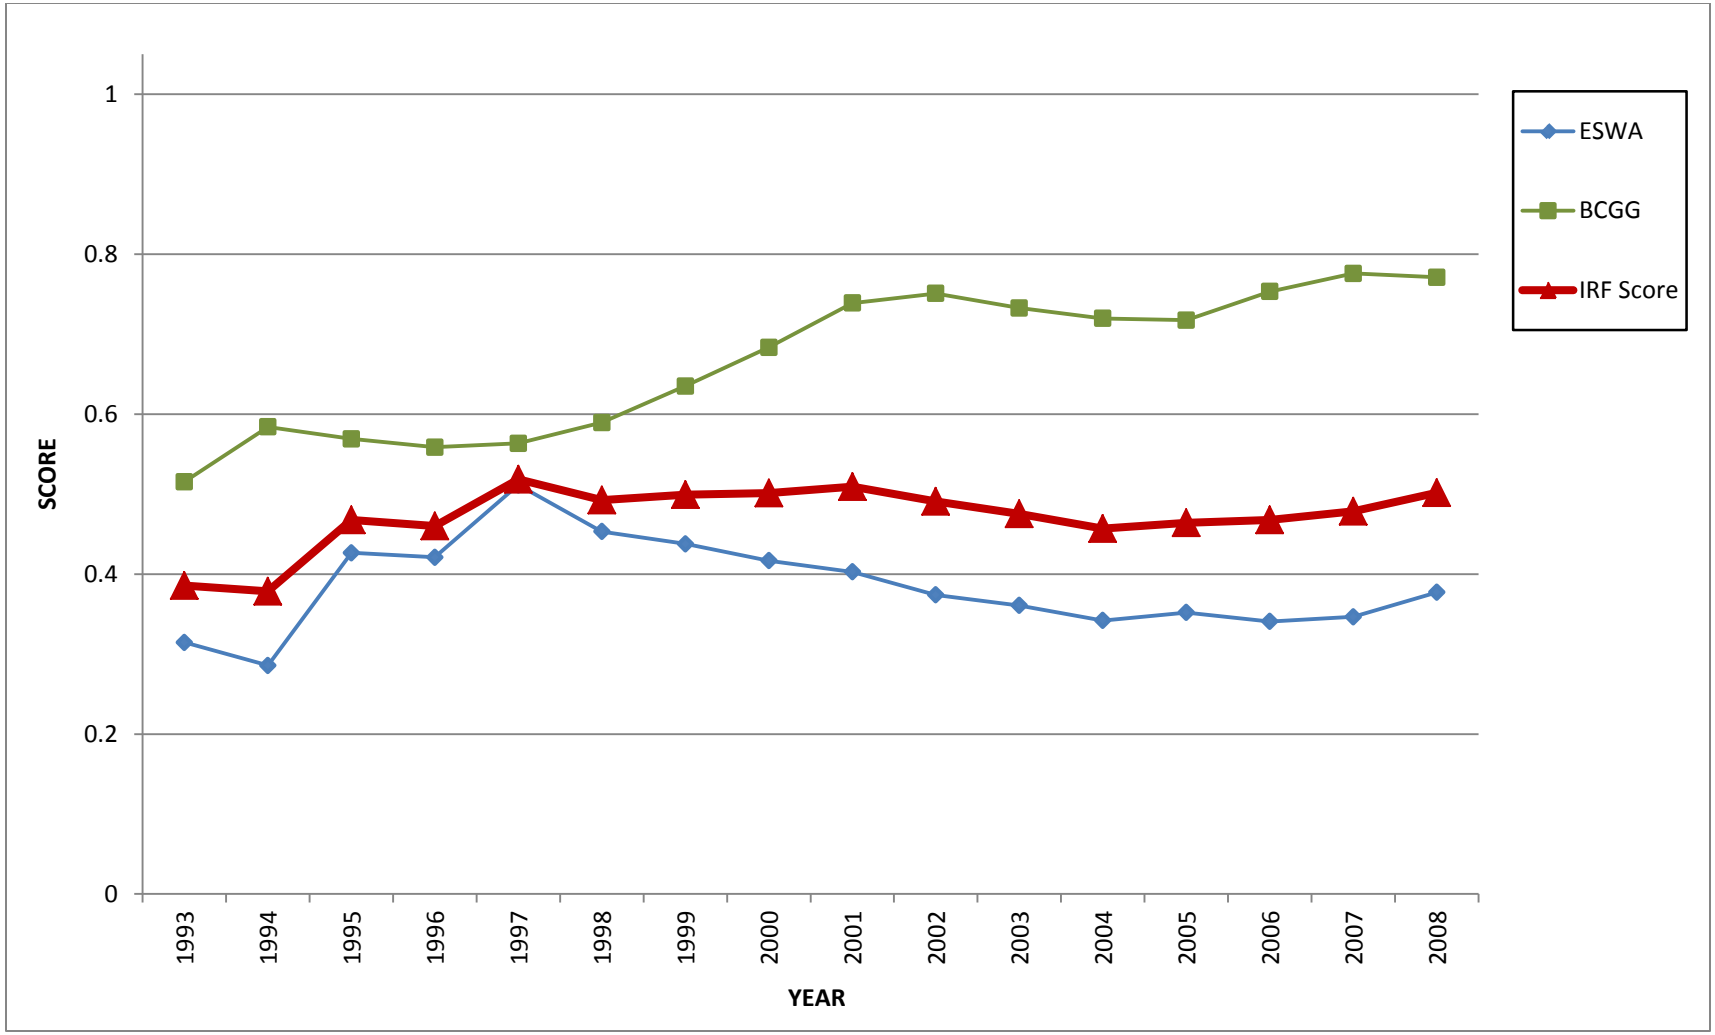

FIGURE 4-44 Objective Environmental Performance Score (IRF Score) and Site Scores for a 16-year Historic Period (1993–2008) for the Gunnison River Using the IRF Historic Tool (The overall historic annual IRF scores for each year are indicated by red triangles.)

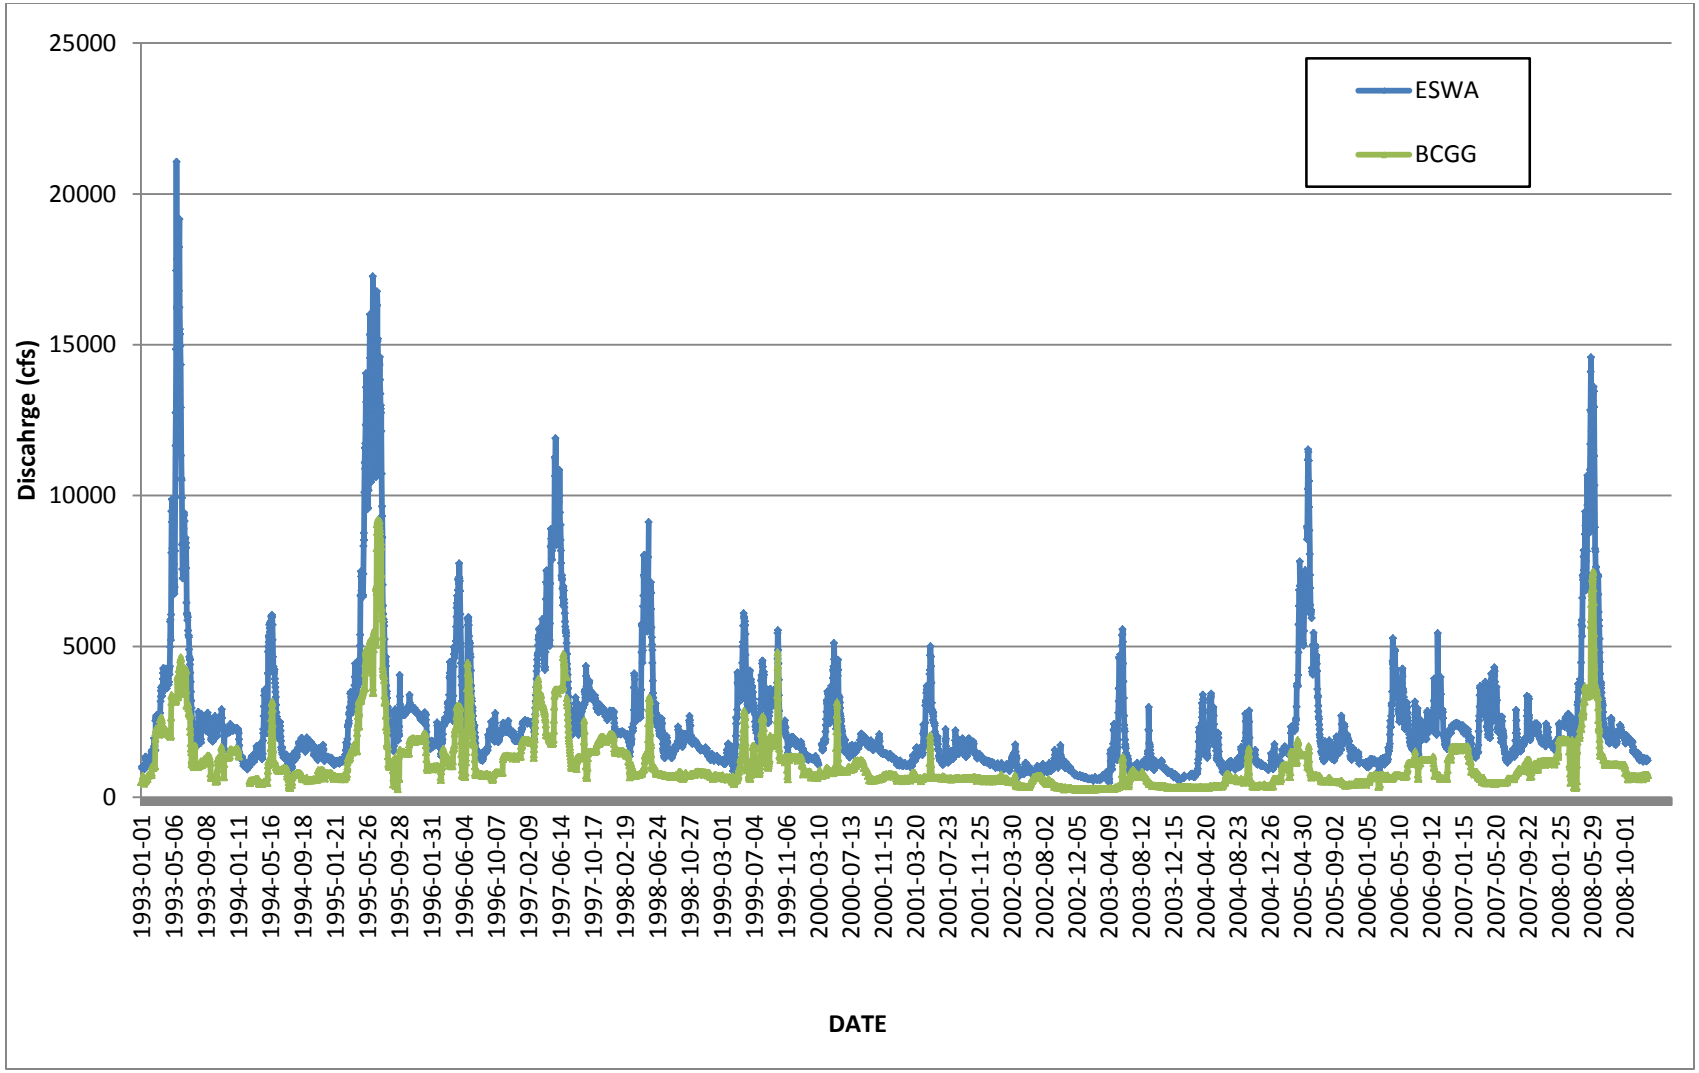

FIGURE 4-45 Historic Hydrographs for the Two Sites in the Gunnison River during a 16-year Period (1993–2008) (Flow data for the ESWA site was recorded by the USGS Whitewater gauge, and flow data for the BCGG site was recorded by the USGS site below Gunnison Tunnel.)

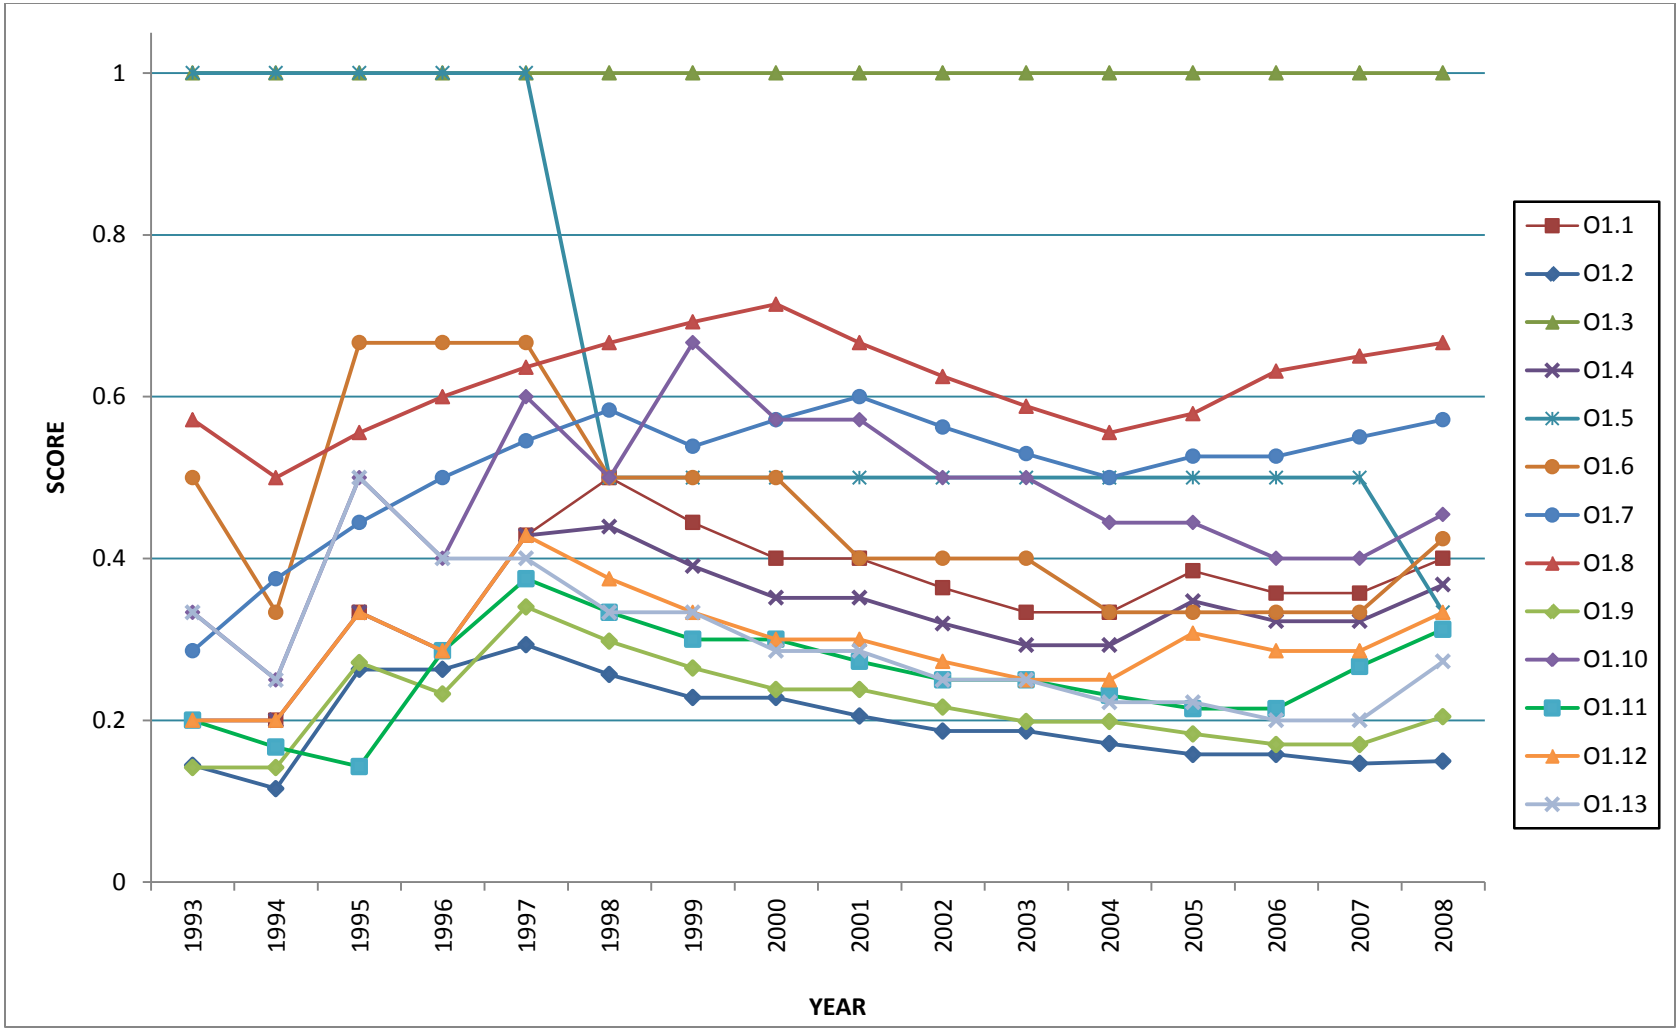

FIGURE 4-46 Objective Performance Score for Objectives 1.1–1.13 for a 16-year Period (1993–2008) for ESWA Site in the Gunnison River Using the IRF Historic Tool (The relevant objectives [labeled O1.1–O1.13 in the graph] are described in detail in Section 4.4.1.3.)

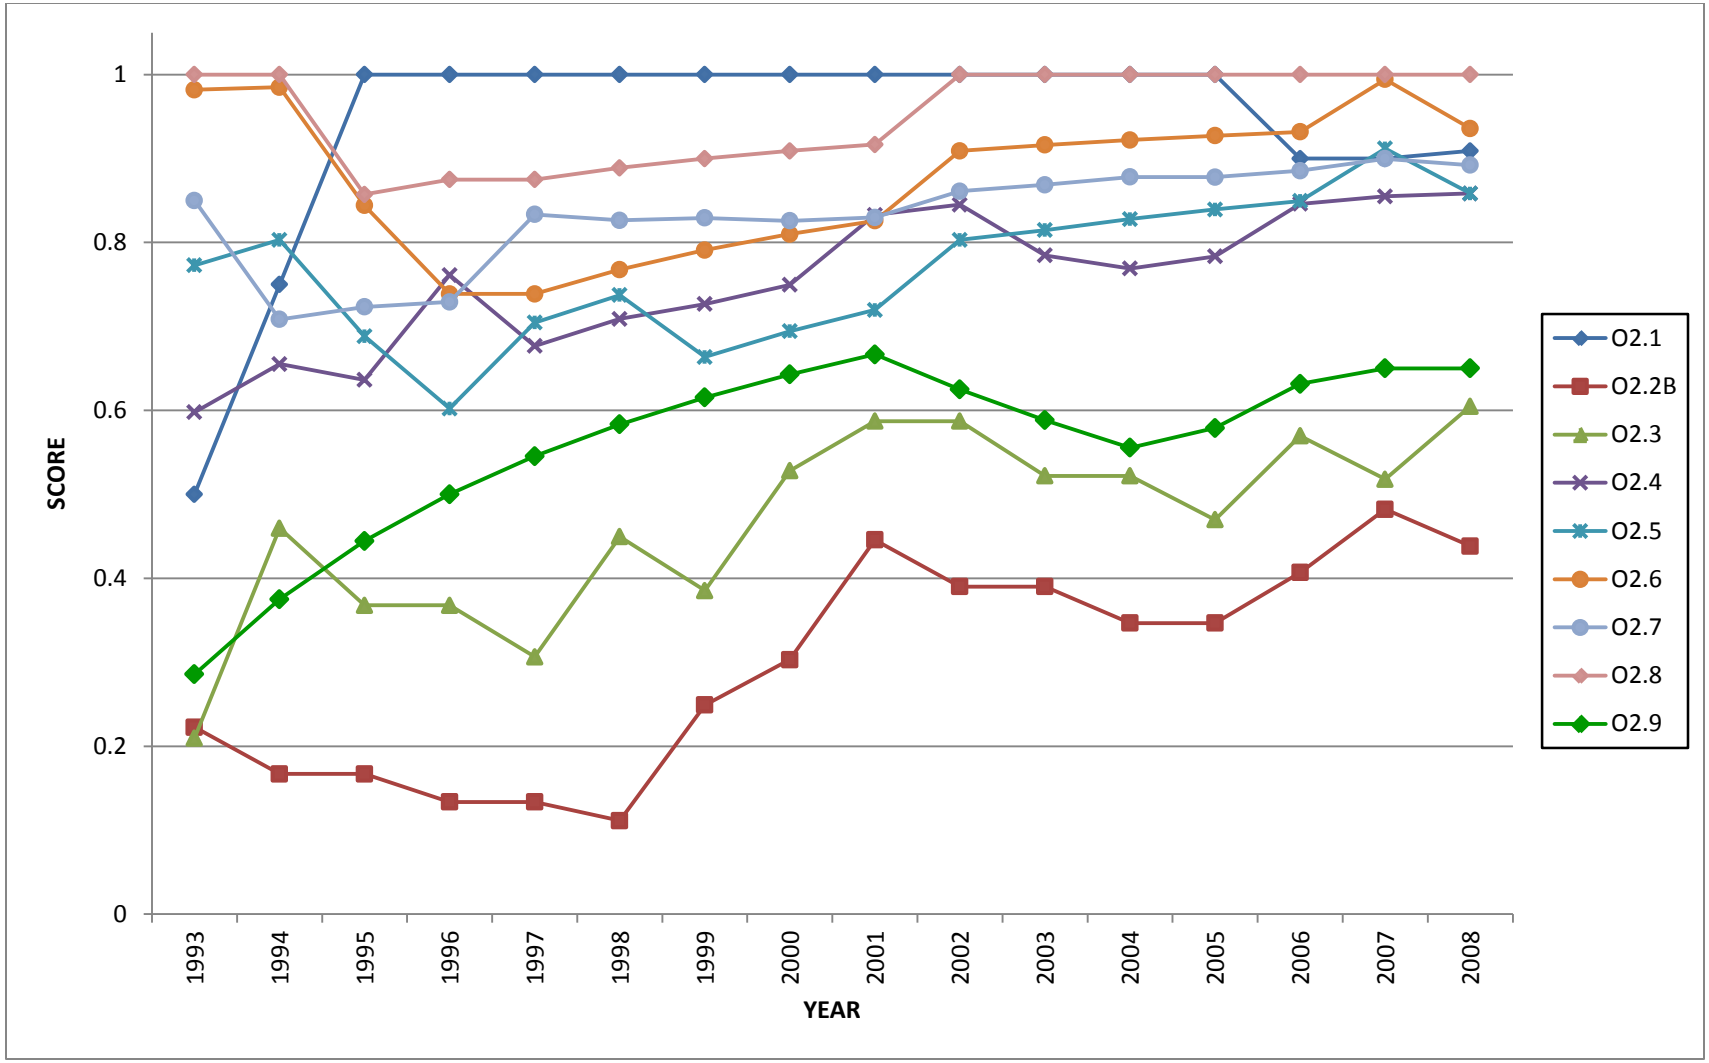

FIGURE 4-47 Objective Performance Score for Objectives 2.1–2.9 for a 16-year Period (1993–2008) for BCGG Site in the Gunnison River Using the IRF Historic Tool (The relevant objectives [labeled O2.1–O2.9 in the graph] are described in detail in Section 4.4.1.3.)

IRF Scores for the Gunnison River Using IRF-Planning Tool Integrated within the WUOT. The planning mode of the IRF tool was used to calculate overall IRF scores at every 6-hour timestep in order to develop the optimized hydrograph for the Gunnison River between April 1, 2009, and September 30, 2009 (Figure 4-48). In this case, the WUOT was operated as if the real time was March 31, 2009; it advanced to the next timestep to develop the most optimized discharge in terms of environmental and economic performance for that time based on inflow and power-price forecasts, while respecting applicable constraints. Once the optimized discharge level was generated by the Seasonal Hydrosystems Analysis tool, the Day-ahead and Real-time Scheduling tool maintained the discharge profile, and as a result the timestep and associated parameter values were pushed into the history table. The simulation then moved to the next future timestep and the process repeated itself while updating the history table at every timestep through the end of the 2009 water-year (i.e., September 30, 2009). In reality, any discrepancy between the long- and short-term inflow forecast would be continuously corrected in the near term as the simulation moved into the future.

The performance of individual sites (i.e., site-specific scores) is also depicted with the overall IRF scores for every time step of the 6-month simulation period (Figure 4-48). Even though the algorithms in the planning mode are somewhat different than those in the historic mode, the definitions used to evaluate the performance of objectives are the same. The aim of environmental performance analysis in planning mode is to demonstrate the functionality of the integrated WUOT by comparing the environmental performance of actual observed operations at the facility to optimized operations, as indicated by the integrated toolset. The score based on observed facility operation at the end of the simulation period (i.e., the assessment date of September 30, 2009), reflects the environmental performance score without the WUOT, and therefore provides a baseline for comparing environmental performance using the integrated toolset.

Daily site-IRF scores for locations 1 and 2 based on actual and optimized operations are shown in Figure 4-48. Similarly, daily scores based on optimized operation for the individual objectives defined for locations 1 and 2 are shown in Figures 4-49 and 4-50, respectively. Figure 4-51 shows flows recorded at the gauge stations for each of the evaluation locations, as well as the simulated hydrograph that would have resulted from the WUOT implementation during the evaluation period.

Comparison between the observed and optimized IRF scores indicates a modest improvement in environmental performance under the implementation of the IRF integrated WUOT. The improvement stems from modification of the releases from the Crystal Reservoir, which distinctly shows delayed peaking for the optimized case in the hydrographs from both locations (Figure 4-51). This further highlights the value of the integrated WUOT, whereby improvements in environmental and economic performance can be made using the same amount of water, but reshaping the hydrograph. Enhancement in the overall IRF score, although modest, resulted mainly from the improvement in the ESWA site score (Figure 4-48). This improved score was due to the improved performance of one particular objective (i.e., Objective 1-12),

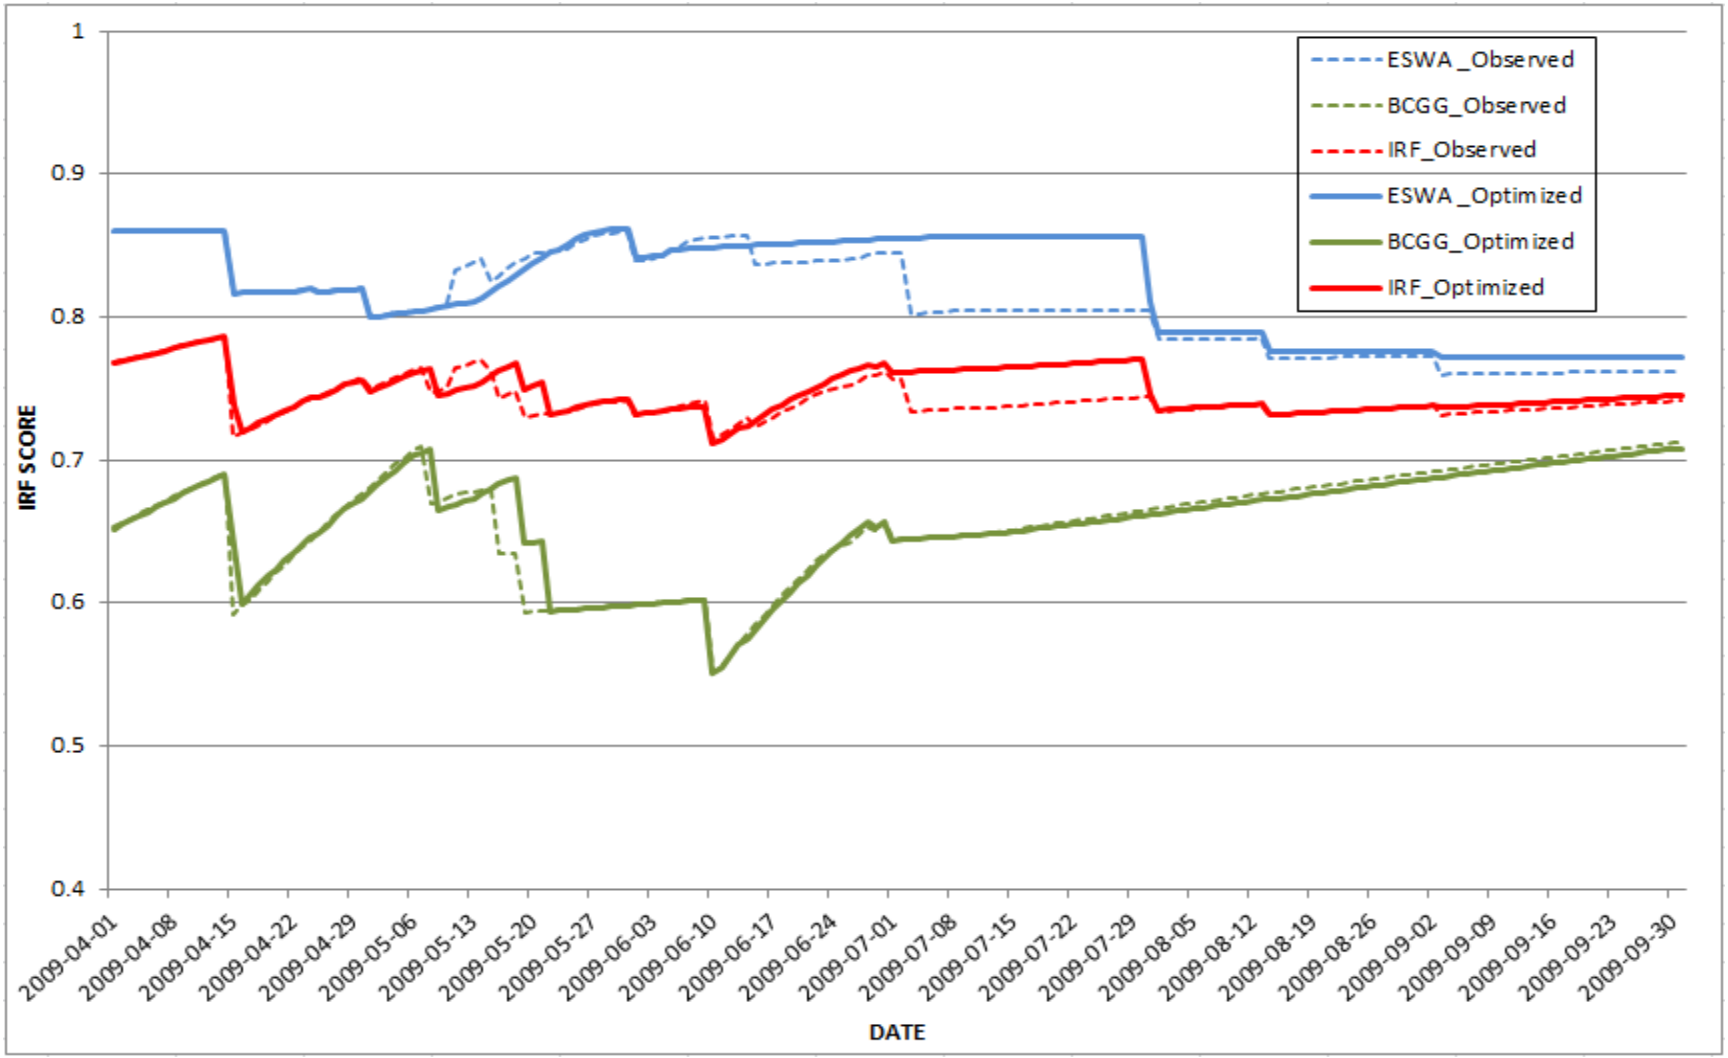

FIGURE 4-48 Daily IRF and Site Environmental Performance Scores Based on Observed Facility Operation (dotted line) and IRF Integrated WUOT Optimized Simulation (solid line) for the Gunnison River Using the IRF Planning Tool for April 1, 2009, to September 30, 2009

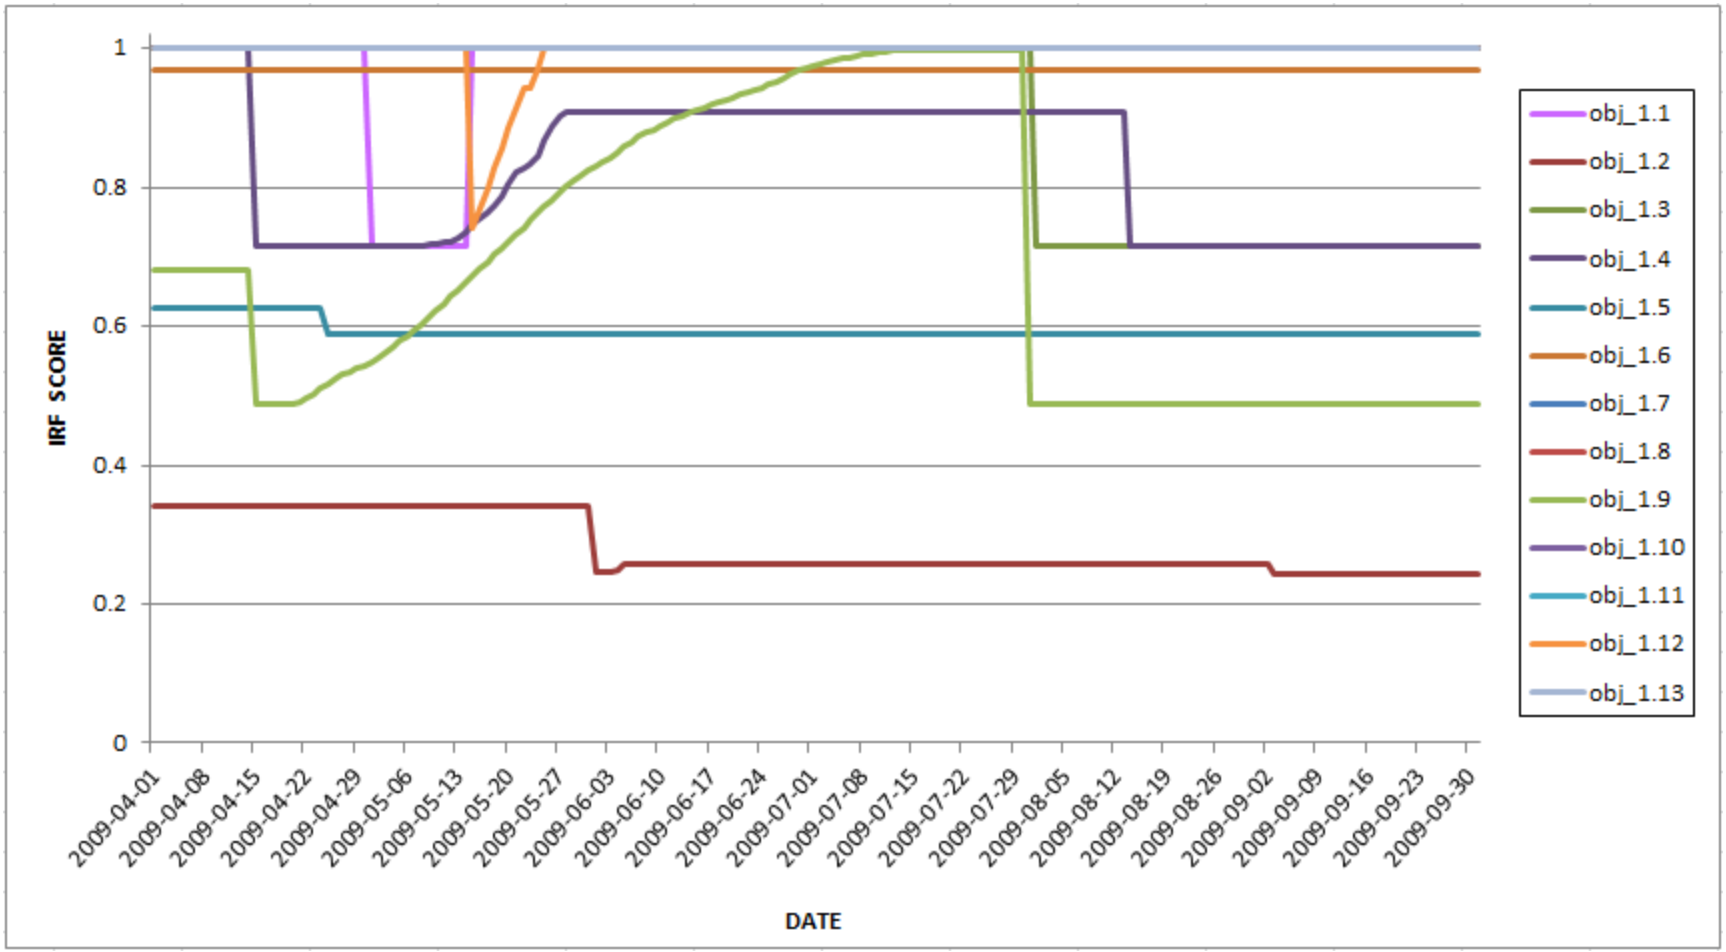

FIGURE 4-49 Daily Objective Performance Score for Objectives 1.1–1.13 based on Optimized WUOT Simulation for the Gunnison River Using the IRF Planning Tool for April 1, 2009, to September 30, 2009, at the ESWA Site (Objectives 1.1–1.13 are described in detail earlier in the result section.)

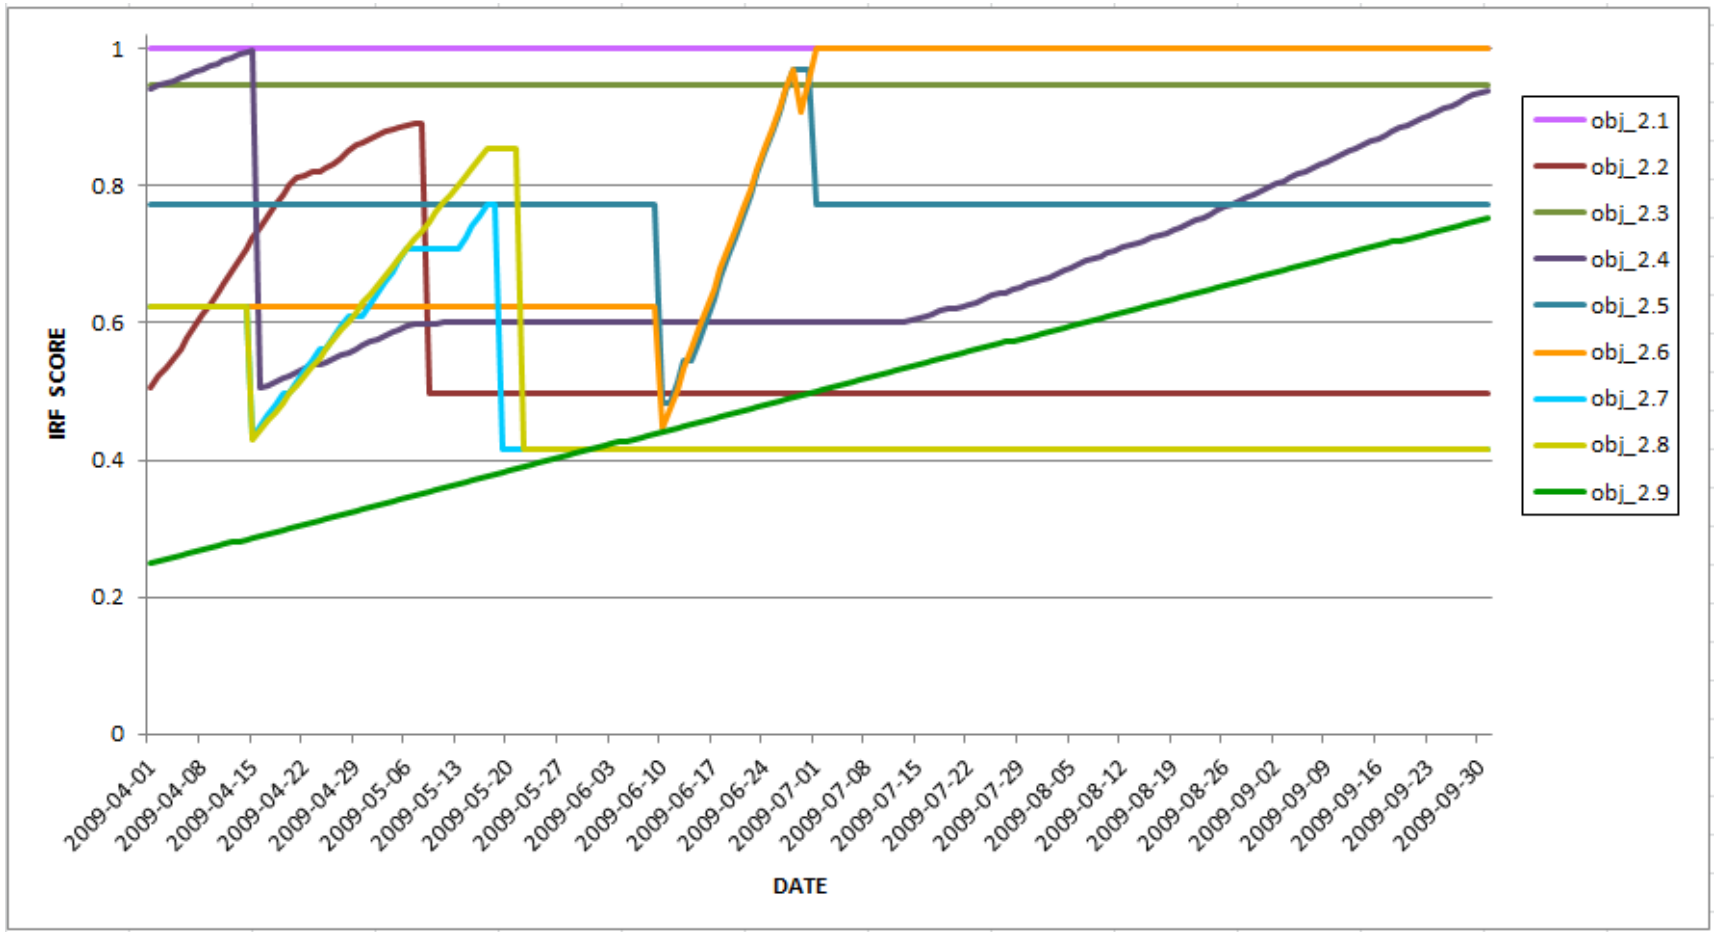

FIGURE 4-50 Daily Objective Performance Score for Objectives 2.1–2.9 based on Optimized WUOT Simulation for the Gunnison River Using the IRF Planning Tool for April 1, 2009, to September 30, 2009, at the BCGG Site (Objectives 2.1–2.9 are described in detail earlier in the result section.)

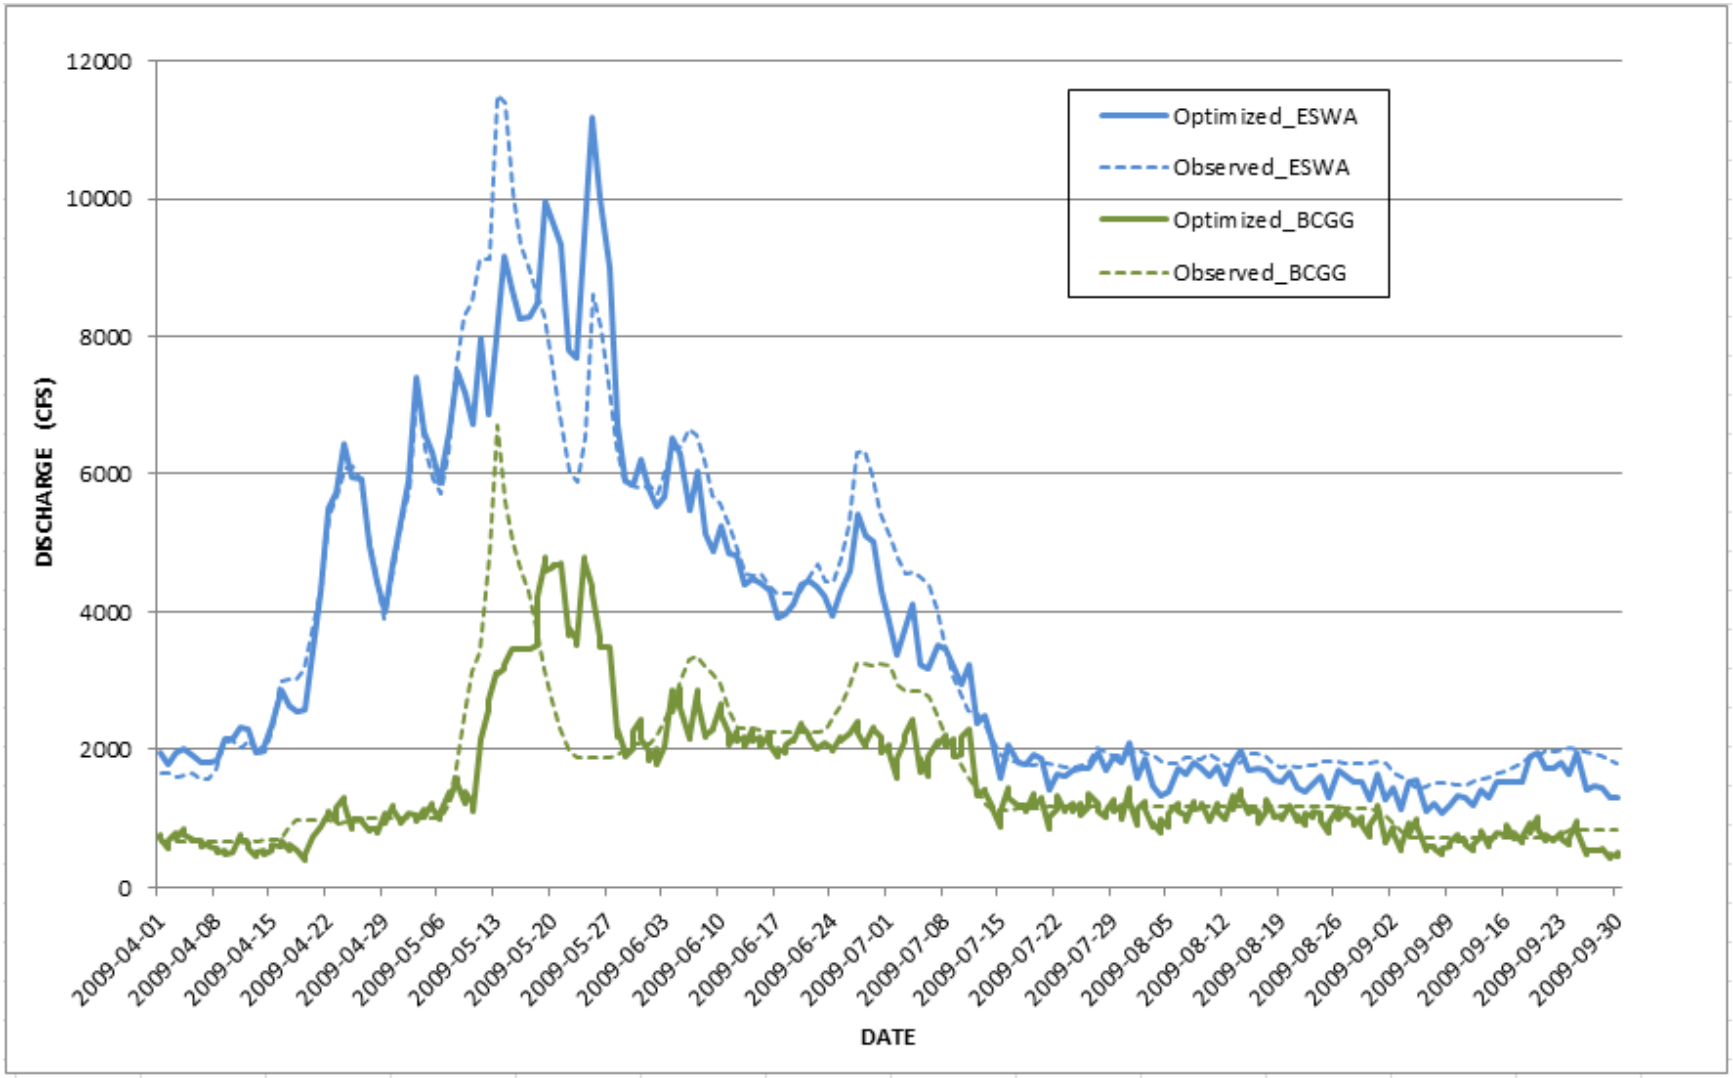

FIGURE 4-51 Actual Observed and WUOT Optimized Hydrographs for the Two Sites in the Gunnison River from April 1, 2009, to September 30, 2009 (Observed Flow data for the ESWA site was recorded from USGS Whitewater gauge, and flow data for the BCGG site was recorded from the USGS site below Gunnison Tunnel.)

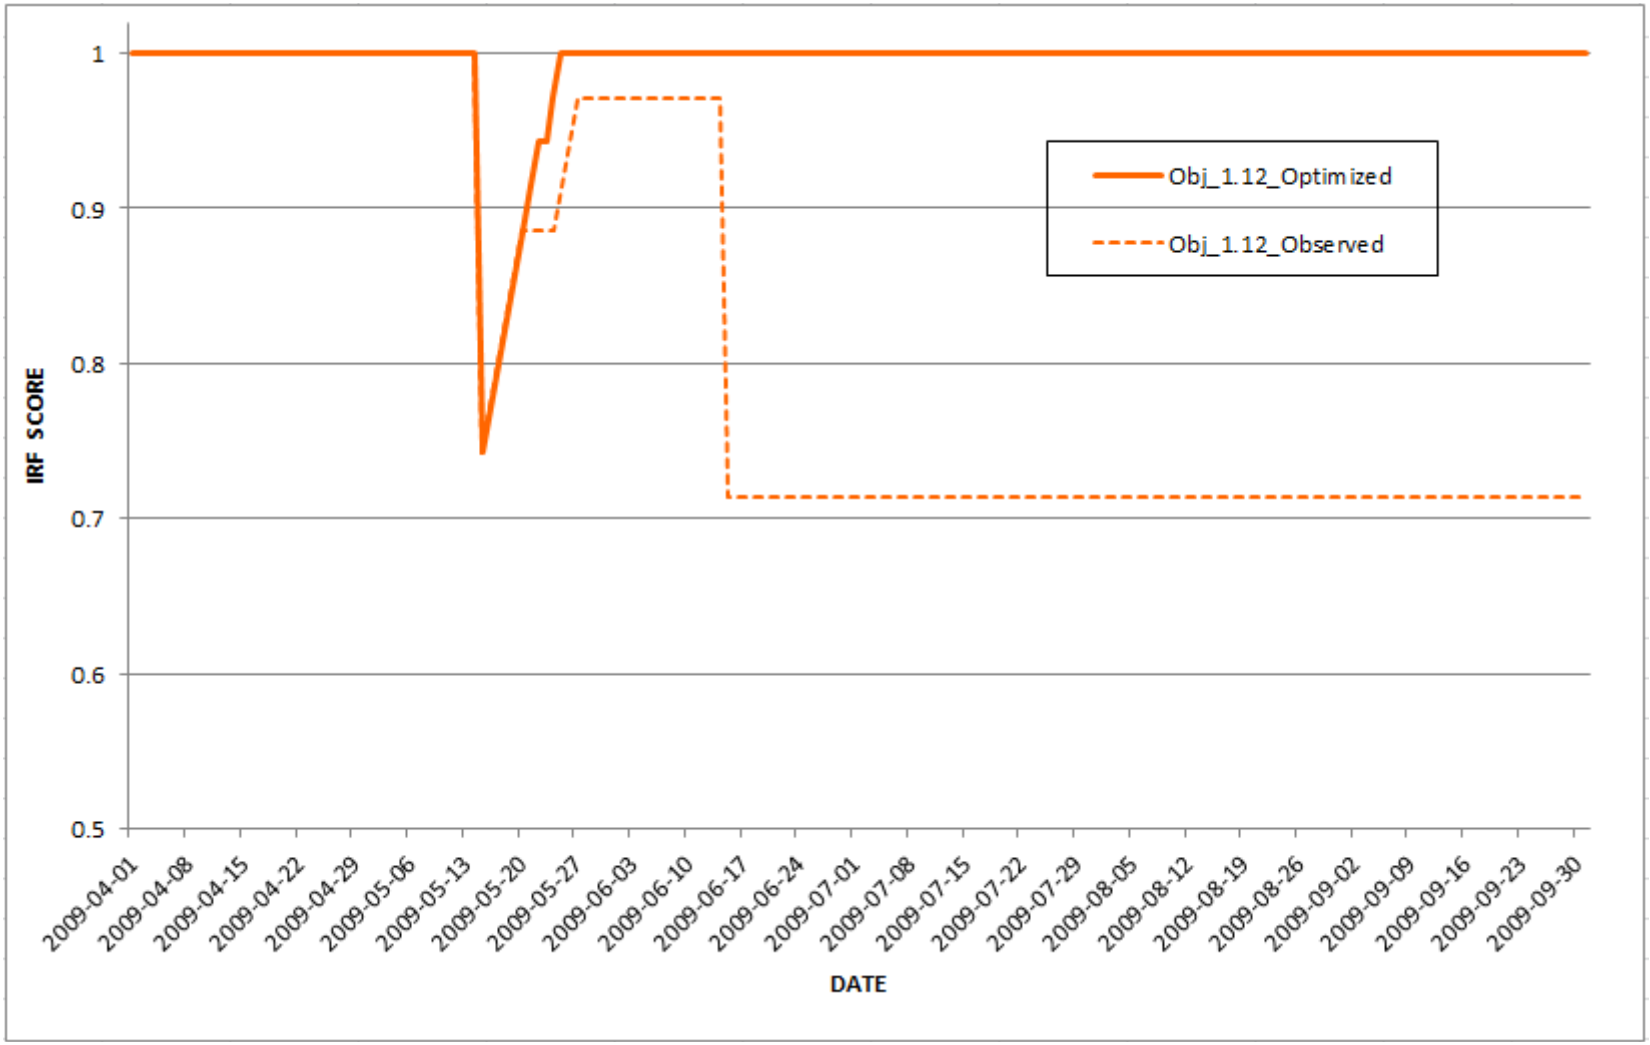

FIGURE 4-52 Comparison of Daily Objective Performance Score for Objective 1.12 between Observed Facility Operation and Optimized WUOT Simulation for the Gunnison River Using the IRF Planning Tool for April 1, 2009, to September 30, 2009, at the ESWA Site (Objective 1.12 pertains to accomplishing half–bank-full discharge, and is described in detail earlier in the result section.)

which is related to achieving flushing-fines flow for an appropriate duration and frequency (Figure 4-52). In fact, scores for some of the objectives were also reduced in the optimized case. However, the significant improvement in the performance of objective 1-12 offset the combined loss from all objectives depicting reduced scores in the optimized case. We would like to point out that we did not assign a higher relative weight to any objective; all objectives had a relative weight of one. The importance of an objective compared to another one varies based on location and timing. We believe that assigning relative importance should be left to the facility operators who are in a better position to understand site-specific concerns and can defend and substantiate the relative importance among various environmental objectives.

Another interesting observation was that unlike in Historical mode, in Planning mode the environmental performance score of the ESWA site was higher than that of the BCGG site (Figures 4-44 and 4-48). This was due to the high spring flows of 2008, when several objectives of the ESWA site were accomplished and therefore were not required to be addressed in water-year 2009. The discrepancy in results between the Historic and Planning modes of IRF should not be a concern, because comparison should always be mode-specific and for the same simulation period. Algorithms in the Planning mode were deliberately designed to help objectives achieve their goals, even before realistic evaluations can be made, by providing fractional points up to the minimum requirement to accomplish objectives that are falling behind on their RF. Historical assessment using the IRF Historic mode, which uses real observations, is a more realistic indicator of environmental performance. It should be noted that any increase in overall, site, or objective scores using Planning mode might not be reflected in the Historic mode of IRF; scores in Historic mode may increase or remain unaffected. However, an increase in environmental performance using Planning mode will not reduce performance using the Historic mode of IRF. Most importantly, with continuous use of the IRF tool, minor improvements in environmental performance in all likelihood will accrue over time.

Even though the modest improvements in environmental performance resulting from WUOT application as indicated by the IRF Planning tool might be insignificant, it is important to note that environmental performance did not deteriorate during the optimization process, and that at the same time a significant economic benefit occurred (~$1 million/yr, in addition to increased unit efficiency). It is also important to understand that in this integrated demonstration exercise, the optimization process initially was somewhat guided by the actual facility operation. While it is true that such guiding baseline conditions will not be offered in a real-life planning scenario, it is also a fact that the IRF incorporated WUOT will not be running in a batch-mode and initial conditions before the start of each day would guide the toolset in moving the solutions to a better space.

In the future, alternative tradeoffs can be explored by placing a comparatively heavier weight on the environmental side than was done in this case. As mentioned earlier, the relative weights of all objectives were designated as 1, specifying that all objectives were equally important. Improvements in the overall IRF score for the facility using WUOT optimization can occur if more important objectives (i.e., those with higher relative weights) show improvement over the observed facility operation. Another reason for lower improvement of environmental performance in the WUOT optimized operation was that, hydrologically, water-year 2009 was a

moderately dry year. Restricted water availability limits the number of ways a given volume can be shaped while meeting regulatory and facility constraints. In this regard, the environmental results would have been significantly better for water-year 2010, which was a hydrologically wet year.

Caution should be used in interpreting IRF scores or relating them to actual environmental performance. As mentioned earlier, these scores indicate relative performance among objectives (including sites or basins) based on known habitat requirements for target objectives. The IRF tool considers parameters of flow (such as discharge, water temperature, stage) that could be controlled in a regulated setting to be the only factor in evaluating environmental performance. Some of the most cited and highly regarded studies related to environmental management of regulated rivers (Richter et al. 1996; Poff et al. 1997) regard stream flow as the "master variable." Even though this rationale appears prudent, ecological systems are often too complex to be explained by a single variable, and natural manifestation at any point in time is an interaction of multiple factors accumulated over time. For example, even though IRF scores indicated favorable conditions for both rainbow trout and brown trout based on flow requirement at the BCGG study site, an epidemic of whirling disease caused significant declines in the rainbow trout population in the mid-to-late 1990s; this did not, however, affect the brown trout. The newly created habitat created by the decline in rainbow trout was well exploited by the brown trout, so much so that their more aggressive piscivorous behavior has been voiced as a genuine concern for the survival of juvenile native fish species and for the ultimate recruitment and extending of their current range. Predicting biotic interactions such as the sudden outbreak of disease and interspecies interactions is beyond the scope of IRF tool. As a result, periodic field surveying is strongly recommended and users of this tool should modify objective definitions to better suit targets based on the best available information, rather than relying exclusively on IRF as an environmental management tool. The IRF methodology thus allows adaptive management capability to meet environmental targets in regulated rivers.

4.4.1.4 Stand-Alone CHEERS Retrospective Analysis of CRSP Day-Ahead and Real-Time Scheduling

The usefulness of the CHEERS model for CRSP applications was demonstrated by a retrospective analysis over a 6-month time span from April 1, 2009, through September 30, 2009. In support of the analysis, CHEERS creates day-ahead schedules and optimizes CRSP operations. For each modeled day, the retrospective simulation process performs two runs. The first CHEERS model run optimizes day-head plans and the second model run optimizes real-time operations. In both cases, physical hydrological and power processes are simulated.

All model runs have the same basic objective subject to a fixed set of operating criteria and management rules. In general, these mimic typical business practices that are currently utilized by staff in the CRSP EMMO. The analysis assumes that Western day-ahead and real-time schedulers would have engaged in market transactions prescribed by CHEERS at prices input into the model. The analysis also assumes that Reclamation hydropower plant operators would have deployed each hydropower unit according to real-time model results. Model results

are compared to observed unit operations and market transactions that actually occurred during the 6-month retrospective study period.

The sections below provide more details about how CHEERS is used for the retrospective analysis when run in stand-alone mode (i.e., without interacting with the other WUOT tools). Later in this document, a CHEERS analysis performed in an integrated mode (i.e., with other WUOT tools) is discussed.

Day-Ahead Optimization. During the day-ahead business process, CRSP EMMO schedulers lay the groundwork for the next day's real-time operations. This entails preparing hourly day-ahead schedules for plant-level operations and entering into short-term firm (STF) bilateral energy purchases and sales contracts for deployments the following day. Day-ahead schedulers also prepare preliminary energy balance purchase and sale amounts for real-time schedulers.

As previously mentioned, day-ahead model runs include two types of market transactions. These include STF contractual agreements and hourly balancing transactions. Firm bilateral agreements are typically made for 16-hour and 8-hour blocks. The 16-hour block occurs during peak load hours. Although the timing of the 16-hour block changes over the course of the year, during the retrospective period it begin at 6 AM and ends at 10 PM, inclusive. The remaining time defines the 8-hour off-peak block period. The modeling process assumes that within each block the purchase or sales level is constant, but the level may differ by block. In contrast, balancing market transaction amounts may vary hourly and are made in real time. For day-ahead planning purposes, these purchases are projections of hourly transactions that may occur tomorrow in order to balance CRSP system supply and demand.

A day-ahead model run of CHEERS is based on projections about events that will transpire from the time day-ahead schedules are prepared through midnight of the next calendar day. These projections do not necessarily come to fruition in the CHEERS modeling process, reflecting schedulers' decision making under uncertainty. For day-ahead CHEERS runs, it is assumed that schedulers will run the model shortly after 10 AM, and therefore "know" operations and system state histories through 10 AM on the morning before the day-ahead optimization time period. CHEERS utilizes the history of actual recorded observations up to this point in time with certainty. From 11 AM of the current day until midnight of the current day, simulated CHEERS schedulers make projections based on initial conditions at 10 AM and information previously recorded in the CRSP scheduling workbook. Identical to actual day-ahead schedulers, CHEERS also knows these schedules with certainty. A more detailed description of this scheduling workbook is provided in the section that follows. Day-ahead projections are to some degree inaccurate.

The CHEERS day-ahead modeling process that is used for the retrospective analysis is shown in the upper portion of Figure 4-53. The process mimics key uncertainties and reflects the current business practices used by Western schedulers. This includes uncertainties associated with both hourly water inflows/side flows into the Aspinall Cascade and Reclamation's mandates

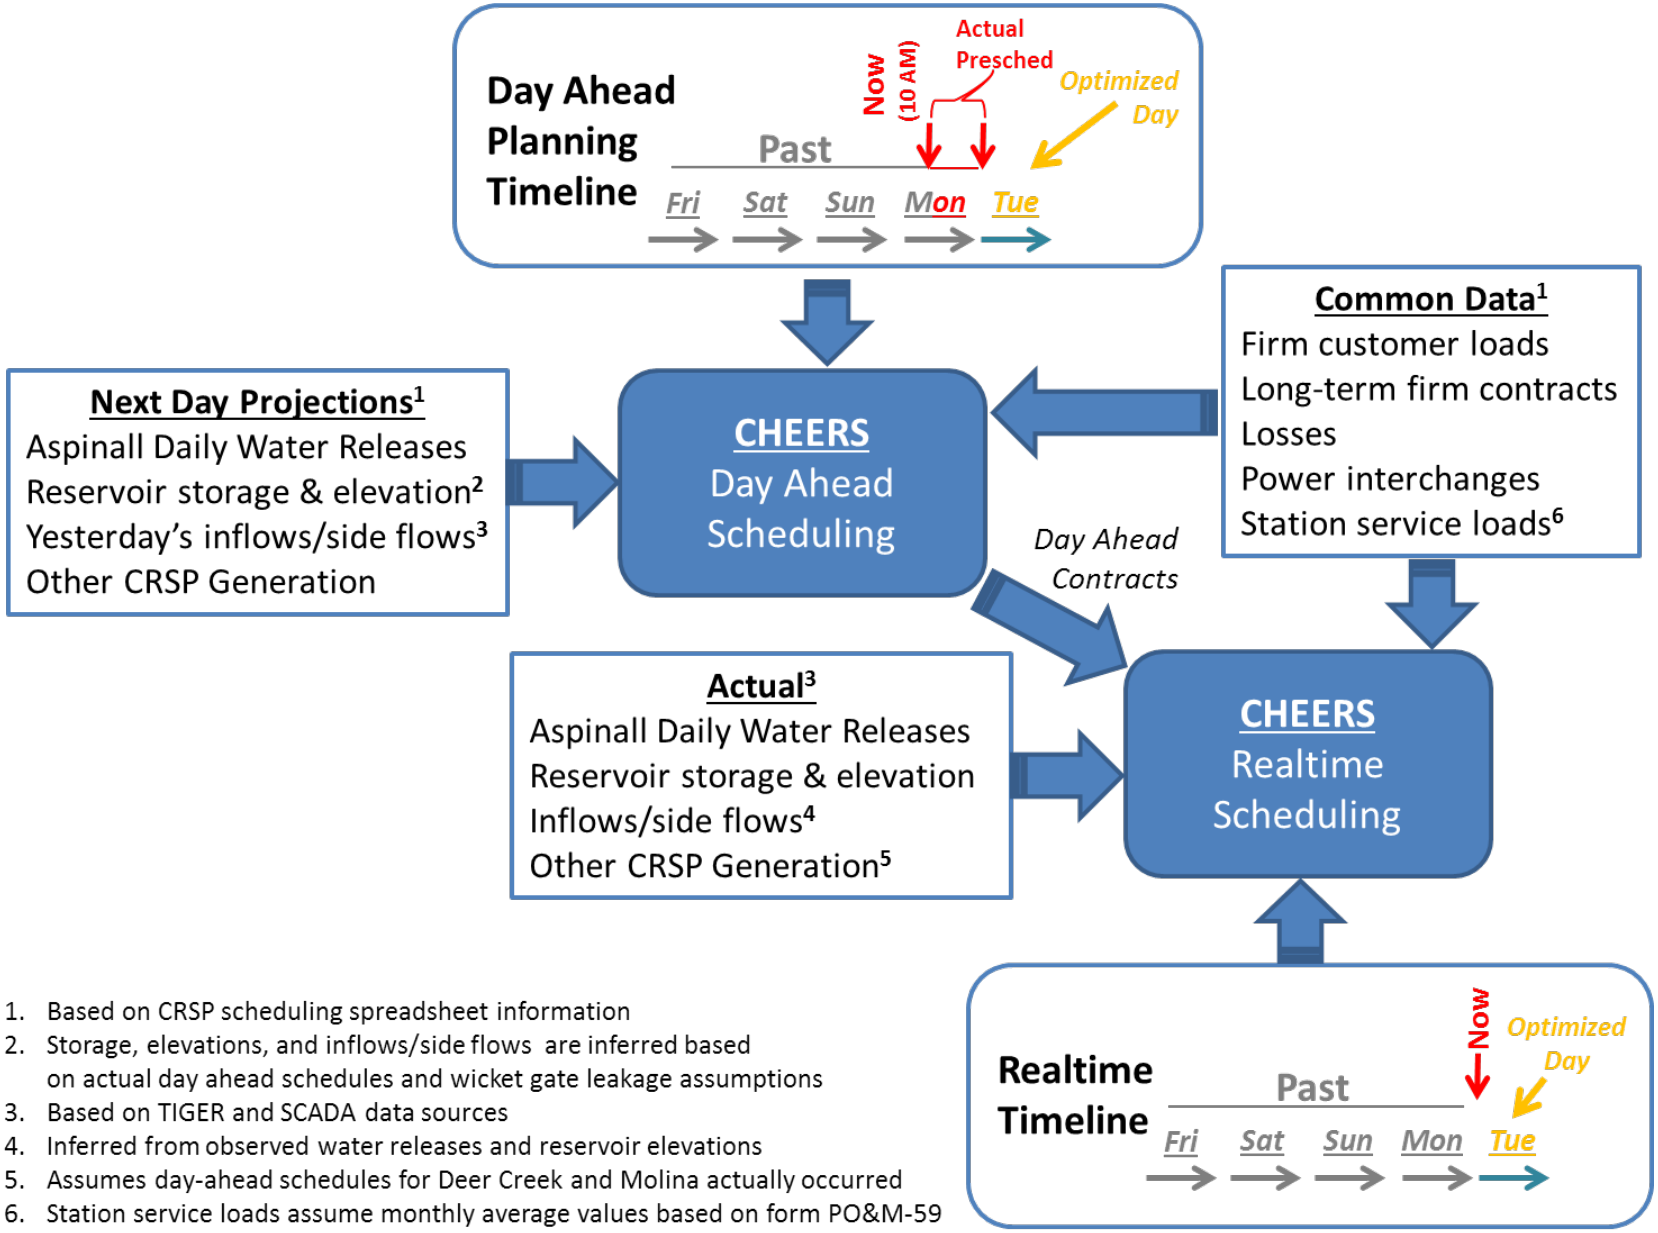

1. Based on CRSP scheduling spreadsheet information
2. Storage, elevations, and inflows/side flows are inferred based on actual day ahead schedules and wicket gate leakage assumptions
3. Based on TIGER and SCADA data sources
4. Inferred from observed water releases and reservoir elevations
5. Assumes day-ahead schedules for Deer Creek and Molina actually occurred
6. Station service loads assume monthly average values based on form PO&M-59

FIGURE 4-53 CHEERS Model Run Sequence and Data Flows for Stand-Alone Retrospective Analyses

for next-day reservoir water release volumes. The retrospective methodology used for stand-alone CHEERS runs involves several steps and assumptions, as outlined below.

1. Project next-day cascade inflows/side flows based on persistence forecasting. Model runs assume that yesterday's average daily observed inflows/side flow rate will occur the rest of the current day and all day tomorrow.

2. As recorded in the actual scheduling workbook, assume that hourly plant-level day-ahead schedules that were prepared yesterday will occur through the rest of the current day

3. Reservoir elevations from the current day-ahead decision hour (between 10 and 11 AM) until midnight of the current day are based on side flow assumptions from step 1 and reservoir water releases associated with yesterday's day-ahead schedules. These water releases are based on scheduled generation levels divided by historical daily water-to-power conversion factors.

4. For day-ahead runs, assume that the next day's total water release volume equals the total *day-ahead schedule generation* levels times the historical daily water-to-power conversion factor.

5. Optimize CRSP schedules for the 24-hour period between midnight of the current day and midnight tomorrow using (a) historical day-ahead schedules for all CRSP power plants except for those in the Aspinall Cascade, (b) projected inflows/side flows from step 1, (c) daily reservoir-level water releases from step 4, (d) existing long-term firm contacts, (e) historical firm customer loads, and (f) all other hourly schedules such as losses and exchanges as recorded in CRSP workbook.

In short, the CHEERS day-ahead run does not use any information that would not have been known to an actual scheduler using the tool at 10 AM on the day in question.

STF bilateral agreements made in the day-ahead optimization run are assumed to be take-or-pay contracts and therefore always occur in real time without alteration. In contract, balancing market transactions produced by this first CHEERS run are only suggestions for real-time schedulers and often change in response to events that differ from day-ahead predictions.

Real-Time Optimization. As simulated time transpires, CHEERS is run a second time. Figure 4-53 shows that real-time runs begin on midnight of the optimized day. At this point, a scheduler using CHEERS knows what actually occurred during the 14-hour period from the time day-ahead schedules were formulated through midnight. For this second run, it is also assumed that total daily water releases and actual hourly water inflows/side flows are known with certainty. Power produced by other CRSP hydropower plants is assumed to be scheduled at actual operating levels rather than at prescheduled levels. In addition, all STF contracts optimized by the day-ahead CHEERS model run are treated as fixed energy transactions. Optimized real-time model results for Aspinall Cascade operations and balancing market

transactions differ from day-ahead outputs when projections made during the day-ahead scheduling process do not come to fruition.

CRSP Daily Scheduling Workbook. The CHEERS model mimics the operational business practices applied by day-ahead and real-time schedulers. Many of these practices were briefly discussed in the previous section. One key component of the scheduling cycle involves the CRSP Excel workbook. EMMO schedulers use this tool to balance CRSP loads and resources by entering hourly values into customized worksheets.

In this section, model results are compared to actual CRSP operations for a few select days that have various hydrological and operational characteristics. Hydropower and reservoir operations, STF market transactions, and power economics are presented.

Figure 4-54 shows an actual day-ahead load and resource balance for Monday April 20, 2009. The graph shows hydropower generation produced by CRSP resources and power supplied through Western's purchasing programs including long-term firm (LTF) purchases connected to its Available Hydropower (AHP) and Western Replacement Power (WRP) marketing plans. Additional STF and LTF purchases are also made on an as-needed or economic basis as market and hydrological conditions evolve. Other sources of energy that feed into the CRSP system include power interchanges sent to Western from neighboring systems and energy payback when Western, on a previous day, generated more power than the amount needed to cover its share of system transmission losses. Note that on this particular day hydropower resources were inadequate to cover all of the loads. Under this energy "short" position, Western needed to make relatively large energy purchases. Some of this shortfall was covered by a LTF WRP purchase that was made long before the day-ahead decision timeframe.

Following general marketing guidelines in the CRSP EMMO, the vast majority of the remaining energy shortfall during this day was covered via 8-hour off-peak and 16-hour on-peak block STF purchases. Much smaller spot market purchases referred to as "Balancing Purchases" were made to fill in the gaps in the remaining shortfall; these are the represented by yellow bars in hours 4–7 and 23–24 in Figure 4-54.

Figure 4-55 is a screen capture of a portion of a worksheet from the CRSP workbook used by schedulers to balance the system. The workbook contains a separate worksheet for each day of the month. Line items in each daily worksheet include schedules for power plant level generation, customer load scheduled under STF and LTF contracts, losses, power interchanges, and energy imbalances. For each line item, schedulers enter 24 hourly values. Values are first entered by day-ahead schedulers, but can be modified in real time as grid and hydrological events unfold. However, long-term bilateral agreements are fixed.

In the CRSP system, actual operations are typically very similar to day-ahead generation patterns. The left half of Figure 4-56 shows that for April 20, 2009, actual unit generation and day-ahead schedules for Morrow Point Power Plant generation are very closely aligned. Note that the blue line that represents the day-ahead scheduled generation in each hour and the stacked unit-level generation are almost identical except for hour 18, and that day-ahead schedules

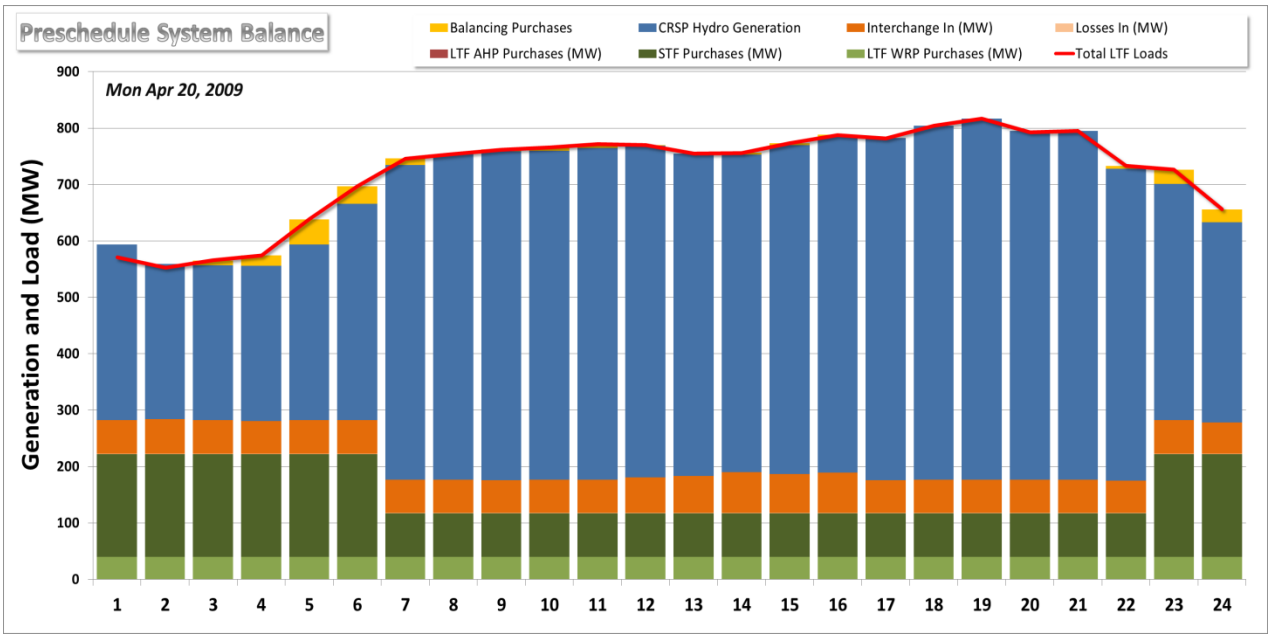

FIGURE 4-54 CRSP Supply and Demand Energy Balance

devised by CRSP EMMO staff are specified at the plant-level while actual operations carried out by Reclamation operators are at the unit level.

Sometimes it is necessary for either CRSP schedulers or Reclamation plant operators to make adjustments in real-time operations. The right half of Figure 4-56 shows that real-time operations were significantly different from day-ahead schedules during several afternoon hours.

As will be discussed in more detail later in this section, the CHEERS Results Report Writer tool summarizes both actual historical events and modeled results. The graphs in Figures 4-54 and 4-56 can be produced with drop-down menus by this post-processer routine for any day for which data are in the WUOT common database.

The CHEERS topology shown in Figure 4-57 was designed by Argonne staff in consultation with Western schedulers. The design is compatible with the structure of the CRSP Excel workbook and the information that it contains. CHEERS scheduling input data will be linked to this workbook via a data bridge that is currently being constructed. This joint Western and Argonne initiative is scheduled to be completed this summer. As required, CHEERS will also utilize information and data contained in Western's PI historical database and TIGER scheduling system.

Many of the nodes in the CHEERS network, which represents the CRSP system, correspond to line items in a daily scheduling sheet. These include CRSP hydropower resources located at dams at Glen Canyon (GC), Flaming Gorge (FG), Fontenelle (FN), the Molina (ML) Project, and Deer Creek (DC). Based on discussions with CRSP EMMO staff, plant-level generation levels for these facilities are held fixed in the CHEERS model runs. For day-ahead

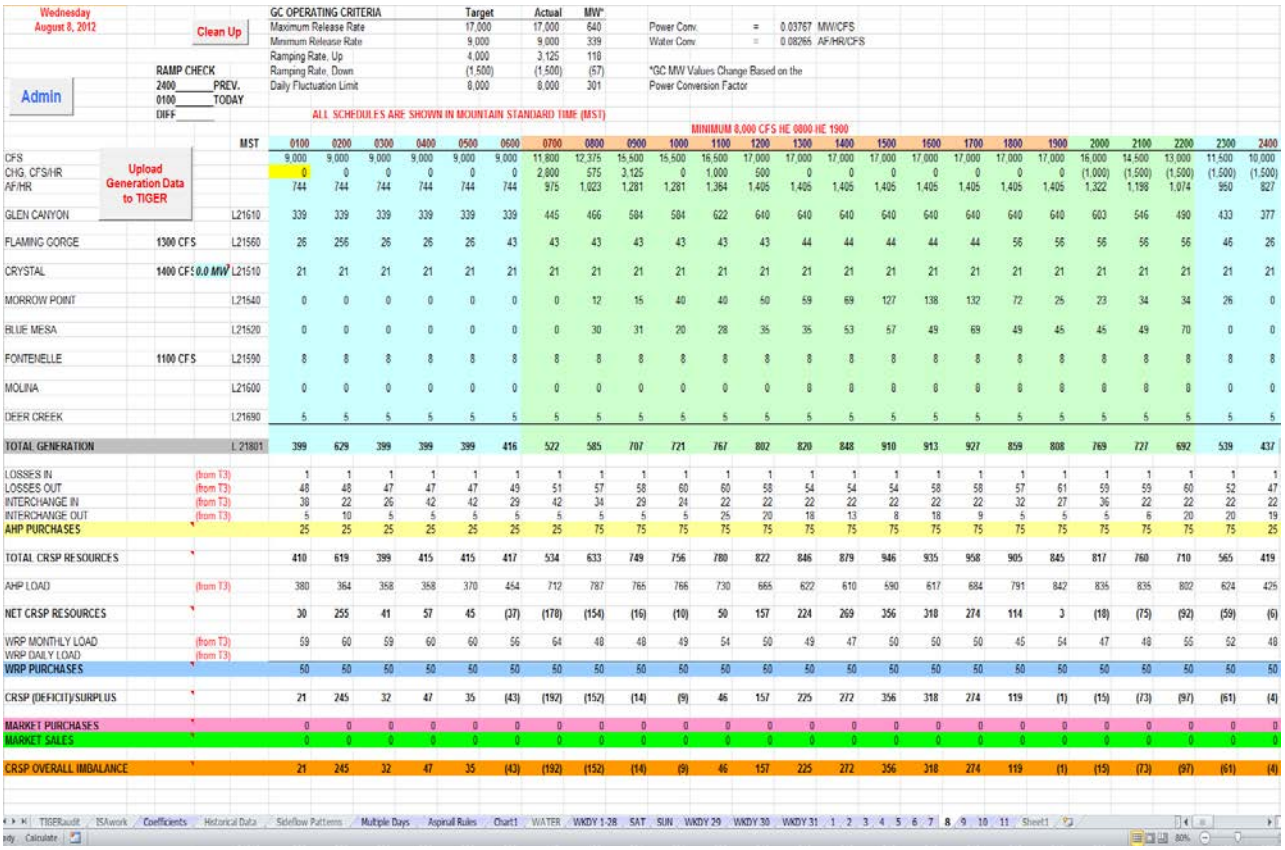

Wednesday August 8, 2012

Clean Up

Admin

RAMP CHECK
2400 PREV.
0100 TODAY
DIFF

GC OPERATING CRITERIA	Target	Actual	MW*
Maximum Release Rate	17,000	17,000	640
Minimum Release Rate	9,000	9,000	339
Ramping Rate, Up	4,000	3,125	118
Ramping Rate, Down	(1,500)	(1,500)	(57)
Daily Fluctuation Limit	8,000	8,000	301

Power Conv. = 0.03767 MW/CFS
Water Conv. = 0.08265 AF/HR/CFS

*GC MW Values Change Based on the Power Conversion Factor

ALL SCHEDULES ARE SHOWN IN MOUNTAIN STANDARD TIME (MST)

MINIMUM 8,000 CFS HE 0800 HE 1900

Upload Generation Data to TIGER

MST		0100	0200	0300	0400	0500	0600	0700	0800	0900	1000	1100	1200	1300	1400	1500	1600	1700	1800	1900	2000	2100	2200	2300	2400
CFS		9,000	9,000	9,000	9,000	9,000	9,000	11,800	12,375	15,500	15,500	16,500	17,000	17,000	17,000	17,000	17,000	17,000	17,000	17,000	16,000	14,500	13,000	11,500	10,000
CHG. CFS/HR		0	0	0	0	0	0	2,800	575	3,125	0	1,000	500	0	0	0	0	0	0	0	(1,000)	(1,500)	(1,500)	(1,500)	(1,500)
AF/HR		744	744	744	744	744	744	975	1,023	1,281	1,281	1,364	1,405	1,405	1,405	1,405	1,405	1,405	1,405	1,405	1,322	1,198	1,074	950	827
GLEN CANYON	L21610	339	339	339	339	339	339	445	466	584	584	622	640	640	640	640	640	640	640	640	603	546	490	433	377
FLAMING GORGE	1300 CFS	26	256	26	26	26	43	43	43	43	43	43	43	44	44	44	44	44	56	56	56	56	56	46	26
CRYSTAL	1400 CFS 0.0 MW	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21
MORROW POINT	L21540	0	0	0	0	0	0	0	12	15	40	40	50	59	69	127	138	132	72	25	23	34	34	26	0
BLUE MESA	L21520	0	0	0	0	0	0	0	30	31	20	28	35	35	53	57	49	69	49	45	45	49	70	0	0
FONTENELLE	1100 CFS L21590	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8
MOLINA	L21600	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
DEER CREEK	L21690	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5
TOTAL GENERATION	L 21801	399	629	399	399	399	416	522	585	707	721	767	802	820	848	910	913	927	859	808	769	727	692	539	437
LOSSES IN	(from T3)	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
LOSSES OUT	(from T3)	48	48	47	47	47	49	51	57	58	60	60	58	54	54	54	58	58	57	61	59	59	60	52	47
INTERCHANGE IN	(from T3)	38	22	26	42	42	29	42	34	29	24	22	22	22	22	22	22	22	32	27	36	22	22	22	22
INTERCHANGE OUT	(from T3)	5	10	5	5	5	5	5	5	5	5	25	20	18	13	8	18	9	5	5	5	6	20	20	19
AHP PURCHASES		25	25	25	25	25	25	25	75	75	75	75	75	75	75	75	75	75	75	75	75	75	75	75	25
TOTAL CRSP RESOURCES		410	619	399	415	415	417	534	633	749	756	780	822	846	879	946	935	958	905	845	817	760	710	565	419
AHP LOAD	(from T3)	380	364	358	358	370	454	712	787	766	766	730	665	622	610	590	617	684	791	842	835	835	802	624	425
NET CRSP RESOURCES		30	255	41	57	45	(37)	(178)	(154)	(16)	(10)	50	157	224	269	356	318	274	114	3	(18)	(75)	(92)	(59)	(6)
WRP MONTHLY LOAD	(from T3)	59	60	59	60	60	56	64	48	48	49	54	50	49	47	50	50	50	45	54	47	48	55	52	48
WRP DAILY LOAD	(from T3)																								
WRP PURCHASES		50	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50
CRSP (DEFICIT)/SURPLUS		21	245	32	47	35	(43)	(192)	(152)	(14)	(9)	46	157	225	272	356	318	274	119	(1)	(15)	(73)	(97)	(61)	(4)
MARKET PURCHASES		0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
MARKET SALES		0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
CRSP OVERALL IMBALANCE		21	245	32	47	35	(43)	(192)	(152)	(14)	(9)	46	157	225	272	356	318	274	119	(1)	(15)	(73)	(97)	(61)	(4)

FIGURE 4-55 Screen Capture of the CRSP Daily Scheduling Sheet

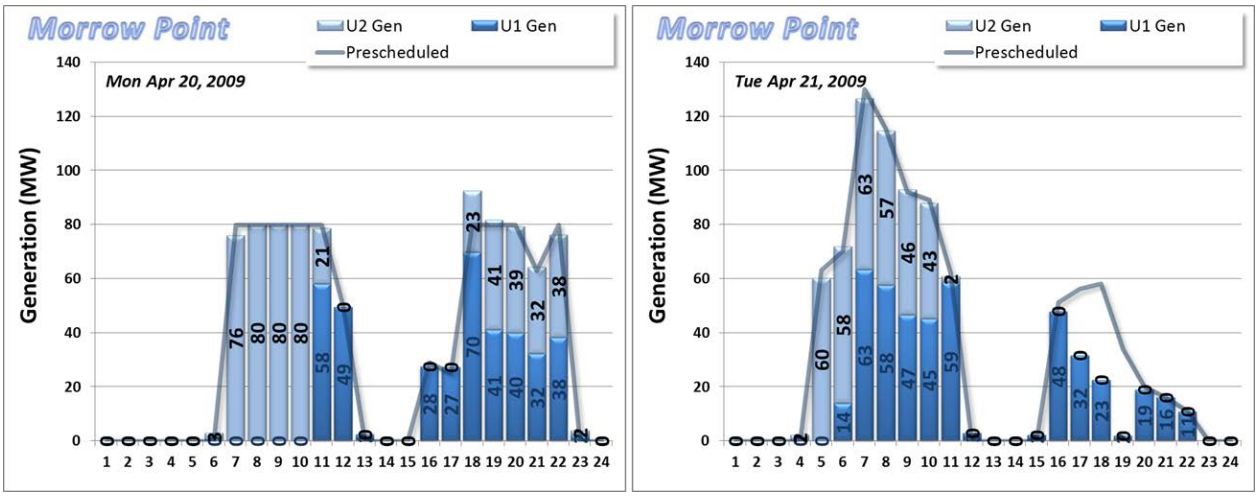

FIGURE 4-56 Morrow Point Power Plant Day-Ahead and Actual Unit Generation Level for April 20, 2009, and April 21, 2009

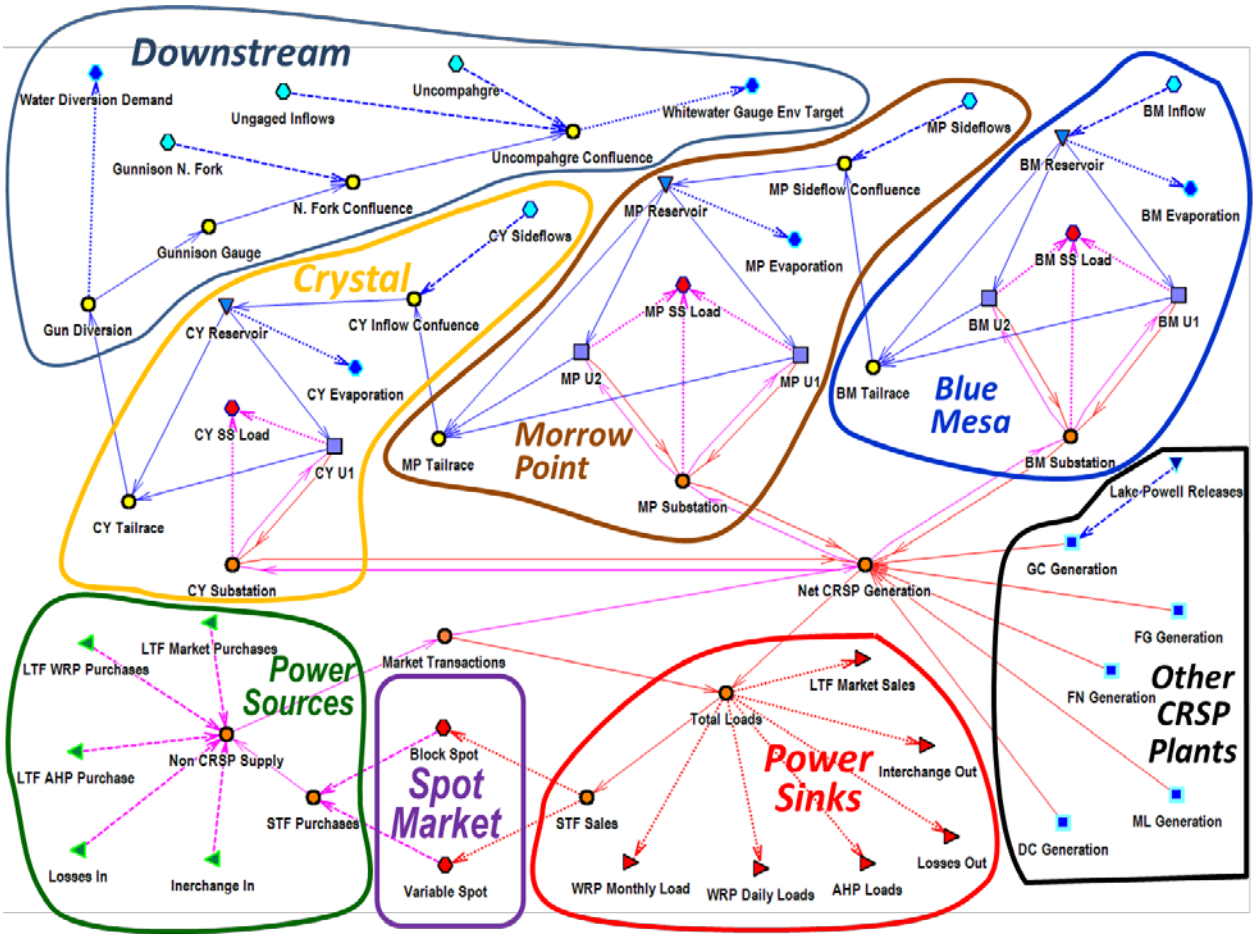

FIGURE 4-57 CHEERS Network Developed for the CRSP Demonstration Site

scheduling, modeled generation levels exactly match day-ahead levels as documented in historical CRSP workbooks. For CHEERS real-time runs, however, plant-level operations at GC, FG, and FN are based on actual operations as posted on the CRSP Web site (http://www.wapa.gov/crsp/opsmaintcrsp/scada.htm). Real-time operations for the ML and DC, both of which are relatively small, were not readily available.

Generating capabilities for plants in the Aspinall Cascade depart somewhat from the CRSP scheduling workbook in that the network contains a more refined level of granularity; that is, the network has a unit level representation of the cascaded hydrological system while the workbook is at the plant level. The requisite for a more detailed representation of specific CRSP resources is discussed in more detail later in this section.

LTF bilateral purchase agreements are typically binding for a month or more and are typically “take-or-pay” contracts that have been struck a month or more in advance. Therefore, in CHEERS, purchases are held fixed in accordance with contracted levels for both day-ahead and real-time runs.

Other sources of energy that feed into the CRSP system include power interchanges with neighboring systems (labeled "Interchange In") and energy payback when Western, on a previous day, generated more power than the amount needed to cover transmission losses (labeled "Losses In"). CRSP hydropower and non-hydropower energy supply resources are used to serve customer loads. This includes loads that Western is obligated to serve under AHP LTF contracts, along with loads under WRP monthly and daily LTF contracts. Other energy sinks include transmission losses (labeled "Losses Out") and energy interchange payback (labeled "Interchange Out"). STF and LTF purchases may also be made to balance the system when CRSP resources exceed demands. All of these CRSP workbook line items are also held fixed by CHEERS in both day-ahead and real-time runs since these are known in advance with certainty.

Day-ahead STF purchases are represented by the "Block Spot" node. STF contract durations are typically one or two days and are, at most, agreed upon between buyers and sellers only a few days in advance. As previously described, these STF transactions are optimized in the CHEERS day-ahead run given projections about what will occur in the near future. During the real-time run, day-ahead STF arrangements are held fixed even though it may be advantageous to alter the contract when inflow and daily water release projections do not come to fruition.

Western may also make spot market purchases and sales to balance the system. In CRSP scheduling sheets, these are computed and displayed in the "imbalance" line item. This corresponds to the node labeled "Variable Spot" in the CHEERS network and is used to represent the energy balancing market. Computations of balancing transactions are made in both day-ahead and real-time runs. However, unlike LTF and STF contracts, these transactions are unrestricted in real-time CHEERS optimizations runs.

Except for water releases out of the Crystal Dam, other portions of the network labeled "Downstream" in Figure 4-57 are not contained in the CRSP spreadsheet. It represents downstream water delivery obligations and environmental water flow requirements. When CHEERS is run in an integrated mode, it utilizes guidance from other WUOT tools to ensure that Aspinall operations are compliant each day with these requirements from both short-term and long-term perspectives.

Model Constraints and Aspinall Cascade Representation. Based on consultation with Western staff, the CHEERS model is constrained to yield specific results for several CRSP network components. As discussed above, some CHEERS nodal results are fixed such that output values exactly match either scheduler entries in the CRSP scheduling workbook or actual operations. All LTF contacts, including loads and both purchase and sales agreements, are inflexible in terms of day-ahead and real-time scheduling decisions. These are all known with certainty a day or more in advance and Western is obligated to abide by all contractual terms.

In addition, other than Aspinall Cascade hydropower resources, there is very little or no operational flexibility at other CRSP facilities. These are tightly constrained by stringent environmental regulations. Therefore, CRSP schedulers requested that CHEERS only allow the model to deviate from scheduler entries for operations at the Aspinall power plants along with day-ahead and real-time market transactions. This includes those purchases and sales that are

made on the STF and balancing spot markets. Transactions made on the spot market include those that are scheduled in terms of constant hourly quantities during either on-peak or off-peak hours.

During the CRSP demonstration process, schedulers expressed the need for an improved methodology for estimating the effects of hourly power schedules on reservoir elevations at the Blue Mesa and Morrow Point Reservoirs in the Aspinall Cascade. As described in previous sections, both reservoirs in the cascade are subject to minimum and maximum elevation constraints. In addition, the Crystal Reservoir has restrictions on the change in reservoir elevations over daily and multiday time spans. To help Western staff improve the scheduling of generation from Aspinall power plants, Argonne staff added line items and associated graphics to the CRSP daily scheduling sheet. A screen capture of the graphics portion of these additions is provided in Figure 4-58.

To more closely correspond to the CRSP scheduling sheets, Argonne staff and CRSP schedulers built more details into CHEERS to represent the Aspinall Cascade compared to other CRSP hydropower resources. Detailed model components shown in the CRSP network shown in Figure 4-57 include the following:

1. Storage nodes (blue inverted triangles) with reservoir storage-elevation functions and hourly forebay elevation computations;

2. Functional boundary nodes (light blue hexagons) that represent hourly water inflows into the Blue Mesa reservoir located at the top of the cascade and hourly side flows between the two lower reservoirs;

3. Junction nodes (yellow circles) that aggregate river channel main stem flows and side flows between reservoirs at Blue Mesa and Morrow Point and between Morrow Point and Crystal Reservoirs;

4. Functional boundary nodes (dark blue hexagons) that represent hourly reservoir evaporation;

5. Conversion nodes (light blue squares) that represent power plant generating units with associated water-to-power conversion functions and hourly generation computations;

6. Junction nodes (yellow circles) that represent the tailrace below a dam with associated tail water elevation functions and hourly elevation computations;

7. Functional boundary nodes (red hexagons) that represent hourly station service load; and

8. Junction nodes (orange circles) that represent the distribution of electricity into and out of the generating units to satisfy CRSP loads/energy sinks and station service load.

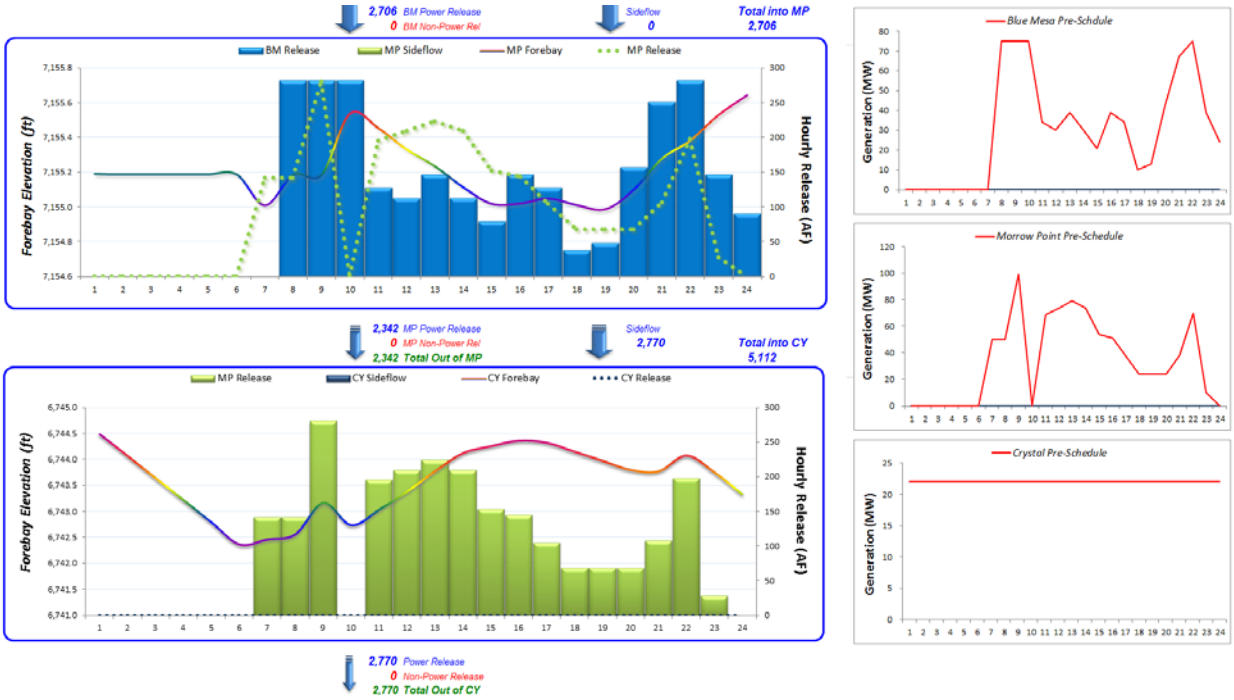

FIGURE 4-58 Aspinall Reservoir Graphics Added to CRSP Daily Scheduling Sheets

The CHEERS model representations of reservoirs and side flows are very similar to the ones imbedded in the CRSP workbook. However, the representation of Aspinall power resources in CHEERS is more detailed compared to the CRSP workbook. This is primarily a reaction to analyses of historical unit-level operations that revealed units have significantly different efficiency levels during the course of a day. Results of these analyses, discussed in more detail in the following sections, show that efficiency is primarily a function of turbine water flow rate.

CRSP Data Sources and CHEERS Model Inputs. Supporting data that were used to prepare inputs for the CHEERS CRSP network were supplied by both Western and Reclamation staff. This section discusses this data, data processing, and model inputs. Data processing involves the use of both linear and nonlinear approximation techniques that are supported by and performed with the following CHEERS Excel workbook tools: (1) Curve Fitting Algorithm, (2) Interpolation Routine, and (3) Piecewise Linear Optimizer.

Storage and Elevation Functions at Aspinall Reservoirs. As explained and discussed in the previous sections, there are a variety of parameter pairs that describe nonlinear functions in both water and power systems. The discussion below presents polynomial approximations that were used to represent these relationships. To find the polynomial fits, the Curve Fitting Algorithm utilizes Excel's "linest" function. Curve Fitting Algorithm results are presented for reservoir elevations approximations for Aspinall Cascade Reservoirs. Raw input data for Blue Mesa, Morrow Point, and Crystal reservoirs were supplied by Reclamation staff. Figures 4-59,

4-60, and 4-61 show the third-order polynomial approximations for Blue Mesa, Morrow Point, and Crystal, respectively. Corresponding polynomial coefficients are given in Table 4-9.

Reservoir water storage-elevation functions are typically concave. At a low reservoir elevation a small change in water storage results in a relatively large elevation change and when the reservoir is full a small decrease in reservoir water storage results in a relatively small reservoir elevation change. For the three Aspinall reservoirs, there is a nearly constant elevation change for each AF change in the reservoir storage over daily operating ranges. Therefore, for this application at each reservoir a single line segment is used compute elevations. The slope of the line is based on two points on the reservoir elevation/volume function. Points represent the approximate reservoir operational storage range that is expected to occur during the CHEERS optimized day.

Water Release Rate and Tail-Water Elevation at Aspinall Dams. For this Aspinall study, it is assumed that the tail water (or tail race) elevation is a nonlinear function of total water releases from a reservoir. This includes the sum of all power and non-power releases. In support of previous CRSP studies conducted at Argonne, Reclamation staff provided tail water elevation data that is consistent with these assumptions.

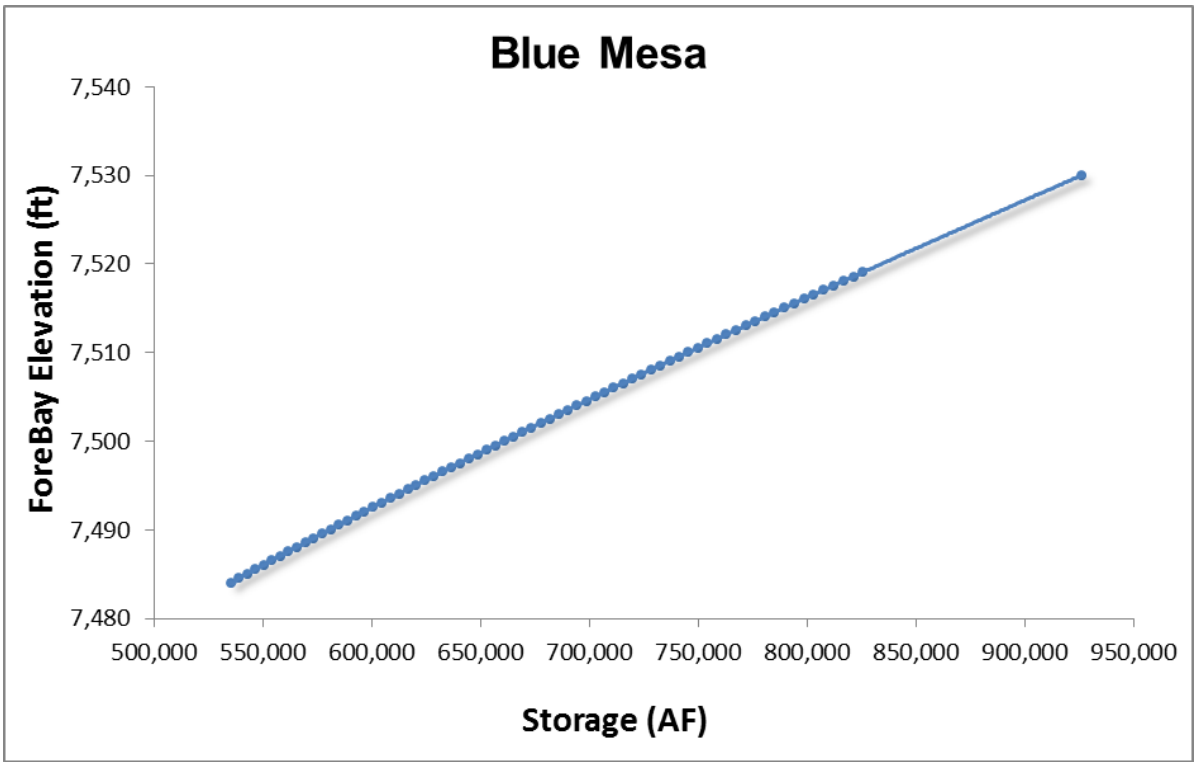

FIGURE 4-59 Blue Mesa Reservoir Elevation as a Function of Water Storage Volume

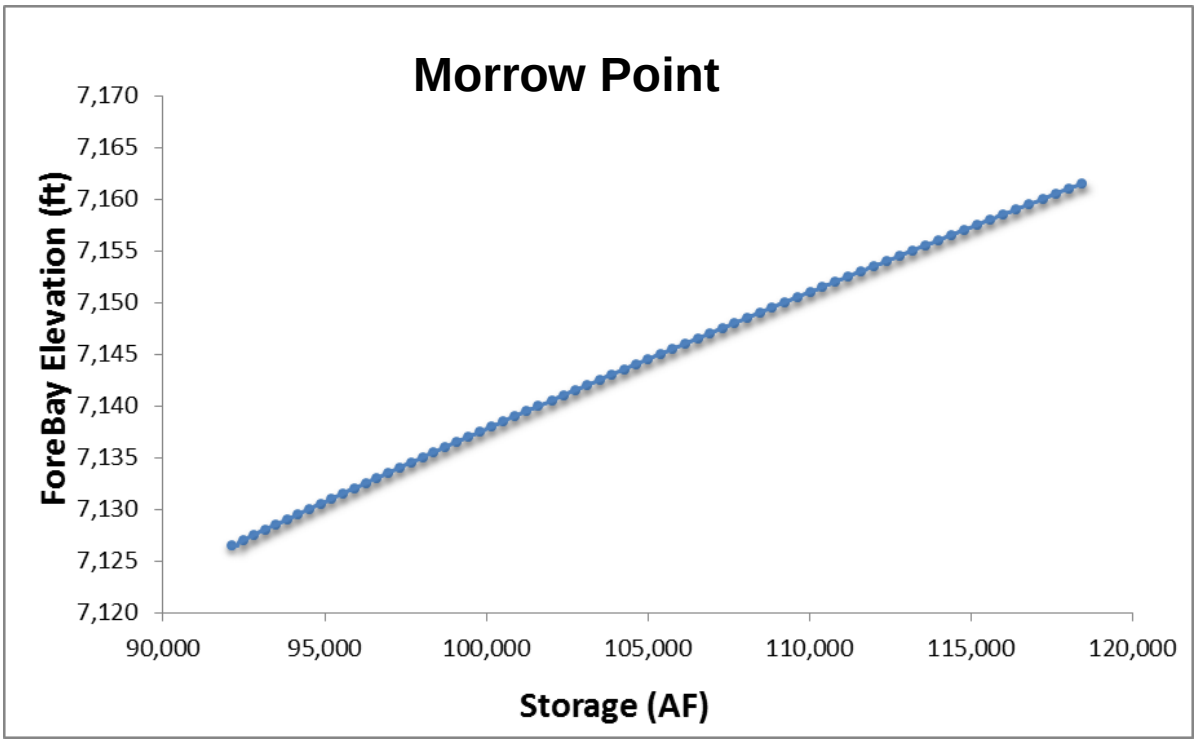

FIGURE 4-60 Morrow Point Reservoir Elevation as a Function of Water Storage Volume

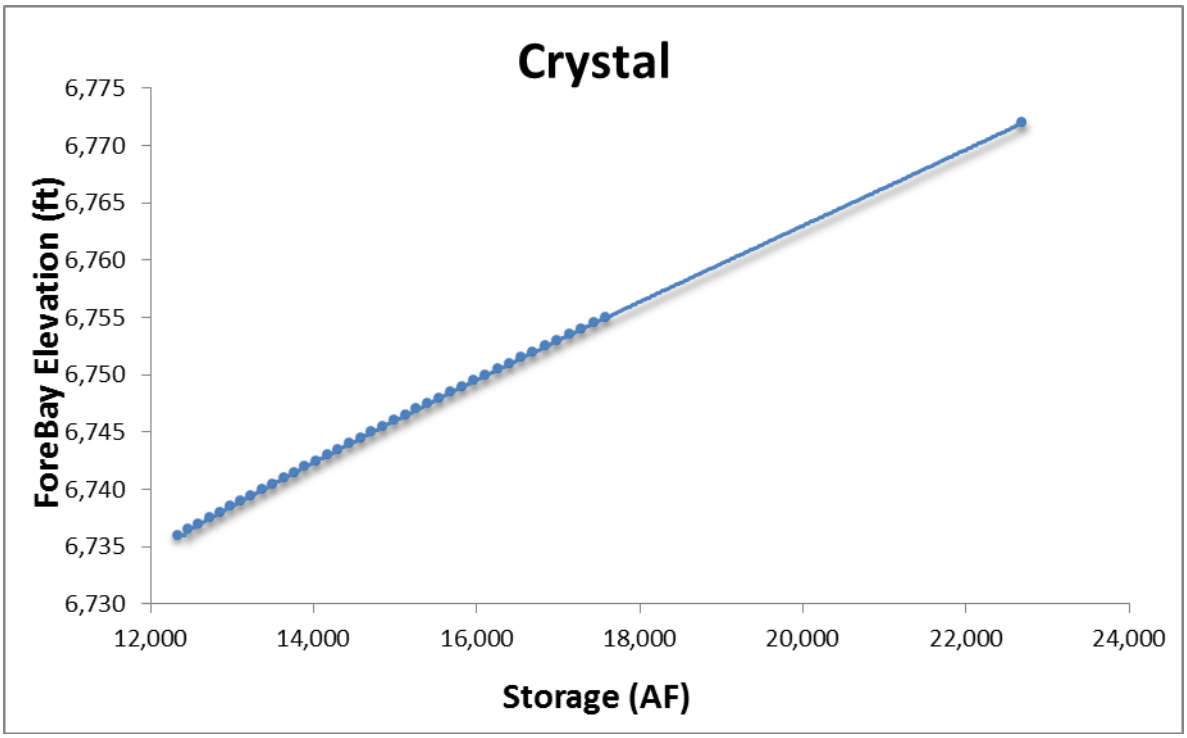

FIGURE 4-61 Crystal Reservoir Elevation as a Function of Water Storage Volume

TABLE 4-9 Polynomial Coefficients for the Reservoir Elevation Functions

Reservoir	Intercept	1st Order	2nd Order	3rd Order
Blue Mesa	7.39E+03	2.24E–04	–1.13E–10	3.64E–17
Morrow Point	6.86E+03	5.04E–03	–3.02E–08	7.97E–14
Crystal	6.67E+03	7.98E–03	–2.33E–07	3.88E–12

However, Reclamation data are not supported by actual tail water observations. Therefore, SCADA data from April 1, 2009, through September 30, 2009, were used to construct tail water elevation curves. These data were trimmed by removing average hourly power releases less than 500 cfs. These observations are generally for either the last or first 10 minutes of an hour as power plant generation is ramping up to or down to a zero output level. Polynomial and linear approximations for Blue Mesa and Morrow Point are presented in Figures 4-62 and 4-63. For Crystal, the tail water observations in Figure 4-64 only contain values with elevations below 6,535 ft. Others were excluded since above this level observations are highly erratic and typically represent spill conditions. Tail water function polynomial coefficients are given Table 4-10.

Reservoir elevation levels minus the tail water elevation are used by CHEERS to compute hourly hydraulic head values. The head value is then utilized in another function that restricts the maximum turbine flow rate as a function of head.

Power Output as a Function of Turbine Water Release Rate at Aspinall Generating Units. Power output at Aspinall generating units is a nonlinear function of water release rate and head. This is evident in the predicted characteristic curves for Aspinall generating units. Although these predictions, made in the early 1970s, provide some useful insights, the curves provided in Appendix G do not adequately reflect actual current operations. Therefore, for this analysis historical information was used to construct turbine flow rate versus power output functions.

Hourly data for unit-level generation is available in Western's data archives. Unfortunately, dam water releases are only reported for aggregate power channels and for each non-power release route including spillway and bypass water. Since both Blue Mesa and Morrow Point have two generating units each, unit level power functions were based on only those observations when one unit or the other operated alone. A separate analysis was also performed on those observations when both units generated power at the same time.

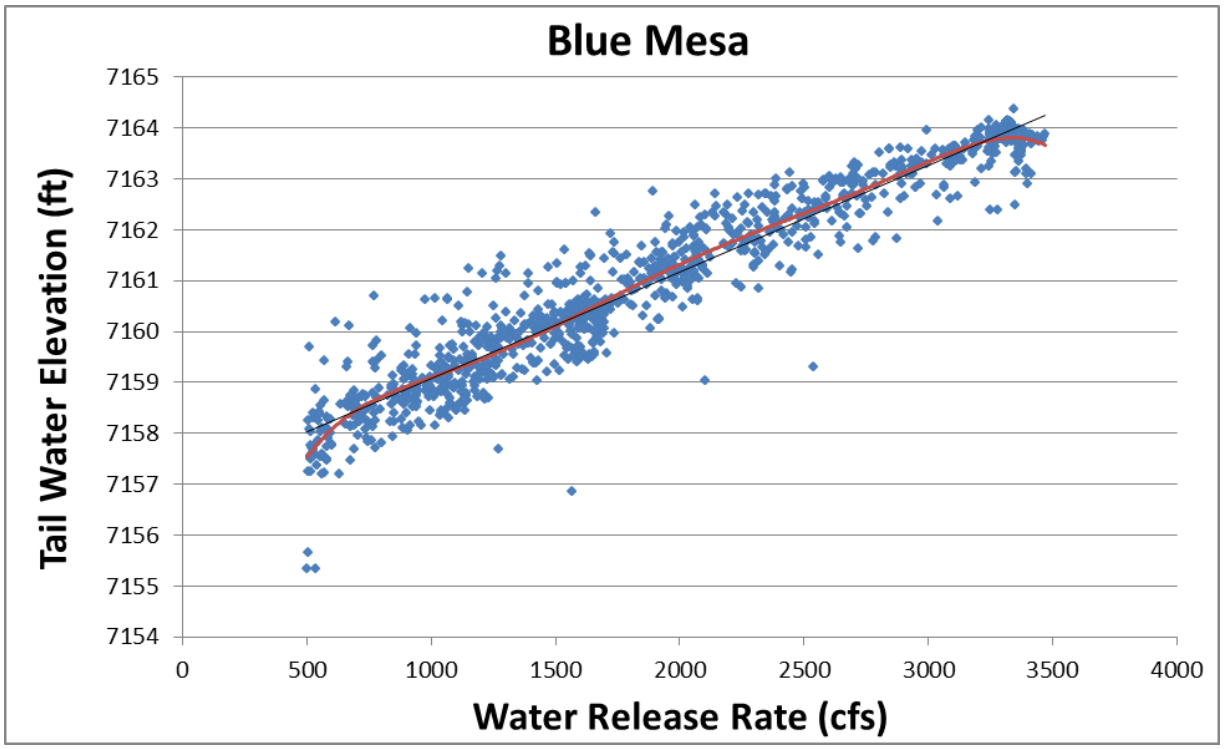

FIGURE 4-62 Blue Mesa Tail Water Elevation as a Function of Water Release Rate

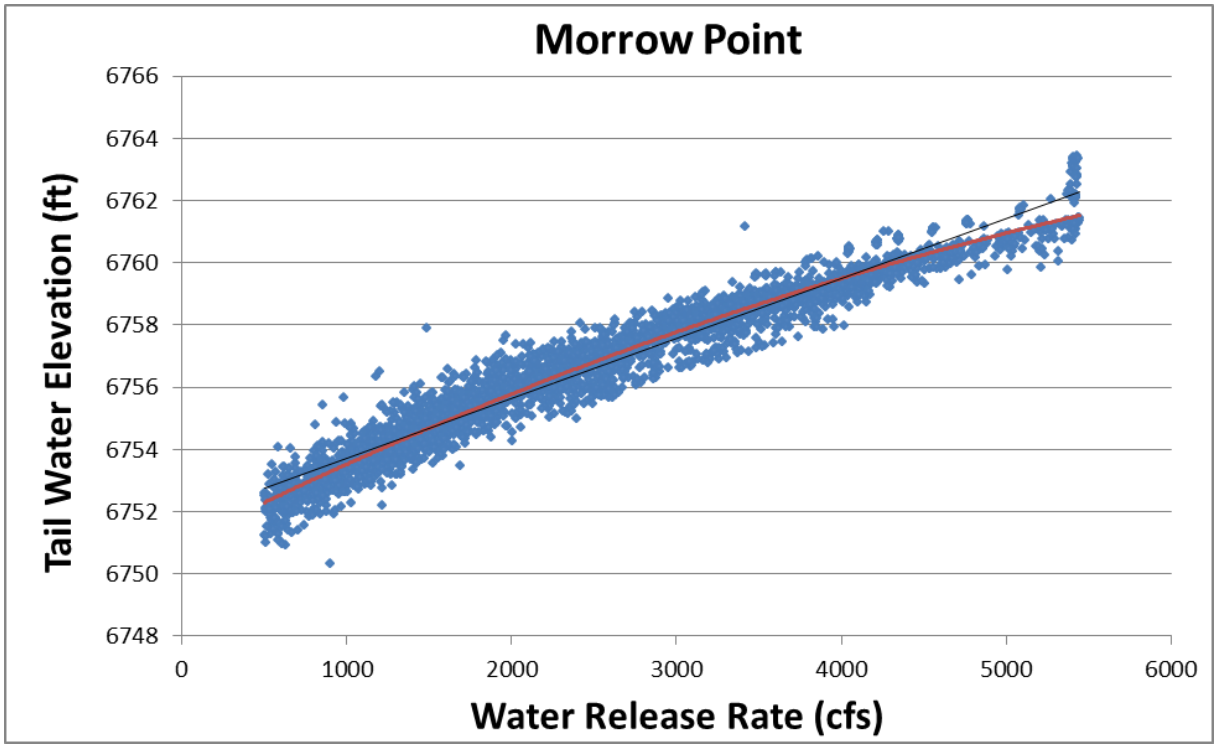

FIGURE 4-63 Morrow Point Tail Water Elevation as a Function of Water Release Rate

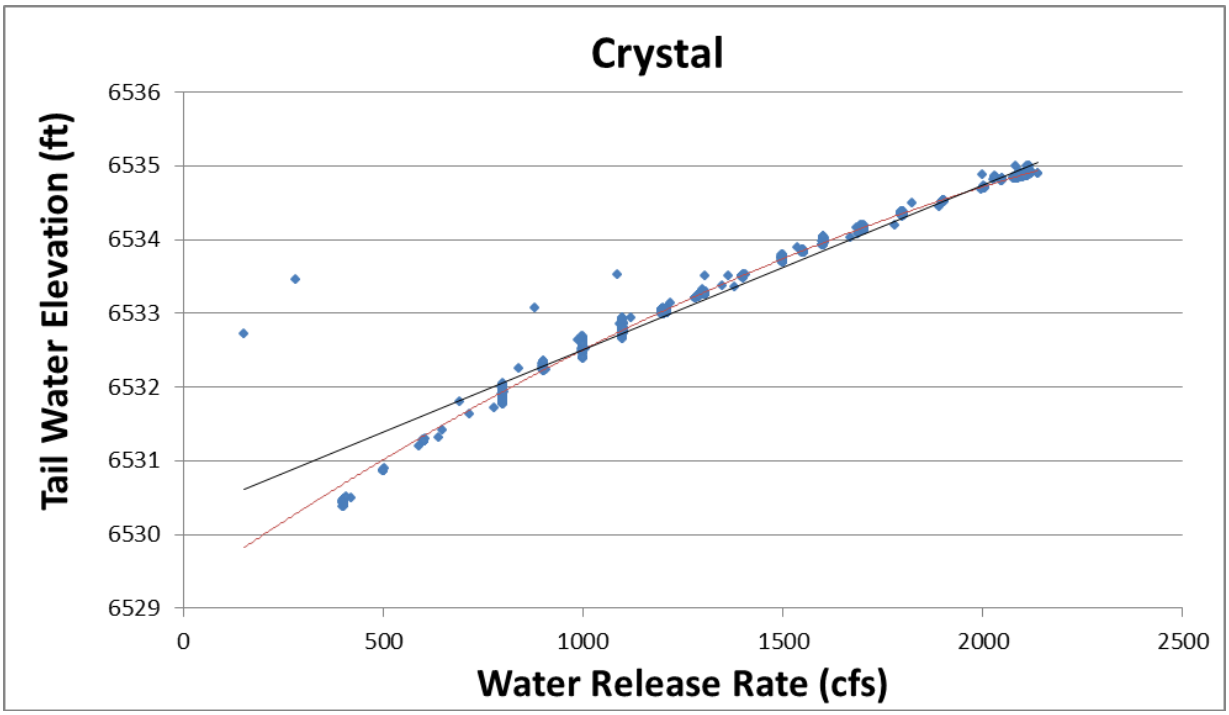

FIGURE 4-64 Crystal Tail Water Elevation as a Function of Water Release Rate

TABLE 4-10 Polynomial Coefficients for the Tail Water Elevation Function

Tail	Intercept	1st Order	2nd Order	3rd Order	4th Order	5th Order	6th Order
Blue Mesa	7.15E+03	3.73E–02	–5.52E–05	4.30E–08	–1.77E–11	+3.67E–15	–3.04E–19
Blue Mesa (Linear)	7.16E+03	2.09E–03	NA[a]	NA	NA	NA	NA
Morrow Point	6.75E+03	2.66E–03	–1.33E–07	NA	NA	NA	NA
Morrow Point (Linear)	6.75E+03	1.93E–03	NA	NA	NA	NA	NA
Crystal	6.53E+03	3.76E–03	–5.15E–07	NA	NA	NA	NA
Crystal (Linear)	6.53E+03	2.23E–03	NA	NA	NA	NA	NA

[a] NA = not applicable.

Observations for Blue Mesa during the month of April, 2009, are shown in Figure 4-65. Orange and blue data points in the figure, along with associated polynomial curves, represent the turbine water release rate and power production relationships when either unit 1 or unit 2 operate alone. Green data points and the green curve show the relationship when both units are running at the same time. Observations with flow rates less that 500 cfs are excluded from this relationship. As discussed previously, such occurrences take place when units are either ramping down to or up to a zero generation set point. Note that curves for unit 1 and unit 2 are almost indistinguishable, indicating that these two units have nearly identical performance characteristics. Table 4-11 shows polynomial coefficients for these three Blue Mesa power functions.

Figure 4-66 shows Blue Mesa turbine water to power conversion efficiencies as a function of turbine water release rate for the same April 2009 period. Turbine efficiency curves for both units display the classic concave pattern in which efficiencies are relatively low at slow water release rates and increase as the turbine flow rate becomes faster. At about 80 to 90% of the maximum turbine flow rate, the efficiency peaks. As water releases increase above this point the power generation increases, but more water is released per incremental MWh of electricity produced. When both turbines operate the maximum tribune efficiency may be slightly lower than when each unit operates alone. Table 4-12 shows polynomial coefficients for the three Blue Mesa efficiency curves.

Turbine power-flow rate curves and associated power efficiency curves were also constructed for units at Morrow Point. These functions are shown in Figures 4-67 and 4-68, respectively. Function polynomial coefficients are provided in Tables 4-13 and 4-14. Note that Morrow Point units convert water to power more efficiently than Blue Mesa units. This is because Morrow Point has a taller dam structure and therefore a higher hydraulic head than Blue Mesa.

Blue Mesa and Morrow Point power functions, along with associated coefficients, were constructed for each of the 6 months of the retrospective study period. For each unit there were only minor differences in curve magnitudes and shapes among study period months. These differences are mainly attributed to changes in monthly forebay elevation levels.

An examination of historical operations during the study period reveals that turbines are frequently operated at sub-optimal efficiency points. Note that in Figures 4-66 and 4-68 Blue Mesa and Morrow Point units operate at either much lower or somewhat higher levels than the most efficient point. There are also times when it would be more efficient to operate only one unit instead of two. A cursory examination of other historical months reveals a very similar pattern.

There are many reasons why turbines are operated at sub-optimal generation points. These include, but are not limited to, minimum release rate requirements associated with environmental operating criteria, energy and ancillary market price drivers, and load and resource balance requirements.

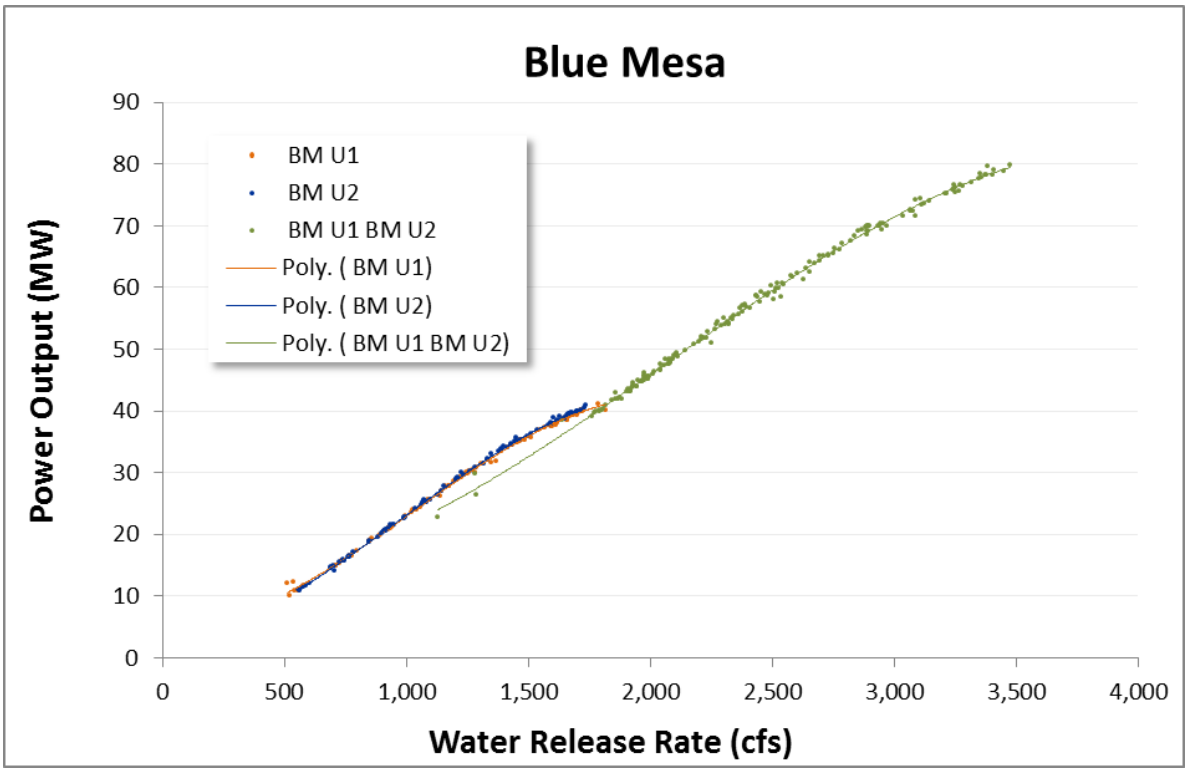

FIGURE 4-65 Power Output as a Function of Water Release Rate, April 2009

TABLE 4-11 Polynomial Coefficients for the Conversion Factor Function, Blue Mesa, April 2009

Power Output	Intercept	1st Order	2nd Order	3rd Order
Unit 1	2.04E–02	–4.59E–06	1.20E–08	–4.87E–12
Unit 2	1.34E–02	1.24E–05	–9.71E–10	–1.65E–12
Unit 1 Unit 2	2.06E–02	–2.05E–06	2.75E–09	–5.70E–13

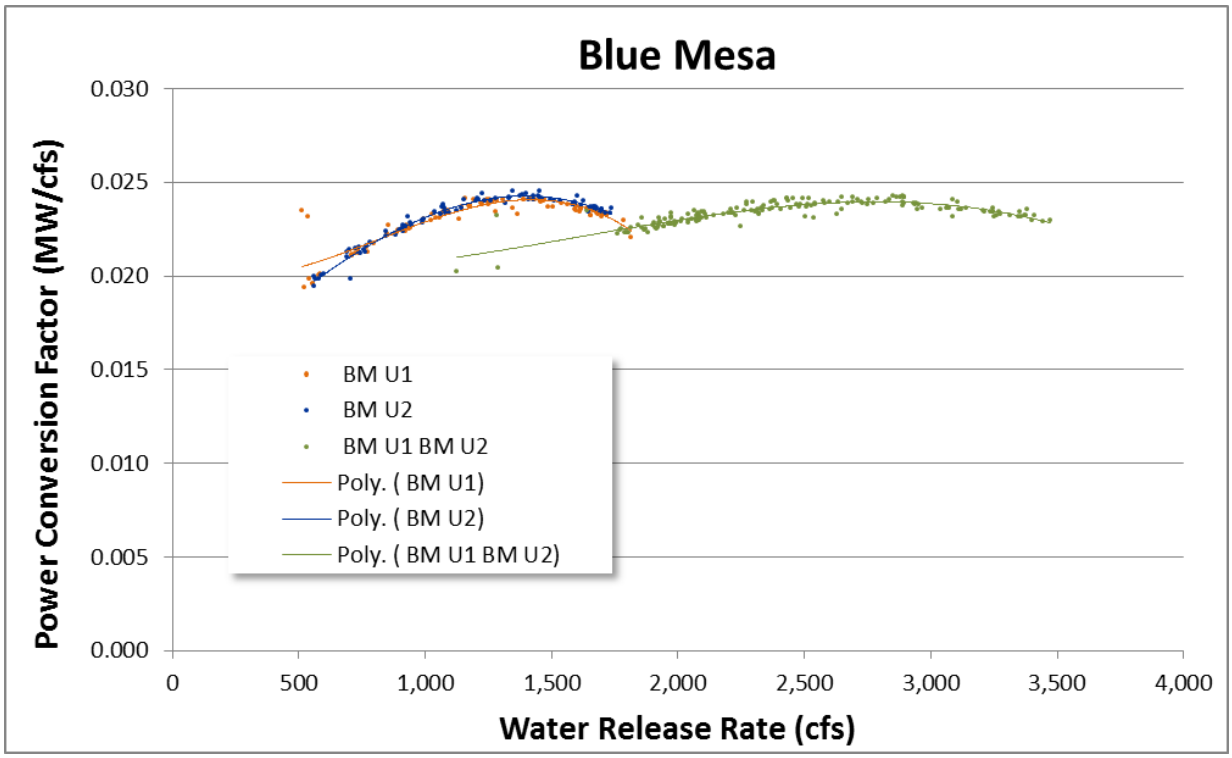

FIGURE 4-66 Power Conversion Factor as a Function of Water Release Rate, April 2009

TABLE 4-12 Polynomial Coefficients for the Blue Mesa Power Output Function, April 2009

Power Output	Intercept	1st Order	2nd Order	3rd Order
Unit 1	5.03E+00	–1.62E–03	+2.90E–05	–9.46E–09
Unit 2	1.89E+00	5.50E–03	2.40E–05	–8.30E–09
Unit 1 Unit 2	1.34E+01	–5.29E–03	+1.59E–05	–2.57E–09

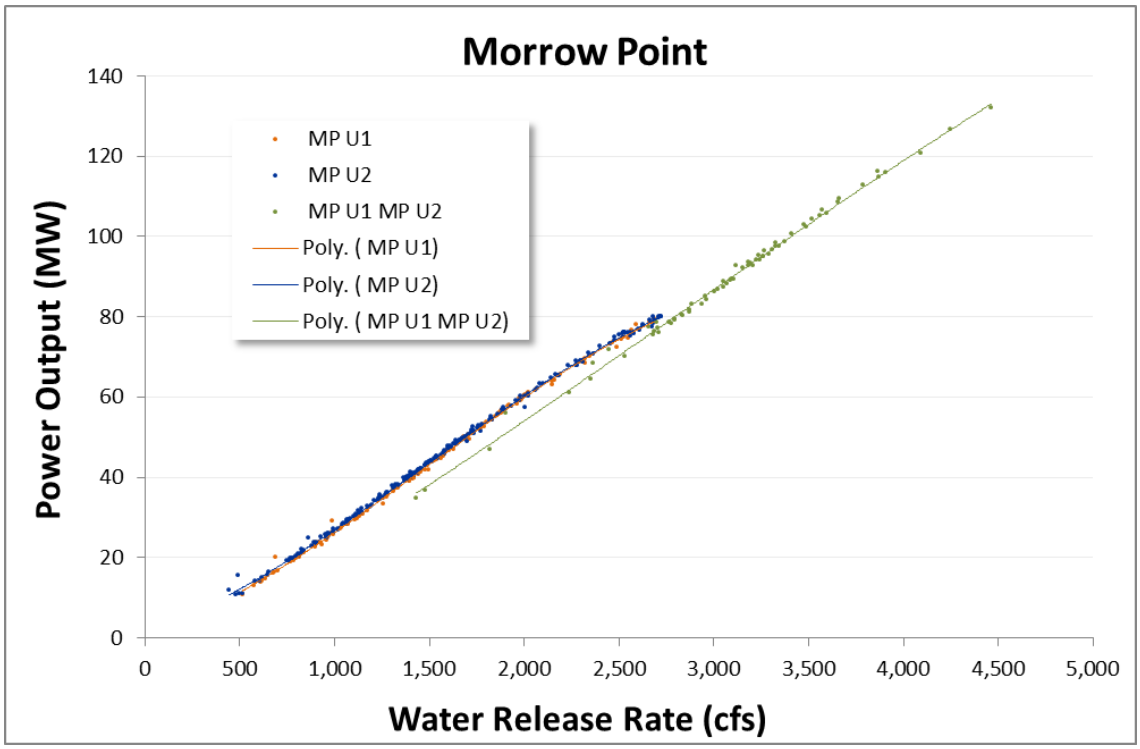

FIGURE 4-67 Power Output as a Function of Water Release Rate, April 2009

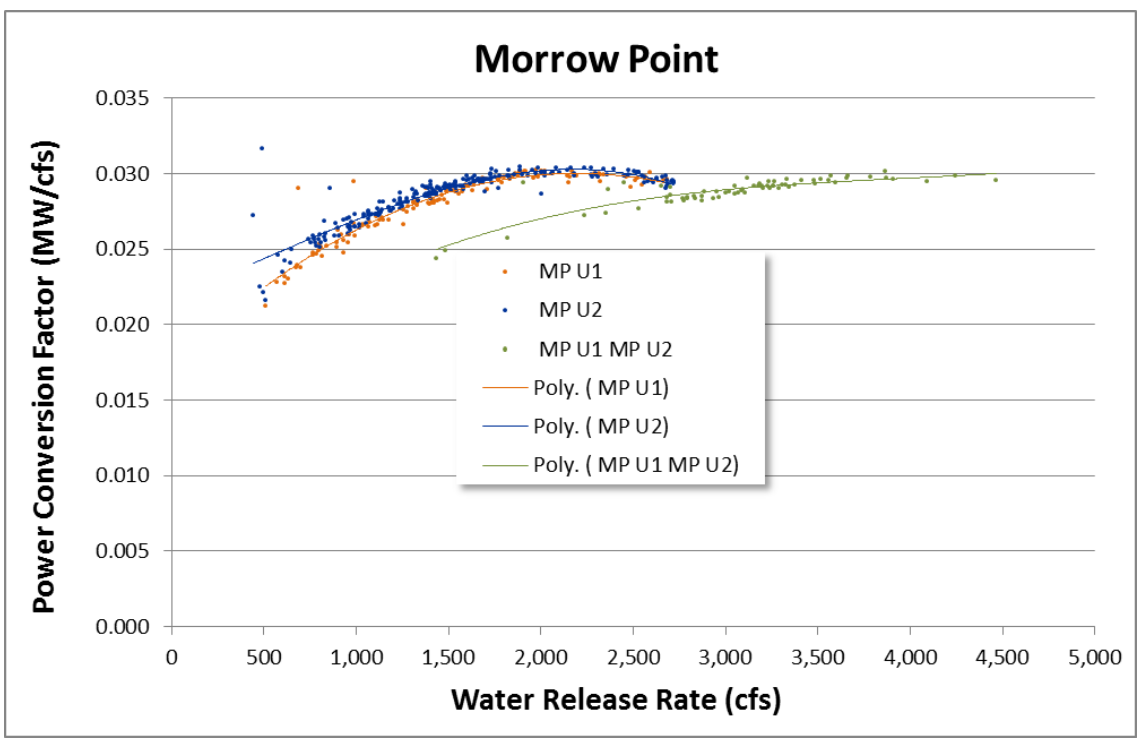

FIGURE 4-68 Power Conversion Factor as a Function of Water Release Rate, April 2009

TABLE 4-13 Polynomial Coefficients for the Power Output Function, Morrow Point, April 2009

Power Output	Intercept	1st Order	2nd Order	3rd Order
Unit 1	4.31E–01	1.63E–02	1.23E–05	–2.79E–09
Unit 2	1.25E+00	1.65E–02	1.18E–05	–2.66E–09
Unit 1 Unit 2	–4.36E+00	2.47E–02	2.97E–06	–3.59E–10

TABLE 4-14 Polynomial Coefficients for the Conversion Factor Function, Blue Mesa, April 2009

Power Output	Intercept	1st Order	2nd Order	3rd Order
Unit 1	1.73E–02	1.16E–05	–2.68E–09	5.91E–15
Unit 2	2.18E–02	4.84E–06	9.25E–10	–6.24E–13
Unit 1 Unit 2	1.51E–02	9.84E–06	–2.32E–09	1.95E–13

One objective of the WUOT toolset is to help schedulers and operators generate more power with the same amount of water resources (i.e., operate at higher efficiency points). Therefore, a total energy production metric is included in the CHEERS Results Report Writer to indicate whether the model is a useful tool for increasing energy efficiency at hydropower plants.

Piecewise Linearization of Generation and Power Conversion Functions. As discussed in more detail in a previous section, due to the complexity of demonstration site operations and the current state of MINLP technology, CHEERS must formulate equations as MILP problems. Therefore, the nonlinear power functions described above are converted to piecewise linear approximations. This is accomplished by the following two-step process:

1. For each monthly turbine water flow rate/power function, a vector of 40 equally spaced *x*, *y* data points are generated by the CHEERS Interpolation Routine; and

2. These points are the input into the Piecewise Linear Optimizer, which determines the slopes and intercepts of line segments that yield the minimum sum-of-square error between the nonlinear function and the piecewise linear approximation.

Currently, the Evolutionary Algorithm option in the Excel Solver Algorithm is used to locate representative breakpoints. As CHEERS development progresses, a second option that performs the same task will be added. Developed by the Mathematics and Computer Science division of Argonne for this project, it uses an adaptively refined dynamic programming approach (see Appendix F).

The piecewise linear curves for power output and power conversion factor are represented in Figures 4-69 and 4-70. In the figures, blue curves represent the power output and red curves represent the power conversion factor relationships. Breakpoints and the number of *x*/*y* points in each segment are shown in Tables 4-15 and 4-16. The same piecewise linear curve is used for both Blue Mesa units. However, two different curves are generated for Morrow Point because its two units operate differently.

In all cases, the first point in the piecewise linear function represents the minimum unit-level generation point. There is also additional information associated with these curves that simulates wicket gate leakage. Based on correspondence with Reclamation staff, leakage is set equal to 0.5% of a unit's maximum turbine flow rate and therefore increases as reservoir elevations become higher.

Typically, water release rates from the Crystal Reservoir are required to remain constant each day. However, in order to attain a different flat flow rate, Reclamation staff requests prescribed changes in hourly flow rates over the course of a day. In the CHEERS model, generation produced by the only unit at Crystal under flat flow conditions is computed by a function that, for a constant flow rate, relates generation to hydraulic head. Figure 4-71 shows a family of curves that relate power output to hydraulic head. When the flat flow rate is at a level that differs from one of these curves, CHEERS interpolates between two adjacent curves. Table 4-17 shows intercepts and slopes for Crystal linear power production functions.

Similar functions constructed for the Blue Mesa and Morrow Point units are presented in Appendix H. However, since reservoir elevations at these two sites do not change significantly over a day and the power plants are located at high head dams, the inclusion of the hydraulic head dimensions was not undertaken at this time. However, refinements that include these functions will be included in a future set of Aspinall Cascade analyses.

Short-Term Market Prices. Short-term market prices input to the CHEERS day-ahead STF and balancing spot market nodes are based on Intercontinental Exchange (ICE) market hub data along with average daily prices for several WECC market hubs as reported on the CRSP Web site (http://www.wapa.gov/crsp/opsmaintcrsp/default.htm). Spot price profiles for the 6-month study period from the ICE contain daily minimum and maximum prices for both on- and off-peak hours. These prices are given on the left side of Figure 4-72.

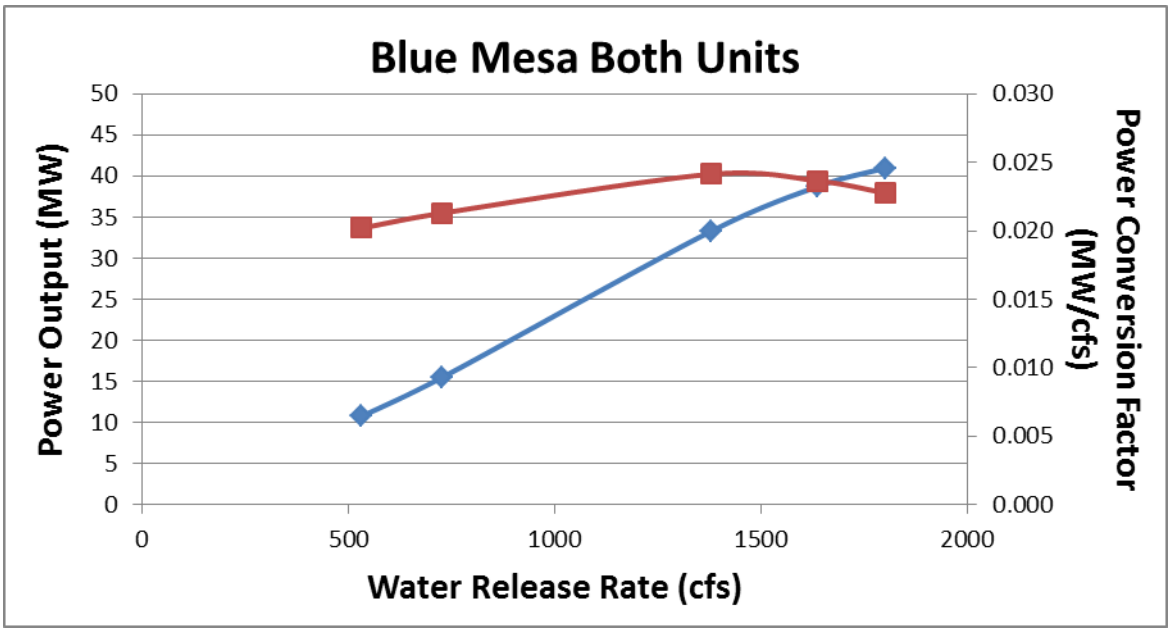

FIGURE 4-69 Piecewise Linear Curves for Power Output and Power Conversion Factor, Blue Mesa

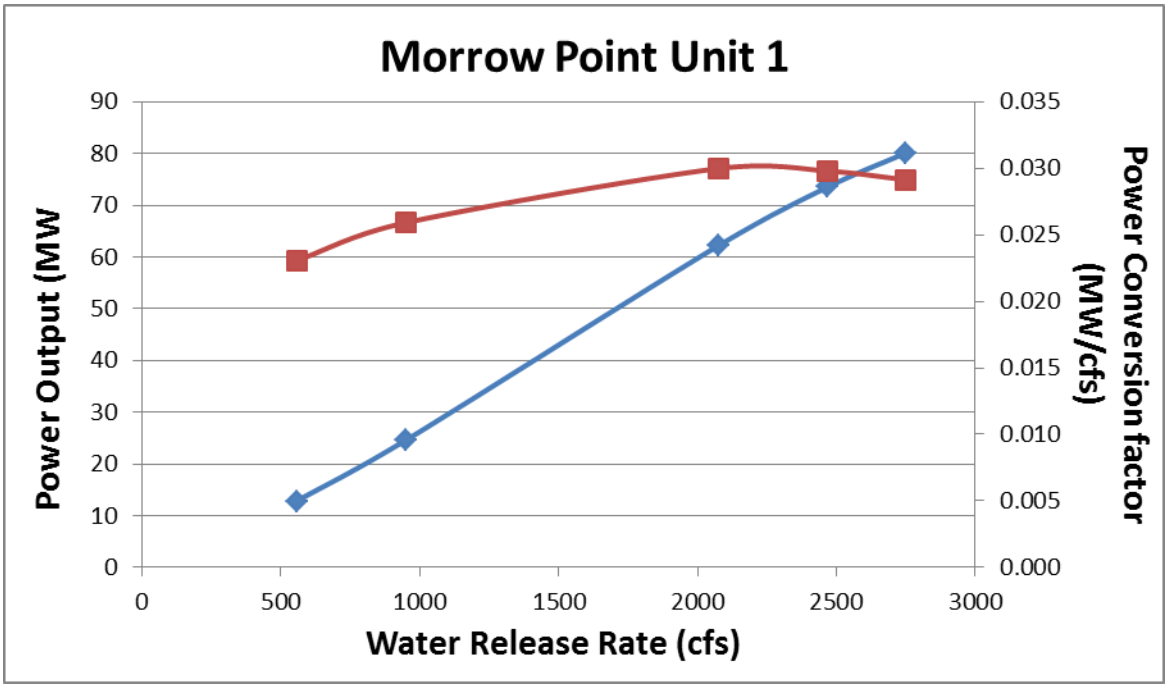

FIGURE 4-70 Piecewise Linear Curves for Power Output and Power Conversion Factor, Morrow Point

TABLE 4-15 Piecewise Linear Curve Breakpoints, Blue Mesa

Number of Points in Segment	Water Release Rate (cfs)	Power Output (MW)	Conversion Factor (MW/cfs)	Incremental Conversion Factor[a] (MW/cfs)
1	532.50	10.75	0.02019	0.02010
6	727.50	15.49	0.02130	0.02317
20	1377.50	33.24	0.02413	0.02771
8	1637.50	38.72	0.02365	0.02193
6	1800.00	40.98	0.02277	0.01435

[a] The incremental conversion factor is computed by taking the ratio of power output difference to water release rate difference for each segment (i.e., it is the slope of power output curve for the given segment).

TABLE 4-16 Piecewise Linear Curve Breakpoints, Morrow Point

Number of Points in Segment	Water Release Rate (cfs)	Power Output (MW)	Conversion Factor (MW/cfs)	Incremental Conversion Factor (MW/cfs)
1	556.25	12.83	0.02306	0.02262
7	950.00	24.62	0.02591	0.02875
20	2075.00	62.21	0.02998	0.03374
7	2468.75	73.54	0.02979	0.02947
6	2750.00	80.11	0.02913	0.02391

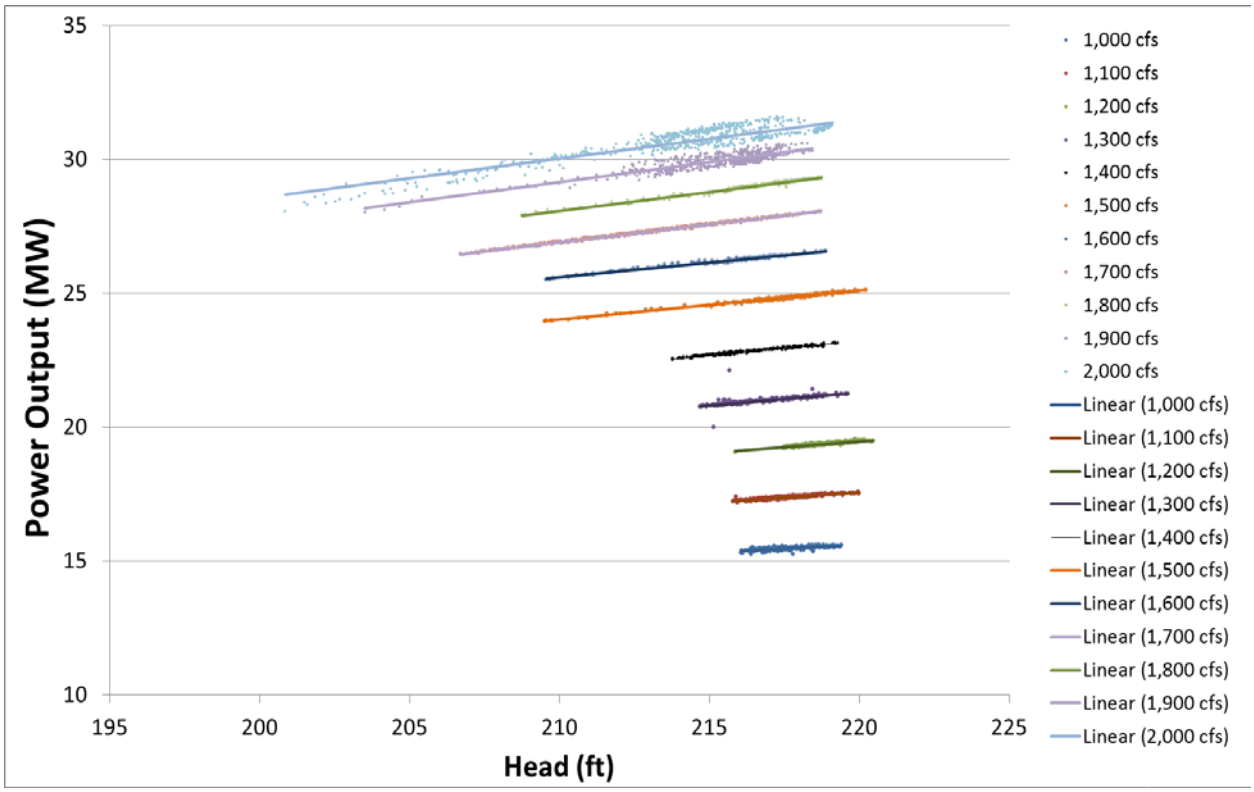

FIGURE 4-71 Crystal Unit 1 Power Outputs as a Function of Turbine Release Rate and Head

TABLE 4-17 Intercepts and Slopes for Crystal Head-Generation Linear Approximations

Water Release Rate (cfs)	Intercept	1st Order
1000	1.98E+00	6.19E–02
1100	9.36E–01	7.55E–02
1200	–7.50E–01	9.19E–02
1300	1.64E–01	9.60E–02
1400	–2.17E–01	1.06E–01
1500	1.30E+00	1.08E–01
1600	3.04E+00	1.07E–01
1700	–9.32E–01	1.33E–01
1800	– 1.58E+00	1.41E–01
1900	– 2.04E+00	1.48E–01
2000	– 8.25E–01	1.47E–01

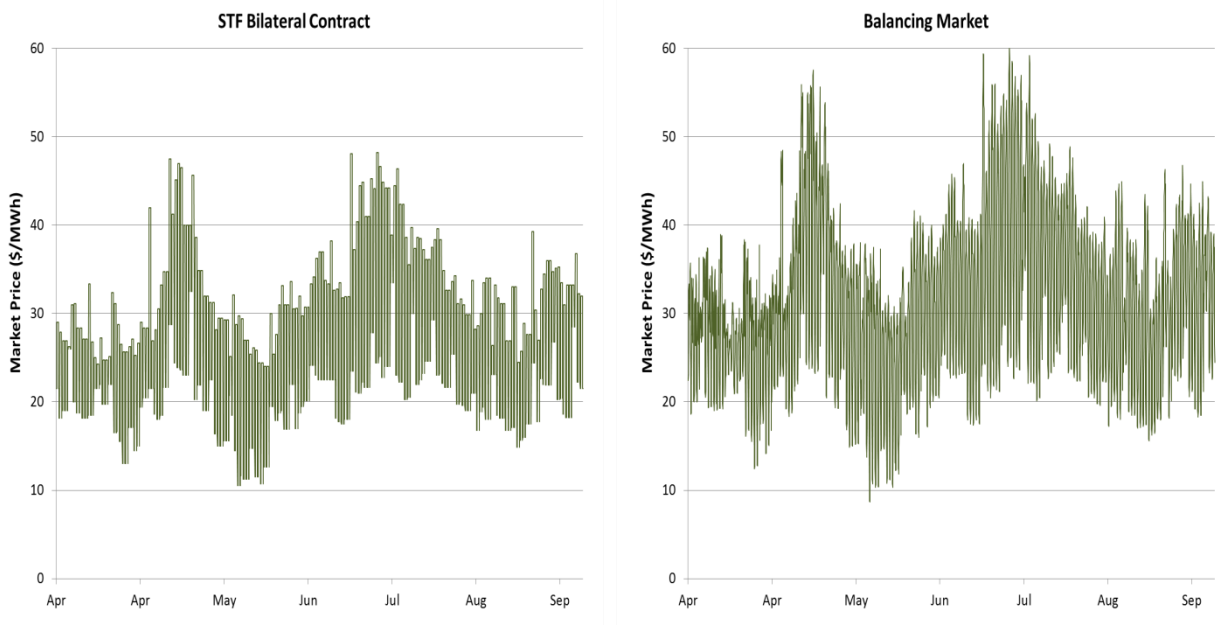

FIGURE 4-72 Market Prices Used for the CHEERS Retrospective Analysis

The hourly electricity prices used for real-time balancing market transactions given on the right side of Figure 4-72 are estimated using a multistep process that produces estimates benchmarked to ICE daily on- and off-peak prices and have a diurnal shape that is consistent with grid loads. Schedulers at the CRSP EMMO indicated that real-time prices tended to be higher than day-ahead STF prices. Therefore it was assumed that on a daily average real-time prices are 10% more expensive than day-ahead prices. Appendix I provides a more detailed description of the method used to perform this hourly price shaping.

It should be noted the staff in Western's EMMO use ICE information along with other information to make daily scheduling decisions. Actual purchase prices may differ significantly from ICE. However, ICE provides a consistent and continuous price benchmark that is indicative of the marginal economic value of energy at specific locations within the western power grid interconnection.

CRSP Comparison of CHEERS Results to Actual Operations. The CHEERS model was run for all hours during the 6-month retrospective study period from April 1, 2009, through September 30, 2009. Model runs implement business practices that were employed during this period. This section discusses some of the model results and compares the results to actual hydropower operations, STF market transactions, and power economics. For this set of runs, it was assumed that Reclamation plant operators limit each unit to a maximum of two starts/stops per day and that a unit start and stop each cost $200. As discussed later in this document, a sensitivity analysis was performed on these assumptions.

Many CRSP planners and schedulers in the Western organization have decades of experience operating the system. Likewise, Reclamation plant operators and dispatchers are highly skilled. This retrospective analysis indicates that in general both organizations have excellent performance records. However, analyses of CHEERS model results also indicate that there are some areas that could be fine-tuned to further improve CRSP operations.

Below, CHEERS results are compared to actual operations under different hydrological and operational conditions. Highlighted are some situations in which operations could be tweaked to gain additional Aspinall energy production and reduce CRSP net purchase power costs.

Early Spring Energy Short Hydropower Condition. On April 20, 2009, before the rapid mountain springtime snowmelt occurred, CRSP hydropower resources plus energy purchased under LTF contracts were insufficient to meet all of Western's load obligations. This situation resulted in very accurate day-ahead predictions and day-ahead scheduling that was very similar to actual real-time operations. Therefore, for this day only actual and real-time comparisons are discussed.

The loads and resource balance graph displayed on the left side of Figure 4-73 shows that block STF purchases were needed in order to fill most of this energy-short void. In addition, spot market balancing energy purchases filled the remaining imbalance gaps. The right side of the figure shows CHEERS real-time model results for that same day. These graphs are very similar. However, there are subtle differences that suggest operations may have been slightly improved upon during that day. Improvements include changes in purchasing profiles, plant level scheduling, and unit dispatch.

Figure 4-74 compares actual purchase and sale transactions (left side) on April 20, 2009, to CHEERS optimization results (right side). The model suggests that higher off-peak STF purchases could have been made in order to nearly eliminate more expensive balancing purchases.

There are a few caveats that should be noted in the construction of "actual" operations. Actual STF purchases quantities are known from historical records. In addition, actual day-ahead balancing quantities are also recorded. These balancing transactions are often very similar, if not identical, to actual spot market balancing purchases. However, as changes occur in real time, balancing purchases need to be adjusted. In this analysis we inferred what changes would have occurred in order to balance the system in real-time. It should also be noted that actual operations for other supply aspects were available and utilized, including plant- and unit-level generation, LTF agreements, losses, and power interchanges.

"Actual" CRSP costs and revenues are rough approximations based on ICE prices. Although actual prices are available, there are also insufficient CRSP price data for modeling purposes since the model suggests transactions that did not actually occur; that is, prices are only available for actual transactions.

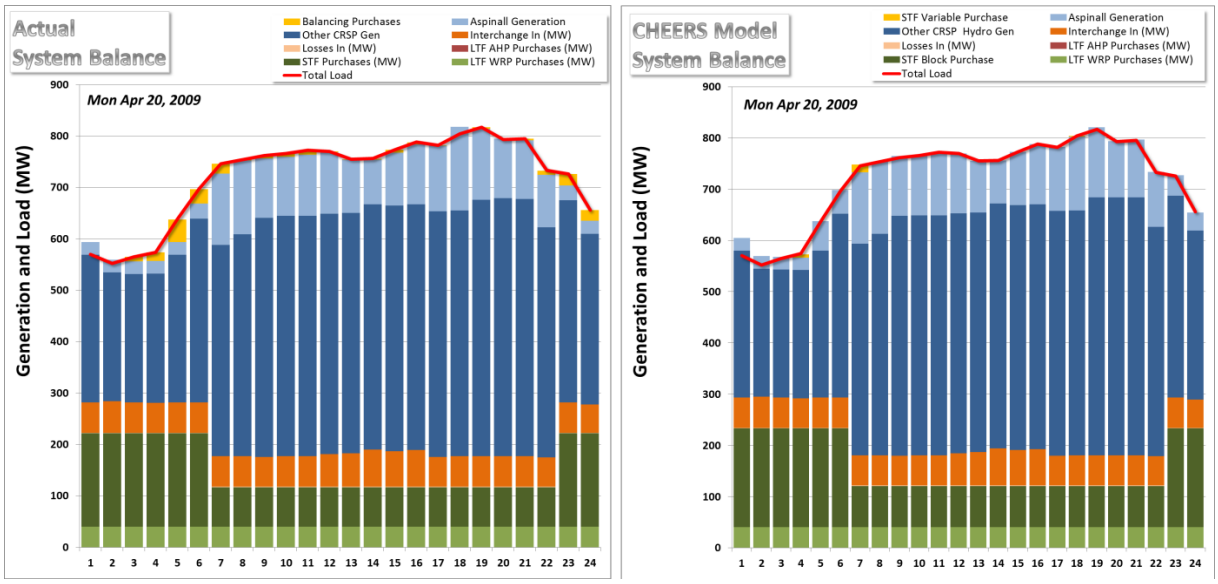

FIGURE 4-73 Actual and CHEERS Real-Time CRSP System Energy Balances (April 20, 2009)

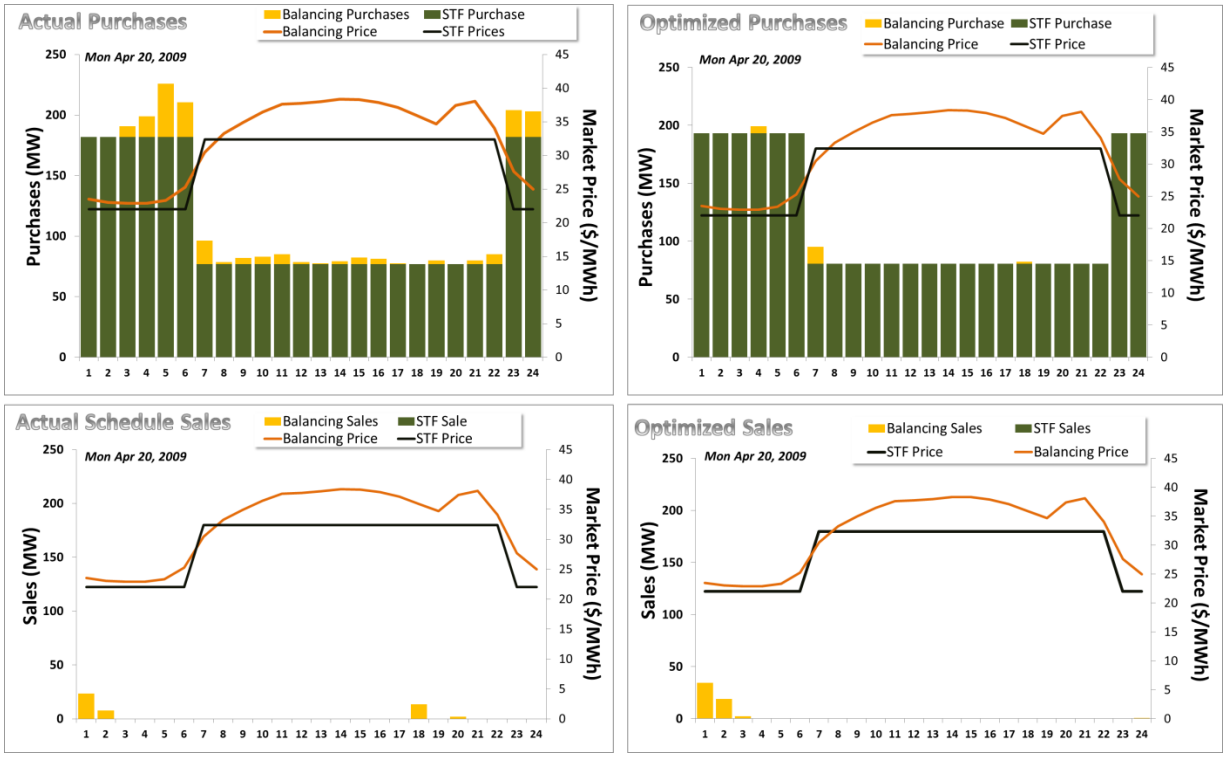

FIGURE 4-74 Actual CRSP and CHEERS Real-Time Power Market Transactions (April 20, 2009)

Tables 4-18 and 4-19 show daily statistics for both actual operations and modeled results. CHEERS purchase amounts and costs are both slightly less than the ones that "actually" occurred. As described in more detail later in this section, lower net purchases quantities are required since slightly more power can be generated from Aspinall power plants with the same quantity of reservoir water releases. In addition, modeled purchase costs are about $1,400 lower; that is a reduction of 1.8%.

Economic computations are based on the caveats discussed above and assume that buyers and sellers of power would have been able to agree upon both the quantities suggested by CHEERS and the assumed prices used for this analysis. Therefore, savings are only rough approximations. Fundamental conclusions are nonetheless valid. Lowering spot purchases by slightly increasing off-peak STF block purchases will in general result in purchase savings. These savings are enhanced by more efficient use of water resources, whereby more power can be generated through a more efficient schedule and dispatch.

Figure 4-75 shows actual and modeled power plant generation. The left side of the figure shows total actual generation from Aspinall hydropower resources. CHEERS model results shown on the right side of the figure suggest a very similar pattern in terms of total hourly Aspinall power generation to actual generation levels. However, CHEERS plant-level production patterns for both Blue Mesa and Morrow Point are very different. Note that in the actual Morrow Point schedule the plant has zero production for 2 hours in the middle of the day. In addition, at the levels scheduled, this will require that both units will need to be turned off at hour 13 and turned back later in the afternoon. The CHEERS model results suggest that some of these unit shutdowns and startups could be avoided by altering the day-ahead schedules. It should also be noted that CHEERS real-time optimization results have significantly less up and down generation ramping than the actual operations.

TABLE 4-18 Comparison of Actual Purchase and Sales Quantities with CHEERS Real-Time Results (April 20, 2009)

	Actual Transactions (MWh)			CHEERS Real-Time Transactions (MWh)		
Market	Purchase	Sell	Net Purchase	Purchase	Sell	Net Purchase
Balancing	210	47	163	22	57	–35
STF	2,688	0	2,688	2,836	0	2,836
Total	**2,898**	**47**	**2,851**	**2,858**	**57**	**2,801**

TABLE 4-19 Comparison of Actual Purchase Costs and Sales Revenues with CHEERS Real-Time Results (April 20, 2009)

	Costs Based on Actual Energy Transaction and ICE Prices ($)			CHEERS Real-Time Transaction Costs ($)		
Market	Purchase	Sell	Net Cost	Purchase	Sell	Net Cost
Balancing	5,849	1,295	4,554	633	1,329	-695
STF	71,918	0	71,918	75,779	0	75,779
Total	77,767	1,295	76,472	76,412	1,329	75,084

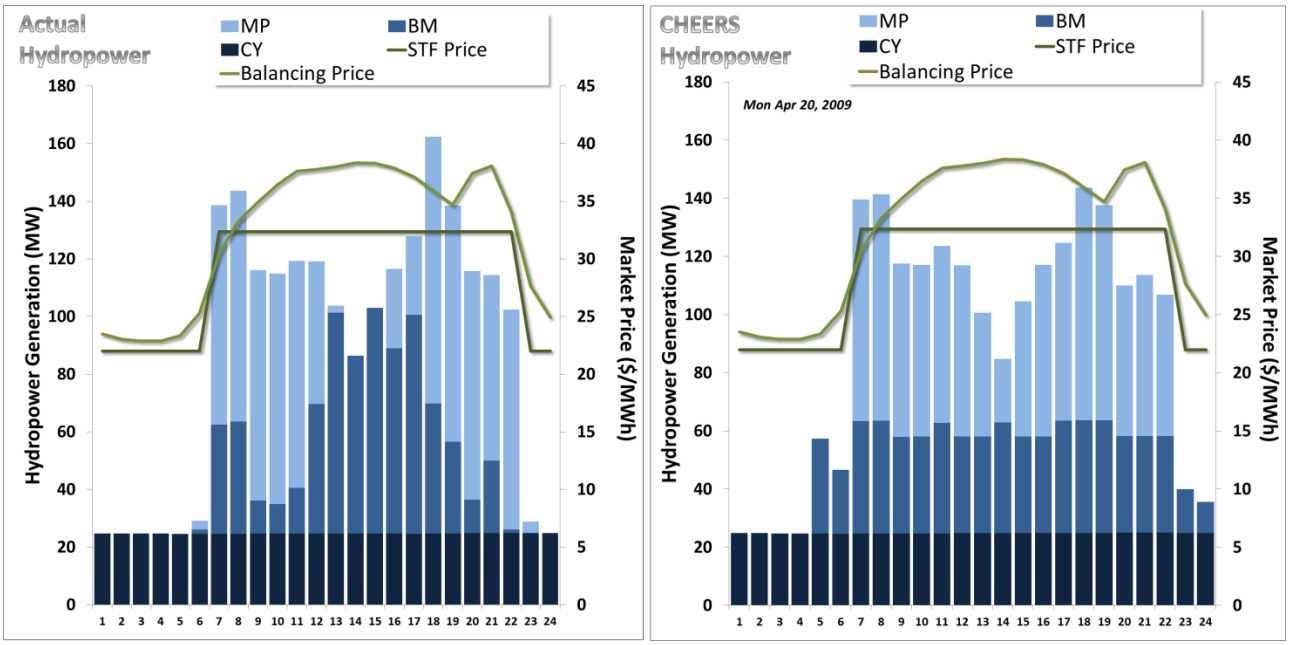

FIGURE 4-75 Comparison of Actual Aspinall Plant-Level Generation with CHEERS Real-Time Optimization Results (April 20, 2009)

A finer granularity view of operations at the unit level rather than plant level provides greater insights into a more efficient unit commitment and hydropower operation. CHEERS model results provided in Figure 4-76 show that much of the time units at Blue Mesa operate at an output level of 33 to 39 MW. As shown in Table 4-15, this is at relatively high conversion efficiencies for Blue Mesa units. During a couple of hours at the end of the day, operations deviate from this optimal efficiency point in order to follow load. Actual Blue Mesa operations show that unit operations are often below or above the most efficient point. Only actual plant-level efficiency is shown in Figure 4-76. Turbine-level flows are not currently measured.

The most efficient operating point for each of the Morrow Point units is about 61 MW, with relatively good efficiencies in the 50 to 75 MW operating range. During hours 7 through 22, CHEERS operates units within this high-efficiency range most of the time. During a few midday hours, CHEERS deviates from this range to follow load. For several hours, CHEERS operates units between 59 and 64 MW; that is, at or very near to the highest efficiency point. Conversely, the actual unit operations were outside the range of high turbine efficiencies.

Both actual and modeled operations at Crystal are nearly identical. Release from this reservoir must be constant at a rate prescribed by Reclamation staff. Note the very slight dip in efficiency shown for both actual and modeled operations in the early morning hours. As described later, this is due to a drop in the Crystal forebay elevation during this time. In actual operations, there is also a power conversion efficiency dip in the afternoon, but this is not nearly as pronounced in the CHEERS results. This may be due to differences between actual side flows between Morrow Point and Crystal and the ones assumed by CHEERS. As noted in earlier sections, measuring these side flows is difficult, especially at sub-daily timeframes.

Table 4-20 provides a daily summary of actual and CHEERS generation for April 20, 2009. Because it tends to operate units at higher efficiencies, modeled generation is about 2.3% higher. If CHEERS is deployed, efficiency gains are expected at Blue Mesa and Morrow Point, but the efficiency at Crystal is expected have little if any gain.

It should also be noted that Aspinall units had undergone a total of 6 start/stop cycles. In contrast the CHEERS model only has two start/stop cycles. Assuming that each start and stop costs $200, CHEERS would save an additional $1,600. CHEERS results also show an operational regime that has significantly less ramping. Whereas Aspinall units actually ramped a total of 832 MWh, CHEERS shows less than half that amount, with 401 MWh of ramping.

In addition to modeling hydropower and economics, CHEERS also simulates water system operations. In doing so, it enforces all operating constraints and meets water delivery obligations. As described in more detail in an earlier section, reservoir elevations are computed by a water mass balance equation in conjunction with an elevation-to-water storage function. Network water channels, shown as blue links in Figure 4-57, were used by CHEERS to route water into and out of reservoir storage nodes (blue triangles). For this retrospective analysis, evaporation was set to zero because historical information is not available and it is assumed to be very small for these reservoirs, which are located in cold high elevation environments. When

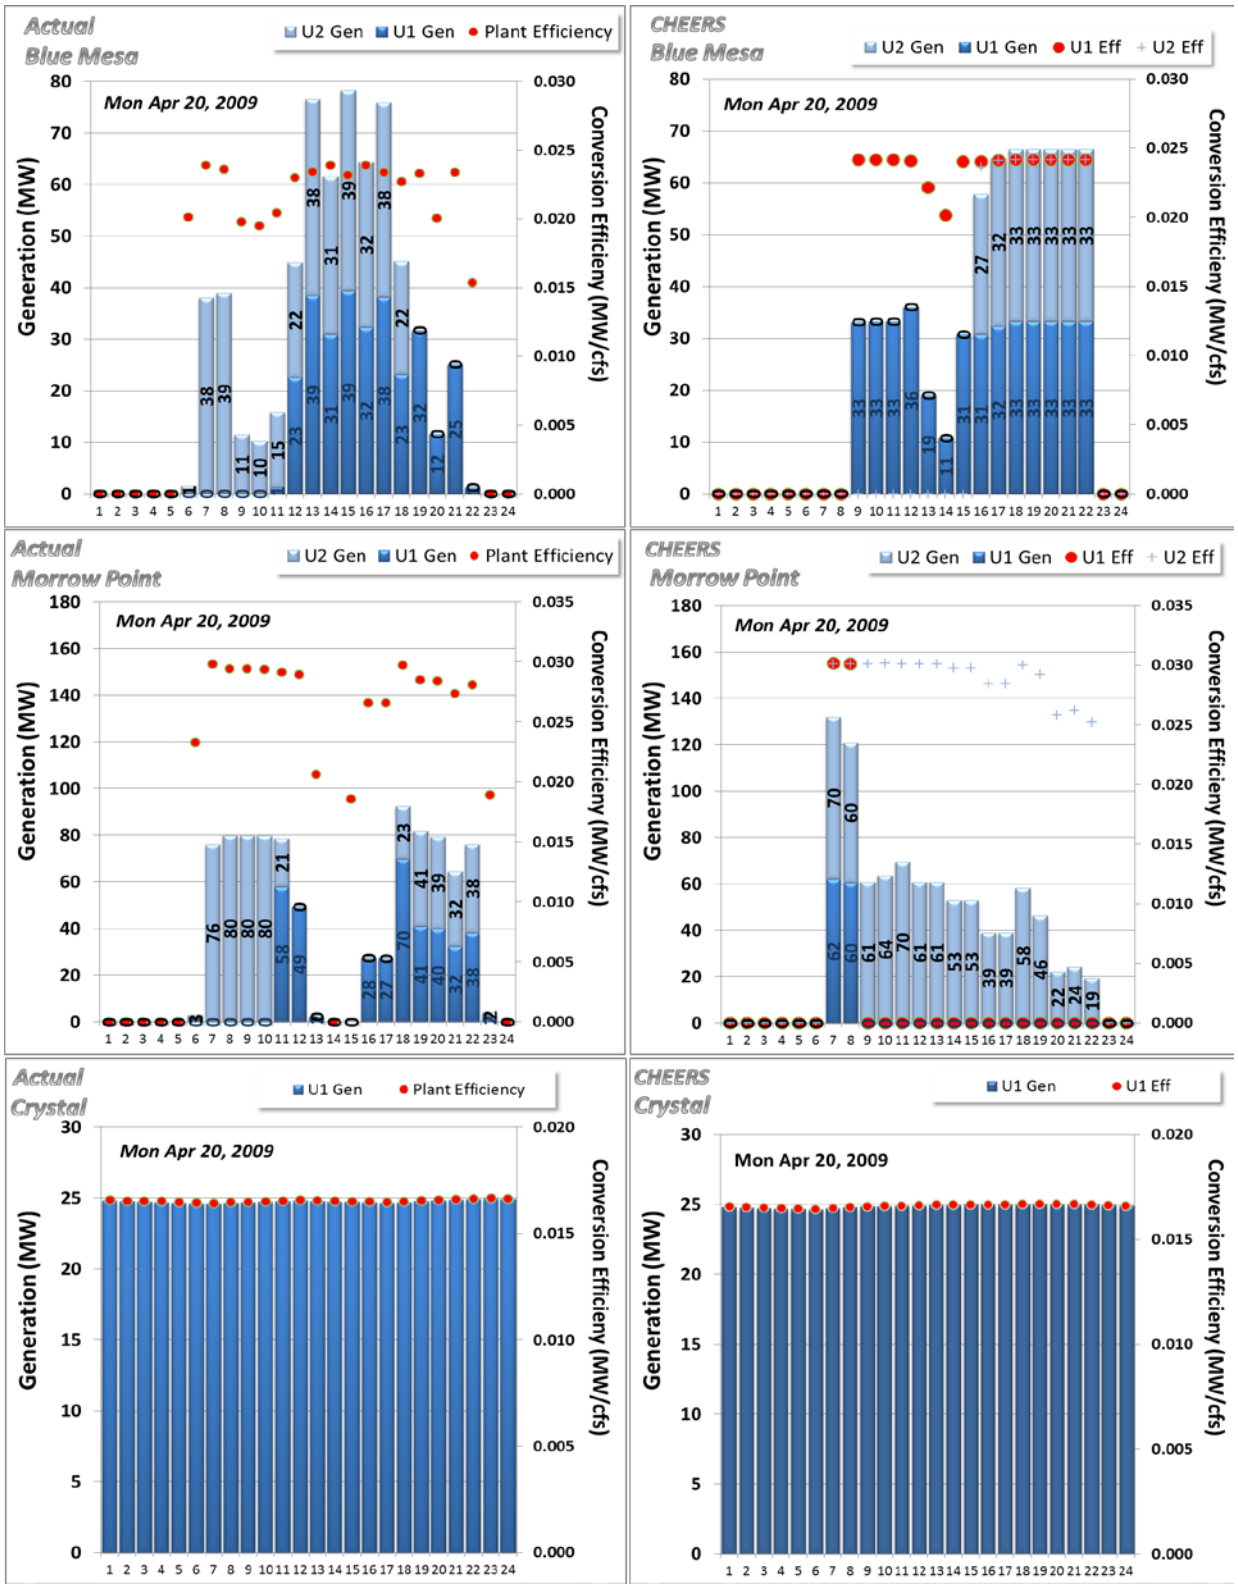

FIGURE 4-76 Comparison of Actual Apsinall Unit-Level Generation with CHEERS Real-Time Results (April 20, 2009)

TABLE 4-20 Daily Summary of Actual and CHEERS Real-Time Generation (April 20, 2009)

	Actual Generation (MWh)			CHEERS Real-Time Generation (MWh)		
Unit	BM	MP	CY	BM	MP	CY
Unit 1	296	515	594	0	0	595
Unit 2	337	388	NA[a]	694	933	NA
Total	633	903	594	694	933	595

[a] "NA" = "not applicable."

CHEERS is run in an integrated mode, evaporation amounts in 6-hour increments are obtained from the seasonal optimization tool.

Figure 4-77 shows CHEERS computed values for water releases, inflows, side flows, and reservoir elevations in the Aspinall water cascade on April 20. The Blue Mesa Reservoir elevation increases throughout the 24-hour period. Since reservoir releases are very small until hour 4, the rise in elevation is relatively fast in the nighttime period. Note the 7 cfs water release per generating unit during this period (the value on the graph for unit 1 is printed on top of the value for unit 2). These water releases do not produce any power, but occur because wicket gates have minor leakages in the locked position. The water elevation rise is much slower through hour 22 because water releases, although much faster than 14 cfs, are less than Blue Mesa inflows.

Morrow Point Reservoir elevations are a function of the quantity water it receives from Blue Mesa, side flows between the two reservoirs, and both power and non-power releases from the Morrow Point reservoir. From midnight through hour 6, the reservoir elevation increases as it receives water from Blue Mesa and about 209 cfs of side flow waters. Wicket gate water leakage is about 10 cfs per unit, for a total of 20 cfs. Once Morrow Point begins to generate power, reservoir elevation levels begin to drop.

Crystal, the last reservoir in the Cascade, has a constant 1,500 cfs water release rate. Since it receives only 218 cfs in side flows, the Crystal Reservoir declines in elevation until a unit at Morrow Point is turned on. At that point, Crystal Reservoir elevation increases through hour 19, at which point side flows plus waters received from Morrow Point are less than the 1,500 cfs Crystal release. During the day, the reservoir fluctuates from a low of 6,749.6 ft to a high of 6,752.8 ft; that is, there is a 3.2 ft change during the day. This level of fluctuation is permitted under Crystal operating rules that allow for a maximum elevation change of 4 ft. In addition, CHEERS operated all reservoirs such that elevations were within minimum and maximum allowable limits.

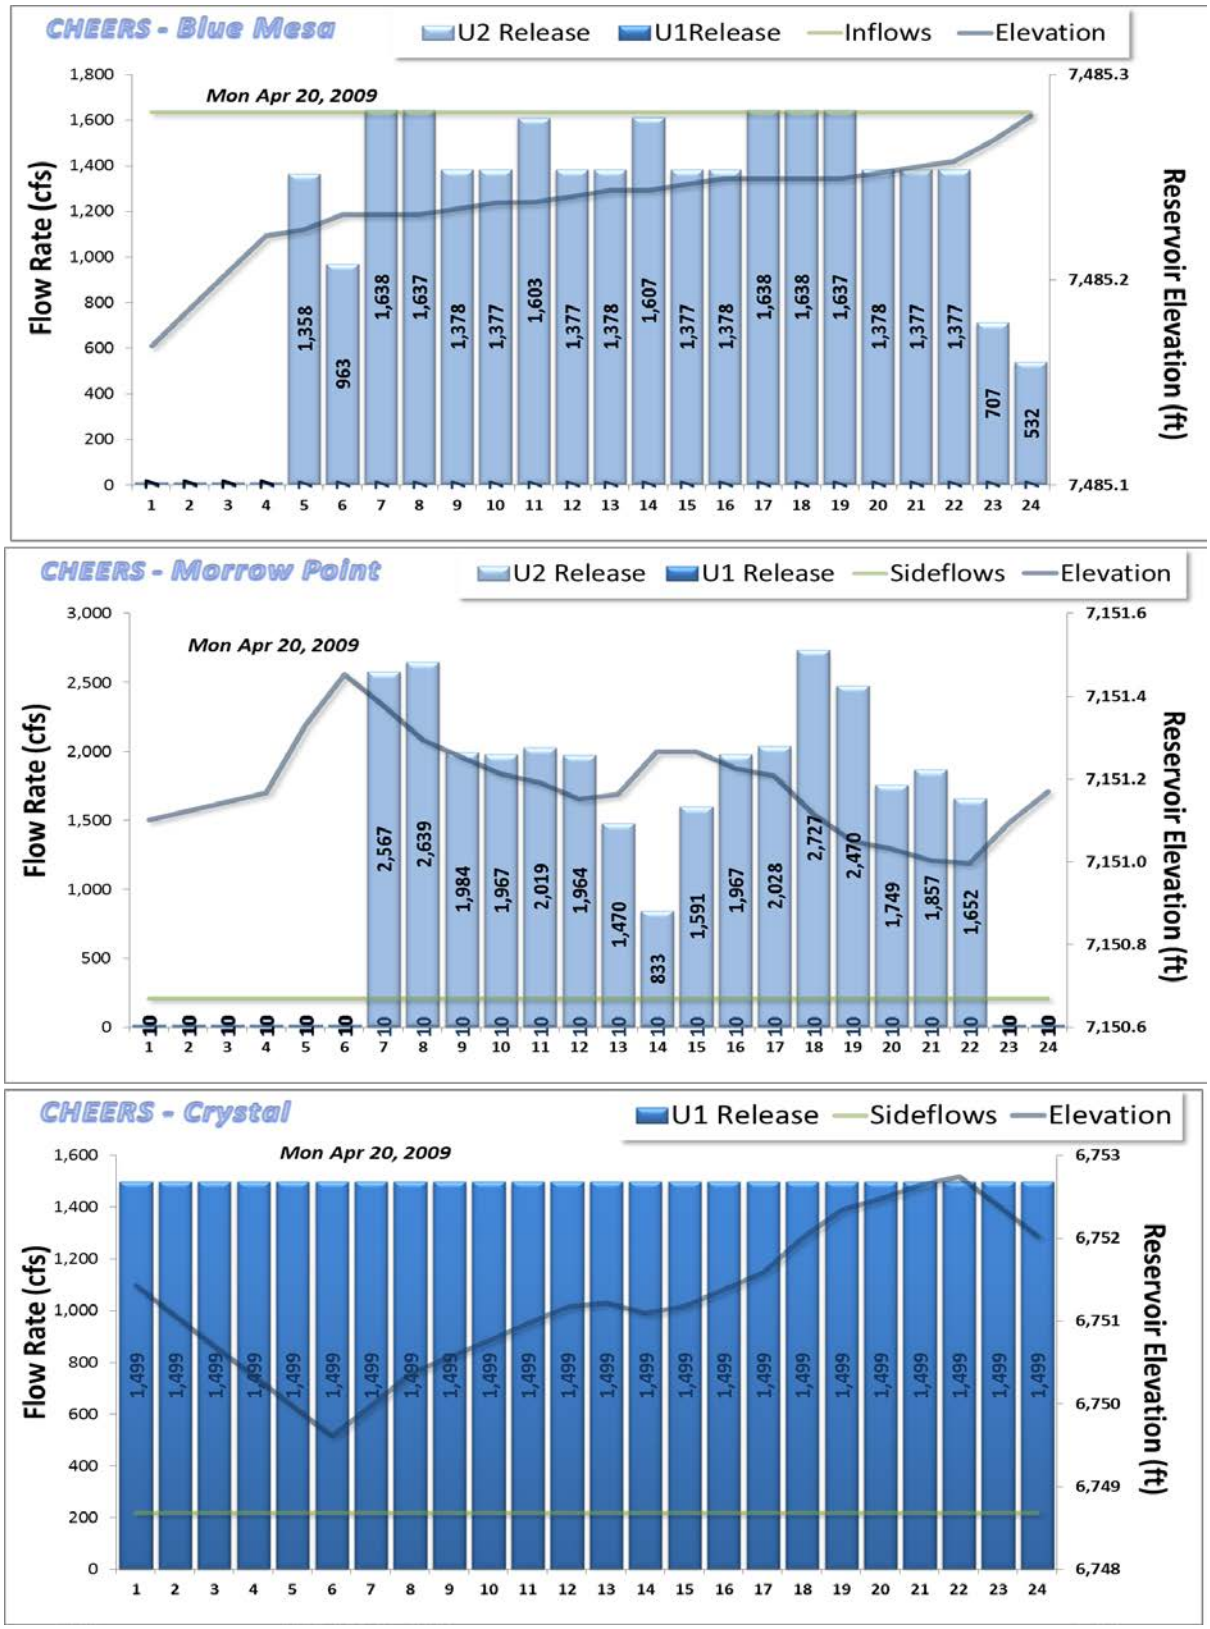

FIGURE 4-77 CHEERS Aspinall Water Releases and Reservoir Elevations (April 20, 2009)

Late Spring Energy Long Hydropower Condition. On June 18, 2009, when there were high water flows in the Gunnison Basin from mountain springtime snowmelt, CRSP hydropower energy resources exceeded all of Western's load obligations. The loads and resource balance graph displayed on the left side of Figure 4-78 shows the hourly level of this "long" energy position, which actually occurred on June 18, 2009; that is, the difference between the top of the stacked bar supply and the red load line. The right half of the figure shows results from CHEERS. These graphs are very similar. However, there are subtle differences that suggest operations may have been slightly improved during that day. Note that the CHEERS model had slightly lower balancing purchases.

Figure 4-79 compares actual purchase and sale transactions on June 18, 2009, to CHEERS real-time optimization results. The model suggests STF sales that are very similar to the actual amount scheduled. However, CHEERS displays slightly higher on-peak STF sales and less balancing market transactions. Tables 4-21 and 4-22 show daily statistics for both actual operations and modeled results. Note that CHEERS sale amounts and revenues are both slightly higher than what "actually" occurred. Modeled net revenues are about $800 higher.

Figure 4-80 shows actual day-ahead schedules and modeled power plant generation for June 18, 2009. The left side of the figure also shows that total actual generation from Aspinall hydropower resources and CHEERS model results have very similar plant-level operations.

Figure 4-81 shows that much of the time units at Blue Mesa actually operated at or very near to maximum output levels during the peak hours. The CHEERS model has a very similar

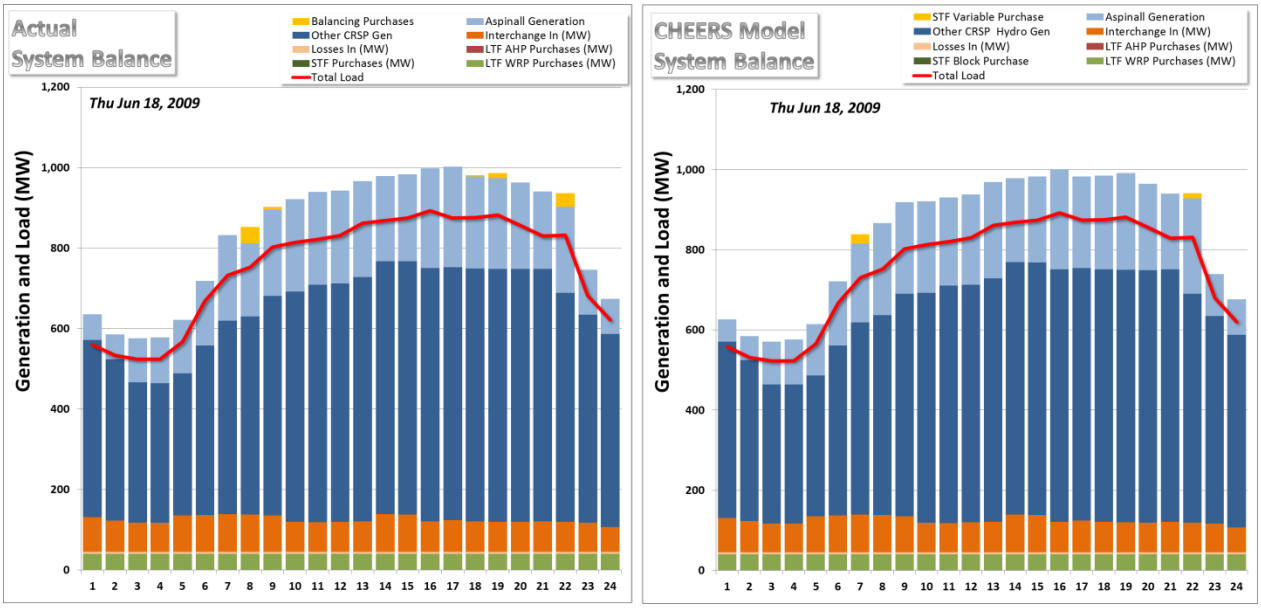

FIGURE 4-78 Actual and CHEERS Real-Time CRSP System Energy Balances (June 18, 2009)

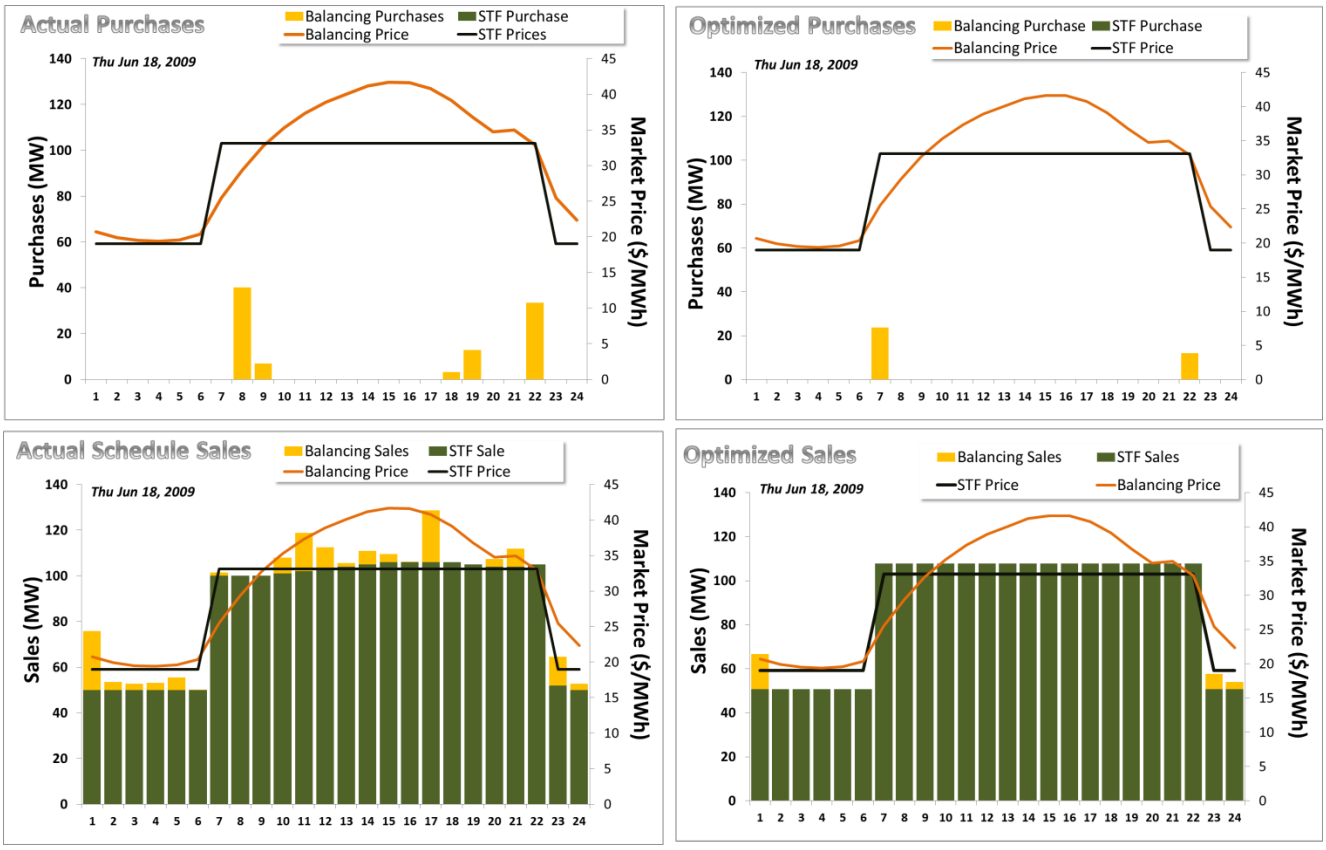

FIGURE 4-79 Actual CRSP and CHEERS Real-Time Power Market Transactions (June 18, 2009)

TABLE 4-21 Comparison of Actual Purchase and Sales Quantities with CHEERS Real-Time Results (June 18, 2009)

	Actual Transactions (MWh)			CHEERS Transactions (MWh)		
Market	Purchase	Sell	Net Purchase	Purchase	Sell	Net Purchase
Balancing	158	197	–40	36	27	9
STF	0	2,059	–2,059	0	2,131	–2,131
Total	158	2,256	–2,099	36	2,158	–2,122

TABLE 4-22 Comparison of Actual Purchase Costs and Sales Revenues with CHEERS Real-Time Results (June 18, 2009)

	Costs Based on Actual Energy Transaction and ICE Prices ($)			CHEERS Transaction Costs ($)		
Market	Purchase	Sell	Net Cost	Purchase	Sell	Net Cost
Balancing	3,877	6,416	–2,539	1,003	589	415
STF	0	61,142	–61,142	0	64,880	–64,880
Total	3,877	67,557	–63,681	1,003	65,468	–64,465

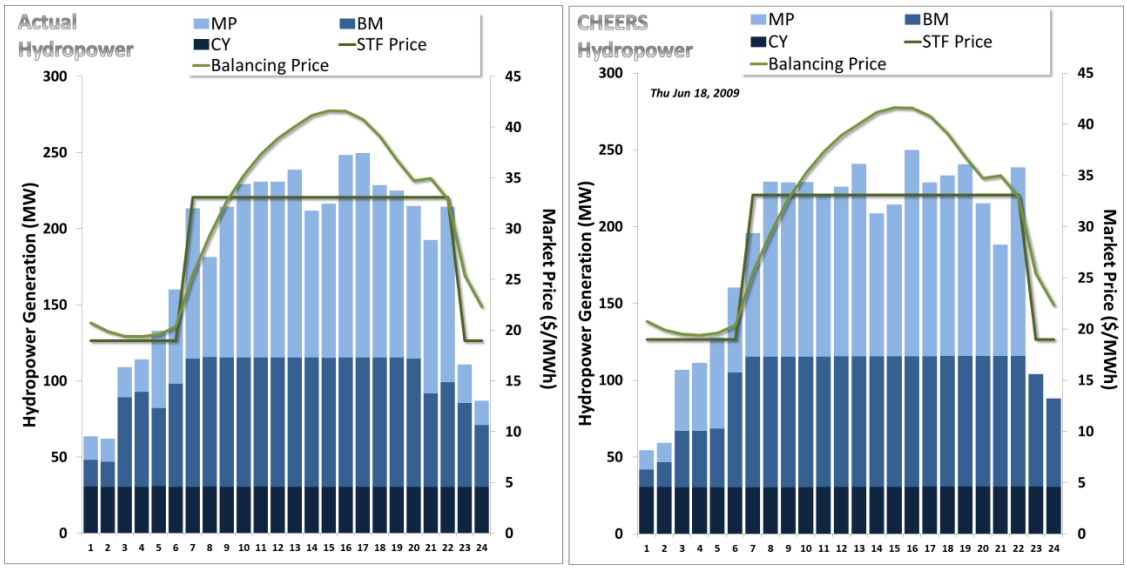

FIGURE 4-80 Comparison of Day-Ahead Scheduled and Actual Apsinall Plant-Level Generation with CHEERS Real-Time Results (June 18, 2009)

pattern. However, during the off-peak period, CHEERS operated just one unit near peak efficiency between hours 3 and 5 while actual operations operated two units at a lower efficiency point. This is reflected in Table 4-23, which shows CHEERS generated slightly more electricity (i.e., about 0.5%) over the day than what actually occurred.

It should also be noted that the number of start/stop cycles and ramping are about the same for both actual operations and CHEERS real-time model results. In general, when water releases are relatively high and units operate at or near maximum levels most of the time, the model has little room to improve operations relative to actual operations.

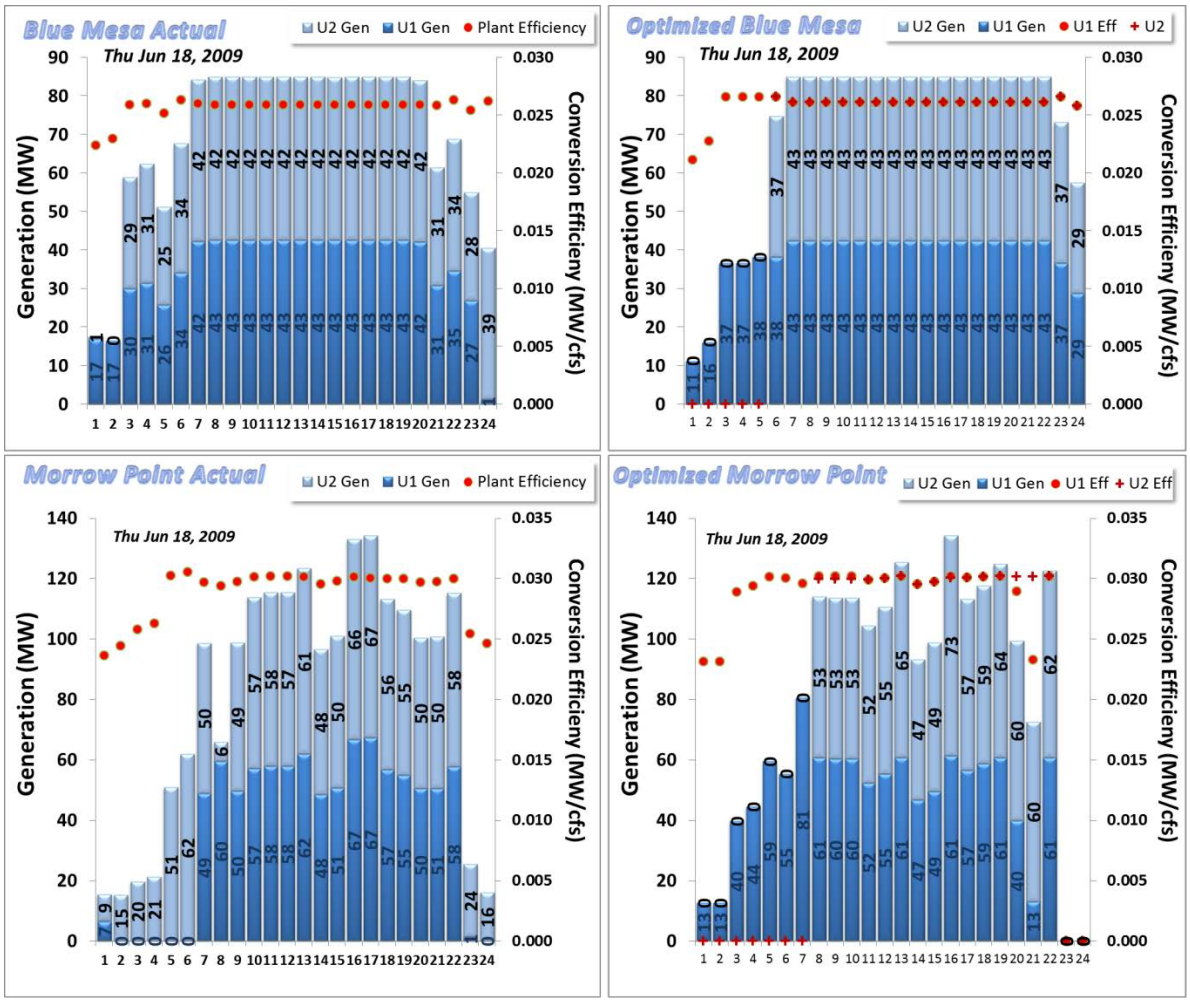

FIGURE 4-81 Comparison of Actual Apsinall Unit-Level Generation with CHEERS Real-Time Results (June 18, 2009)

TABLE 4-23 Comparison of Actual Generation with CHEERS Real-Time Results (June 18, 2009)

	Actual Generation (MWh)			CHEERS Generation (MWh)		
Unit	BM	MP	CY	BM	MP	CY
Unit 1	845	1,057	731	922	1,101	734
Unit 2	844	904		782	862	
Total	1,689	1,961	731	1,704	1,963	734

Real-Time Deviations from Day-Ahead Schedules. In the two cases discussed above, actual real-time operations were very similar to day-ahead schedules. However, sometimes events occur that make it necessary for operations to deviate from day-ahead schedules. This section presents a situation that occurred on Monday, May 4, 2009, that required real-time schedulers to adjust both generation and balancing market purchases. At the time day-ahead schedules were formulated, it was anticipated that generation from Aspinall power plants would be substantially higher than the levels that actually occurred. This most likely was due to an inaccurate inflow forecast. The top left graph in Figure 4-82 shows day-ahead schedules that were made under the assumption that Aspinall generation levels would be relatively high on May 4, 2009.

Mimicking actual operations, CHEERS day-ahead model results under this projection are shown in the upper right graph of the figure. Based on projected high Aspinall generation levels, CHEERS made small STF purchases.

Predicted levels of Aspinall generation did not come to fruition, leaving the system short of energy. Real-time schedulers needed to buy energy on the balancing market to fill the shortfall. The lower half of Figure 4-82 shows scheduling changes that were actually made (lower left graph) and those made by CHEERS (lower right graph). Note that it was too late to make additional STF purchases in real time, so the ones previously made in the day-ahead market were fixed in both actual operations and in the CHEERS simulated reality. Whereas actual purchases were made in small amounts over several hours, the CHEERS model concentrated purchases over a fewer number of off-peak hours when prices were lower.

Tables 4-24 and 4-25 summarize daily energy transaction quantities and dollar amounts, respectively. Since simulated Aspinall units are operated more efficiently, net purchases are lower than actual amounts. Because CHEERS purchases are less and made primarily in off-peak hours, total net purchase costs are about 6.0% less than actual net expenditures.

Actual and CHEERS results for Aspinall unit-level operations are shown in Figure 4-83. At Blue Mesa, CHEERS operates a single unit at levels that are at or very near to the most efficient generation point while actual operations had two units generating power at lower efficiency points. In addition, each set of units had two starts and stops, while CHEERS results have only one start and stop each.

CHEERS operations at Morrow Point are also much more efficient; they generate more power with the same amount of water. The model shows near peak efficiency operations during a smaller number of hours. For several hours, actual operations were at levels that are significantly lower than the turbine peak efficiency point. Total daily generation levels provided in Table 4-26 show that Aspinall production levels may have been up to 4.0% higher if CHEERS suggested hourly generation profiles were implemented.

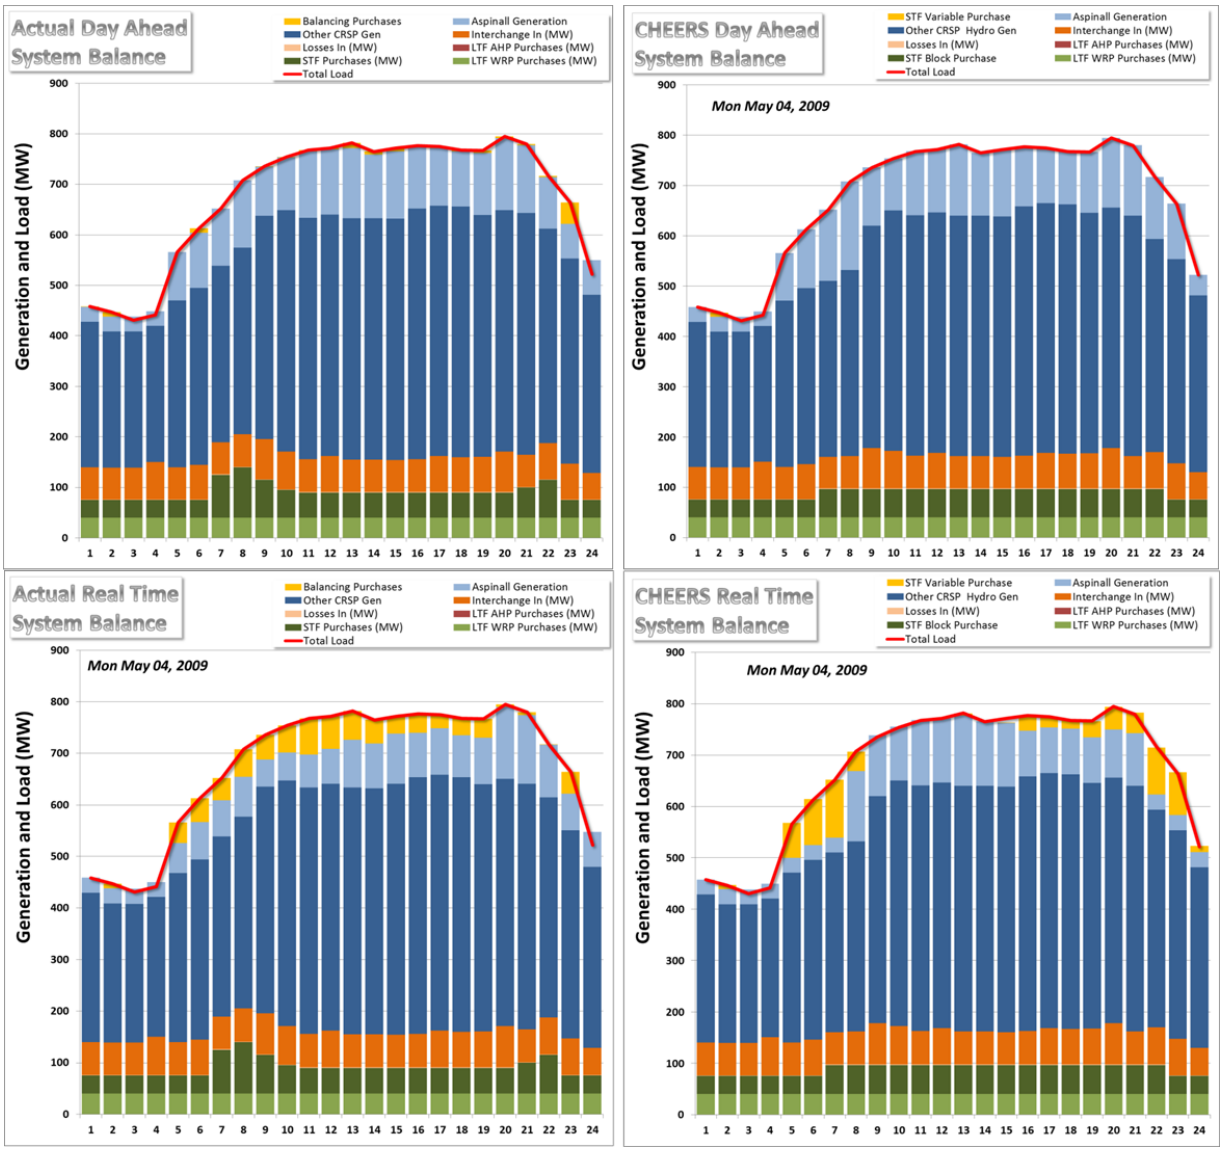

FIGURE 4-82 May 4, 2009, Actual and CHEERS CRSP System Energy Balances under Prescheuled Hydropower Projections (top graphs) and Actual Real-Time Conditions (bottom graphs)

The CHEERS model also suggests that operators could start and stops units less frequently. On May 4, 2009, Aspinall units had five start/stop cycles, while CHEERS real-time model results have only two start/stop cycles. Assuming each start and stop costs $200, CHEERS would save $1,200 in on/off cycling costs. CHEERS also has about 14.8% less ramping.

Monthly Summary. Daily CHEERS model optimizations runs were performed for the CRSP system over the 6-month retrospective study period. The previous sections described detailed model results for a few of those model runs to illustrate how CHEERS compares to actual operations. Although each daily situation is somewhat unique, presenting each day in this

TABLE 4-24 Comparison of Actual Purchase and Sales Quantities with CHEERS Real-Time Results (May 4, 2009)

	Actual Transactions (MWh)			CHEERS Transactions (MWh)		
Market	Purchase	Sell	Net Purchase	Purchase	Sell	Net Purchase
Balancing	748	41	708	690	18	673
STF	1,230	0	1,230	1,192	0	1,192
Total	1,978	41	1,938	1,882	18	1,865

TABLE 4-25 Comparison of Actual Purchase Costs and Sales Revenues with CHEERS Real-Time Results (May 4, 2009)

	Costs Based on Actual Energy Transaction and ICE Prices ($)			CHEERS Transaction Costs ($)		
Market	Purchase	Sell	Net Cost	Purchase	Sell	Net Cost
Balancing	21,473	955	20,518	18,862	402	18,460
STF	31,551	0	31,551	30,507	0	30,507
Total	53,025	955	52,069	49,369	402	48,967

report would be repetitious and somewhat overwhelming for the reader. Therefore, in this section aggregate monthly CHEERS model results are presented.

As with any retrospective analysis, it is very difficult to model what the past would have been under an altered reality. In this case, the altered reality is how the system would have been scheduled and operated had the CHEERS model been available and applied to the CRSP system. While critically reviewing the retrospective methodology and analyzing much of the detailed modeled output using the CHEERS Results Report Writer, it was observed that some aspects of the modeling system may lead to an overestimate of the value of CHEERS, while others may underestimate its value.

One major aspect of this CRSP study that may lead to an underestimate of the model's value is that it has very rigidly defined rules. For each daily run, CHEERS consistently applies the same set of rules and objectives that are based on direction from CRSP EMMO staff. However, the historical data reveal that schedulers' actions display some flexibility as schedulers adapt to prevailing conditions that occur when making scheduling decisions. To some degree this

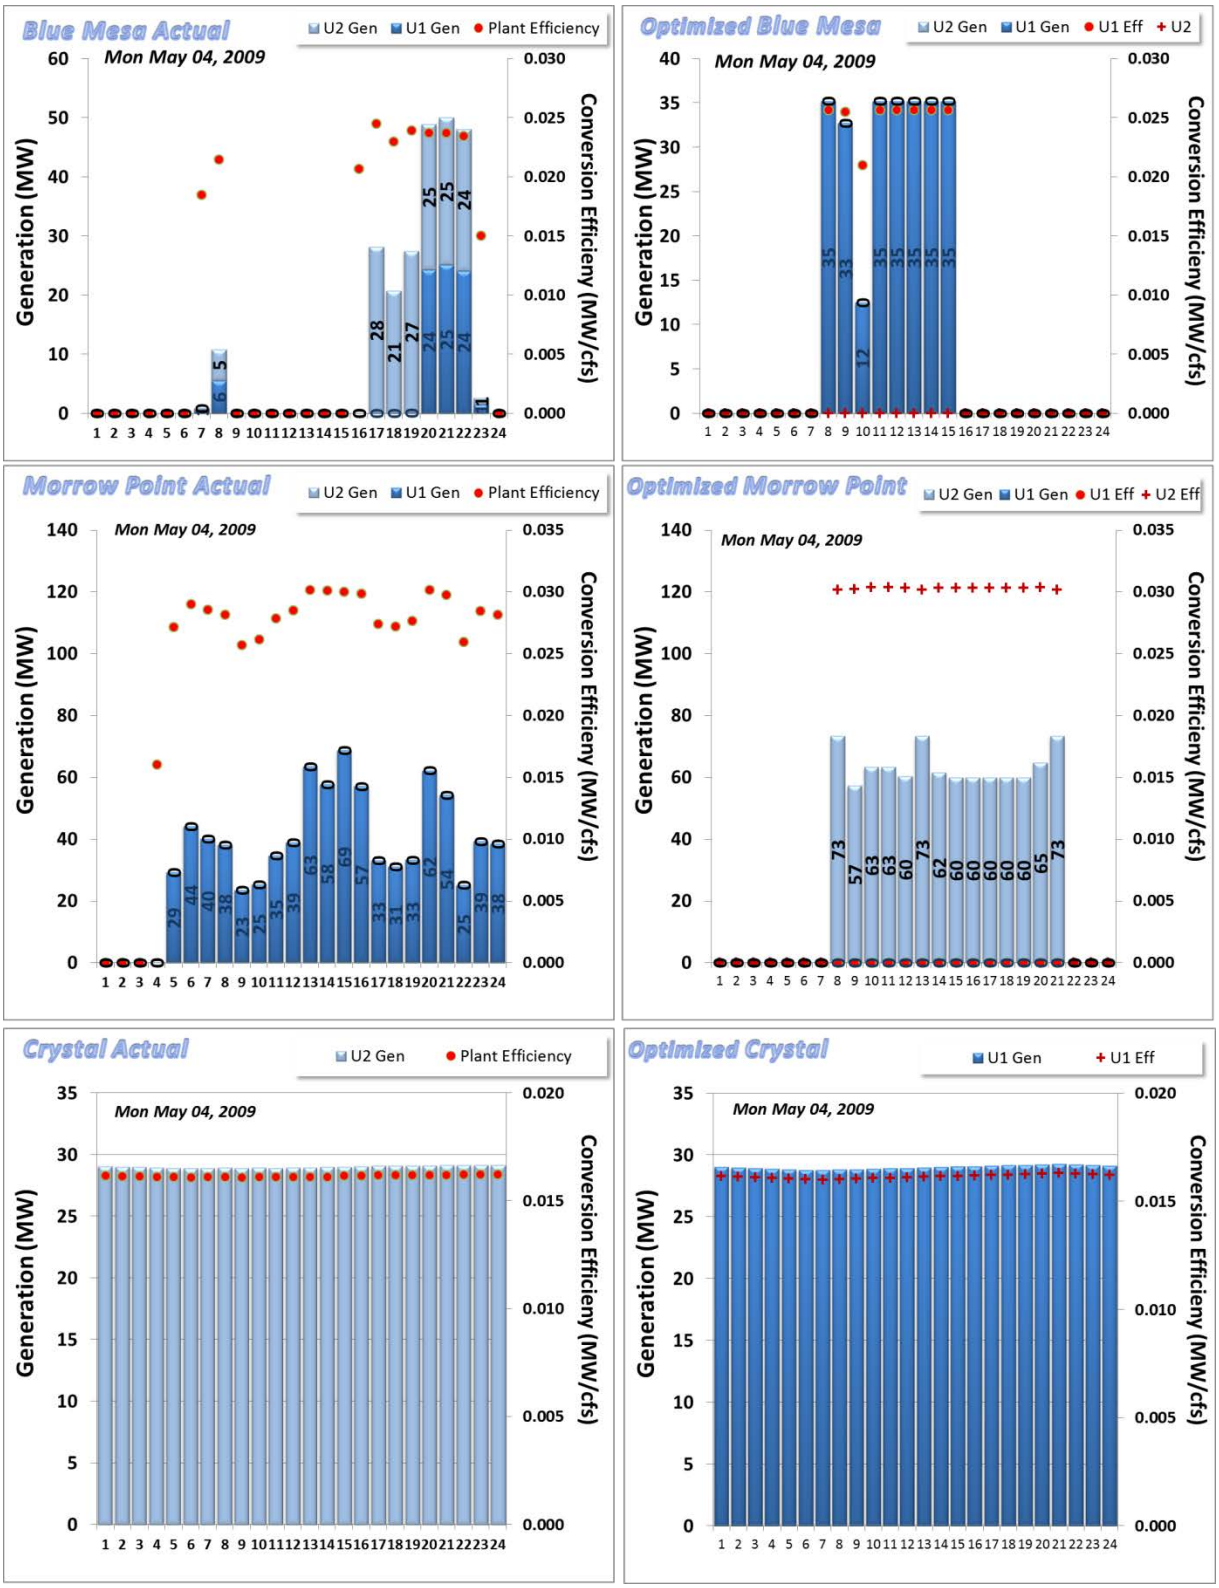

FIGURE 4-83 Comparison of Actual Apsinall Unit-Level Generation with CHEERS Results (May 4, 2009)

TABLE 4-26 Comparison of Actual Generation with CHEERS Results (May 4, 2009)

	Actual Generation (MWh)			CHEERS Generation (MWh)		
Unit	BM	MP	CY	BM	MP	CY
Unit 1	81	0	697	256	0	697
Unit 2	157	837	NA[a]	0	890	NA
Total	238	837	697	256	890	697

[a] "NA" = "not applicable."

flexibility allows for better utilization of CRSP resources. Under some circumstances, CHEERS' rigidly enforced rules tend to reduce estimates of the model's economic and efficiency gains. Due primarily to the model's rigidly enforced rules there are a few rare cases where CHEERS underperformed when compared to actual operations. Some of these rules include the following:

1. CHEERS always follows a strict 16-hour on-peak and 8-hour off-peak day-ahead STF block purchase pattern. Whereas this is generally the case in actual operations, there are numerous cases in which the actual STF on-peak and off-peak purchase pattern is not flat, but patterned, presumably, to meet schedulers' needs. One such case was described above for the May 4, 2009, run.

2. There are also numerous occasions where it appears that STF purchases were made in excess of total energy demands. This enables schedulers to sell the excess in the real-time market.

3. CHEERS only makes operational adjustments to the Aspinall Cascade holding model input values for all other CRSP hydropower plants to day-ahead and actual operating levels. In actuality, operations schedulers make adjustments within operational and environmental limitations to all CRSP plant schedules to best meet system objectives. Allowing CHEERS to modify these operations to some degree would allow it to find better systemic solutions.

4. Due to inflow/side flow forecast errors, operators at times violate reservoir elevation constraints. As discussed in more detail in Appendix K, reservoir operations most frequently exceed limits at the Crystal Reservoir. To the extent possible, CHEERS either eliminates or reduces these errors below levels that actually occurred.

On the other hand, there are some aspects of the retrospective modeling process that may lead to an overestimate of the potential gains that may be realized by CHEERS. The paragraphs below outline some modeling challenges and the rational for the methods employed.

In real-time operations CHEERS has more "foresight" in terms of inflows/side flows into the system. Note that real-time runs are made using information that is available at midnight. The model assumes that future inflows/side flows for that day are known with certainty. In reality, however, this is not the case. On the other hand, it is very difficult, if not impossible, to measure inflows/side flows even when actual retrospective observations are available. Inflows/side flows are not measured, but computed by a mass balance equation using observed hourly reservoir elevation levels and water releases. Both of these measurements are subject to significant error, and when computing inflows on an hourly basis, estimates often fluctuate unrealistically from large positive values to large negative values between consecutive hours. These unrealistic fluctuations are largely attributed to reservoir elevation measurement error. A very small elevation measurement error can lead to a large inflow/side flow error. However, over a longer time period, such as 24 hours, it appears that these errors typically tend to cancel each other and the results appear to be reasonable. However, even daily average inflows occasionally appear to be suspect.

Due to the difficulties listed above, the modeling methodology uses a single average daily inflow/side flow value for all hours in the day. The use of daily averages is consistent with the current day-ahead business process in which Reclamation sends the CRSP EMMO daily water values for the previous day. These data are input into the "water" sheet of the CRSP scheduling workbook and serve as projections for the following day. Under most conditions, this simplification of reality has little impact on the results since reservoir constraints are typically not binding in the model's solution. However, inflow/side flow assumptions during the rapid snowmelt period are sometimes binding at the Crystal Reservoir. The perfect real-time inflow foresight assumption from midnight onward may potentially lead to an overestimated value of CHEERS in these situations. On the other hand, this condition occurs when generation levels are very high. The modeled and actual operations for Morrow Point and Crystal power plants are very similar since the only operational solution is to release water at or near turbine capacity during most, if not all, hours of the day.

CHEERS also has more foresight than schedulers in terms of predicting market prices. However, operational patterns are not driven primarily by absolute market prices, but rather by a relative price pattern. For example, if the market energy price is always $1,000/MWh, power plant operations would tend to be flat around the clock with turbines operating at the highest possible efficiency points. However, if prices are $10/MWh at night and $35/MWh during the day, operations are maximized by generating to the greatest extent during on-peak hours at points that might exceed the peak efficiency level. Although the absolute values of prices are sometimes difficult to predict, the price pattern is fairly stable and consistent during a specific situation. The pattern can be predicted with a relatively high degree of confidence. That is, prices are relatively low during the night and more expensive during the day. Prices follow a single-humped pattern during the warm months of the year and a double-humped pattern with a midday price dip during other times of the year.

Instead of making a single real-time run of CHEERS each day, the model can be run each hour when mimicking real-time operations. In such an arrangement each hour's modeled "historical" information would be updated, and initial conditions would change as simulated time

unfolds. This would be a very time-consuming exercise and would preclude parallel model runs. Since it was judged to change the results very little in most situations, this methodology was not utilized.

A monthly comparison of CHEERS real-time model results and actual operations are summarized in the following tables. Table 4-27 shows that CHEERS net costs over the 6-month retrospective study period are about a half a million dollars less than the amount computed for actual operations using ICE derived market prices. Monthly variations are a function of numerous factors including market price profiles, daily water release volumes, and the skill and objectives of schedulers and plant operators.

Although CHEERS results have lower net costs on a monthly basis, on a few days the CHEERS model produced real-time operations that had higher net costs, underperforming actual operations. Several factors that produced negative results on these few days are listed above. It indicates that CHEERS model inputs and objectives should be tailored to specific conditions. For example, when real-time prices are consistently higher than day-ahead prices and the system has excess energy, it may be prudent for CHEERS to withhold some energy sales for the real-time market instead of selling all of the excess on the day-ahead market.

Monthly monetary savings are partially attributed to more efficiently committing and utilizing Aspinall power plant turbines and generators. Table 4-28 shows that about 1% more energy may have been generated from Aspinall Cascade power resources if the CHEERS model suggestions had been implemented. Note that the model releases exactly the same amount of water on a daily basis as historically observed.

As described previously, CHEERS uses estimated turbine efficiency curves based on historical relationships between turbine flows and power generations. Given all of the constraints input into the model, CHEERS, to the extent possible, maximizes the value of hydropower production, trading off the economics of generating at peak power at lower turbine efficiency when prices are high and backing down power generation to a lower output level at the peak efficiency point. Over the 6-month study period, Blue Mesa has an increased efficiency of about

TABLE 4-27 Comparison of Monthly Net Costs

Month	Net Costs Based on Actual Energy Transaction and ICE Prices ($1000)	CHEERS Net Costs ($1000)	Net Savings ($1000)
April	2,170	2,107	63
May	–788	–899	111
June	–1,529	–1,579	50
July	–1,856	–2,018	161
August	–1,231	–1,305	74
September	642	600	42
Total	–2,592	–3,092	501

TABLE 4-28 Monthly Comparison of Actual Generation with CHEERS Results

	Actual Generation (GWh)				CHEERS Real-Time Generation (GWh)			
Month	BM	MP	CY	Total	BM	MP	CY	Total
April	17.1	23.9	15.8	56.8	17.5	24.5	15.8	57.8
May	32.6	54.8	22.7	110.1	33.8	55.6	22.5	111.8
June	53.9	66.2	21.9	142.0	54.2	66.4	21.9	142.5
July	45.1	53.0	22.4	120.5	45.4	53.8	21.9	121.2
August	39.3	45.9	22.7	107.9	40.5	46.6	22.2	109.3
September	27.7	34.5	19.7	82.0	28.1	35.1	19.7	82.9
Total	215.8	278.4	125.2	619.3	219.5	282.0	124.0	625.5

1.6% and Morrow Point's efficiency increases by 1.3%. Ironically, the efficiency at Crystal decreases by 0.9%. In the standalone retrospective analysis, CHEERS has little control over the operation of Crystal. That is, water is released at a constant rate through the power plant's single turbine as prescribed by historical water release volumes. CHEERS' lower generation may stem from to the water-to-power conversion function that the model uses to estimate generation as a function of water release and reservoir elevation. Perhaps the model input efficiency function should be slightly higher. However, it is difficult to obtain the exact efficiency curve given the somewhat erratic behavior of observed conversion efficiencies in the historical data and possible water flow measurement errors. Regarding the point, CHEERS efficiency and economic gains may potentially be slightly higher the ones reported in this section.

Two other metrics that are used to measure the potential benefits associated with the use of CHEERS are Aspinall unit start/stop cycles and unit generation ramping. Table 4-29 shows

TABLE 4-29 Monthly Comparison of Actual Start/Stop Cycles CHEERS Results

	Actual Number of Start/Stop Cycles				CHEERS Real Time Number of Start/Stop Cycles			
Month	BM	MP	CY	Total	BM	MP	CY	Total
April	72	76	NA[a]	148	41	36	NA	77
May	53	41	NA	94	43	42	NA	84
June	18	38	NA	56	19	48	NA	67
July	40	50	NA	90	46	63	NA	109
August	63	68	NA	131	63	70	NA	133
September	62	62	NA	124	52	60	NA	112
Total	307	334	NA	641	263	318	NA	581

[a] "NA" = "not applicable."

that over the 6-month study period the number of start/stop cycles is about 9.4% less under the CHEERS operational regime compared to actual operations. Much of this reduction occurs at the Blue Mesa Power Plant in the month of April, which is characterized by a two-humped load and price pattern. During early May, this pattern begins to transform into the summer single hump pattern. On the other hand, CHEERS displays a slightly higher number of start/stop cycles during peak load and priced summer months. Monthly differences between actual and modeled ramping (Table 4-30) follow a similar pattern, with CHEERS displaying significantly lower ramping during April.

CHEERS model results are sensitive to model inputs and assumptions. The results presented above are based on a limit of two start/stop cycles per unit each day at an assumed cost of $200 per unit start and stop. Two additional CHEERS model runs were conducted with different start/stop assumptions. The first sensitivity run also assumes a limit of two starts/stops per day at each unit. However, the cost of each unit start and stop was reduced to $50. A second sensitivity run is identical to the previous one except the unit start/stop limit is reduce to one per day. Table 4-31 shows that financial benefits and generation levels would be higher if the unit start/stop cost assumption is lowered. However, this increases both the number of start/stop cycles and ramping above actual operations. When the number of start/stop cycles is reduced to one per day per unit, CHEERS economic savings are slightly lower.

As discussed in previous sections, CHEERS MILP results for water release rates through various channels at a reservoir can be input into the Nonlinear Post-processor Routine. Appendix K discusses the application of the routine to one of the sensitivity model runs. It shows that the nonlinear approximation is very similar to results produced by nonlinear set water and power equations.

Future CHEERS Model Role in the CRSP EMMO Office. An earlier version of the CHEERS model was installed and demonstrated at CRSP EMMO. In the late summer of 2013, the most recent version of the model will replace the older one and automate connections between CHEERS and EMMO data sources. In addition, during this time Argonne and EMMO staff will jointly perform day-ahead and real-time model runs and compare the results to decisions that will be made using current business practices. The intent of this exercise is to train EMMO staff in the use of CHEERS. The plan is to eventually incorporate the model into EMMO daily business practices to support scheduling staff decision making and to improve operations. Argonne, EMMO schedulers, and Reclamation staff will also meet to discuss coordination strategies among CHEERS model results, plant scheduling, and unit level operations.

Through the retrospective analysis, valuable insights were gained into both historical operations and the behavior of the CHEERS model. One potential role of the model at the EMMO is to provide schedulers with several potential operational and purchasing options instead of a single schedule. Different solutions would be driven by variations in model inputs such as those used in the retrospective sensitivity analysis and tweaks in operational guidelines.

TABLE 4-30 Monthly Comparison of Actual Ramping with CHEERS Results

	Actual Ramping (MWh)				CHEERS Real-Time Ramping (MWh)			
Month	BM	MP	CY	Total	BM	MP	CY	Total
April	7,083	11,233	NA[a]	18,316	4,560	8,813	NA	13,373
May	5,352	9,010	NA	14,362	4,187	11,603	NA	15,790
June	3,626	12,311	NA	15,937	3,507	13,640	NA	17,147
July	5,856	12,945	NA	18,802	6,045	12,116	NA	18,161
August	7,757	13,017	NA	20,774	6,950	13,098	NA	20,048
September	5,690	10,337	NA	16,027	5,232	9,952	NA	15,184
Total	35,364	68,853	NA	104,217	30,480	69,222	NA	99,701

[a] "NA" = "not applicable."

TABLE 4-31 CHEERS Stand-Alone Mode Run Sensitivity Analysis

		2 Cycles, $200		2 Cycles, $50		1 Cycle, $50	
Analysis Metric	Actual	CHEERS	Difference from Actual	CHEERS	Difference from Actual	CHEERS	Difference from Actual
System Cost ($1,000)	–2,592	–3,092	501	–3,129	537	–3,085	494
Generation (GWh)	619.3	625.5	6	626.3	7	625.5	6
No. of Start/Stop Cycles	641	581	–60	792	151	618	–24
Power Ramping (MWh)	104,217	99,701	–4,516	112,410	8,192	99,849	–4,369

For example, one run may minimize balancing market transactions and another run would allow for more balancing purchases and sales to occur. Schedulers could examine and compare results in terms of not only net costs, but also other metrics such as unit starts and stops, ramping, and risks associated with day-ahead STF and balancing transactions. Schedulers would then implement and perhaps refine the solution that is deemed best given the current EMMO marketing and operational situation. Since CHEERS is designed to take advantage of a computer's multiple processing capabilities, several model perturbations can be run in parallel with little impact on total elapsed run time.

4.4.1.5 CHEERS Retrospective Analysis in an Integrated WUOT Framework

This section examines the advantages of applying CHEERS as an integrated component in the WUOT framework. In the integrated mode of operation, CHEERS is run in much the same way as described above for the stand-alone retrospective analysis. However, there are a few key differences that involve using results from other WUOT tools instead of relying solely on historical information. For use in day-ahead scheduling, information produced by the WUOT HydroSCOPE tool is used by CHEERS for forecasts of daily water release volumes for Aspinall reservoirs. The WUOT EHFS tool provides inflow/side flow forecasts. For real-time optimization runs, HydroSCOPE provides updated daily water release volumes. Because HydroSCOPE and the IRF tool are intertwined, the HydroSCOPE daily water volumes keep CHEERS short-term operations in line with long-term goals involving not only water deliveries, but also power and environmental objectives. These volumes typically depart from actual operations because HydroSCOPE attempts to simultaneously enhance the environment and increase the value of hydropower above historic levels.

The configuration of data flows into CHEERS along with the sequence of day-ahead and real-time runs used for the integrated retrospective analysis is shown in Figure 4-84. This configuration is similar to the one used for the stand-alone mode, but there are several differences, as described below.

1. The integrated model runs create an "alternate reality" that is recorded and stored as simulated time unfolds.

2. In accordance with Montrose office operational goals and guidelines, CHEERS maximizes day-ahead STF transaction quantities such that balancing purchases and sales are minimal. However, there are occasions when the schedulers' actual operations appear to have deviated from this basic guidance. When CRSP resources are long (i.e., supply is greater than firm load obligations), it was observed that operators sometime sold less STF energy in the day-ahead market and sold more energy during high price hours in the real-time balancing market. This operational regime is not allowed by CHEERS and therefore the model gave a result which was economically inferior to the actual operation.

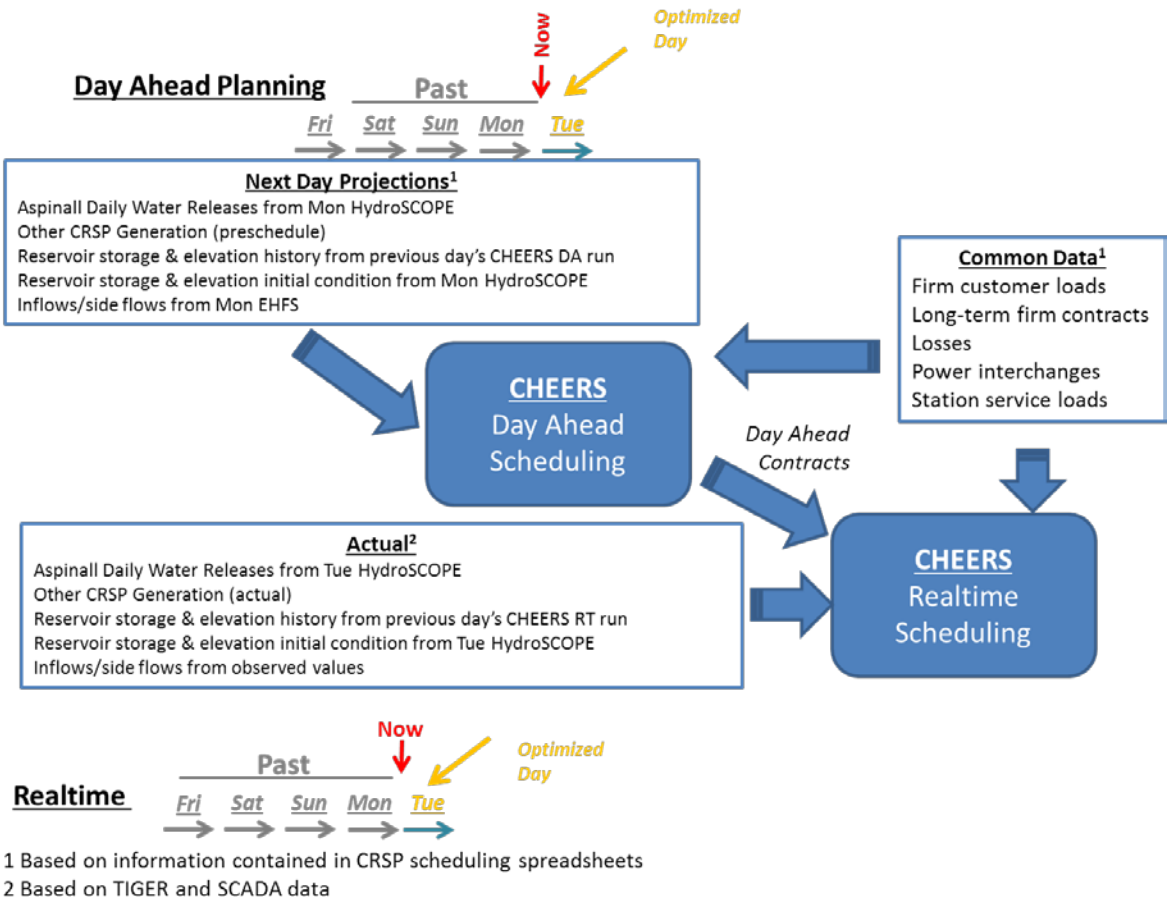

FIGURE 4-84 CHEERS Model Run Sequence and Data Flows for WUOT Integrated Retrospective Analyses (DA = day-ahead)

3. For day-ahead runs, daily inflow and side flow forecasts are based on EHFS results. For real-time runs, actual daily observed inflows/side flows are utilized since it is assumed that flows into the cascade are "natural" uncontrolled events.

4. Projected daily water release volumes for day-ahead runs are based on HydroSCOPE projections for the next day.

5. The day-ahead optimization run is assumed to be performed at midnight rather than at 10 AM. This is because in the integrated retrospective analysis the model does not have access to the actual day-ahead schedule. If the model were to be used by actual schedulers on a daily basis, then such data would be available and therefore the day-ahead optimization could be run at 10 AM or for any other hour of the day.

6. Initial reservoir conditions for each day's run are obtained from HydroSCOPE. The model reconciles the optimal reservoir release volume with the actual inflows and side flows as it steps through the simulated 6-month period.

7. Daily water release volumes for real-time runs are based on HydroSCOPE release volumes for the current day.

From a power resource scheduling perspective, the most significant difference between CHEERS stand-alone runs and WUOT integrated runs is that the total daily water volumes released from each reservoir do not track the actual historic operations. Figure 4-85 shows daily water release volumes compared to day-ahead, actual, and HydroSCOPE amounts. Note that expected day-ahead and actual water release volumes track closely, especially for the Crystal Reservoir. In contrast, HydroSCOPE water releases at times depart significantly from actual release volumes. Although HydroSCOPE daily water releases differ from actual, the total amount released from the cascade over the 6-month study period is very similar. Note that in Table 4-32, despite significant differences between actual and HydroSCOPE releases on a monthly basis, over the entire retrospective study period releases are within about 0.5% of the actual amount for each reservoir. Because HydroSCOPE used initial reservoir conditions on April 1 that exactly matched historic values, HydroSCOPE ending reservoir water storage volumes and elevations at the end of September are very similar.

Another key difference between CHEERS stand-alone and integrated model runs is the use of inflow/side flow projections made by the EHFS tool rather than the use of persistence forecasting. Figure 4-86 compares actual inflow/side flows against a day-ahead persistence forecast and against the EHFS forecast. Note that the persistence forecast is simply the actual flow shifted 2 days into the future.

It should also be noted that in Figure 4-86, in the latter part of September, actual inflows at Morrow Point are erratic. These values are based on water mass balance equations as previously described and were used in the stand-alone analysis. These abrupt side flow fluctuations may not have actually occurred, but instead may be an artifact of elevation measurement errors at the Morrow Point Reservoir. As previously discussed, small reservoir elevation measurement errors can lead to large errors in hourly inflows. Although suspect, there is no definitive proof in the data series that the recorded elevations in the historical dataset are incorrect and they were therefore retained in the modeling process.

Whereas a meaningful comparison of operations can be made between the CHEERS model and the actual historic operation when running in the standalone mode on any given day, such a comparison would be inconsistent for integrated mode runs. Note that as WUOT steps through time it creates an alternate path in which the day's starting and ending reservoir conditions and daily water release volumes can differ markedly from the historic record. Therefore, comparisons are made over the entire study period because HydroSCOPE and actual total water releases over the entire study period are similar (i.e., within 0.5%).

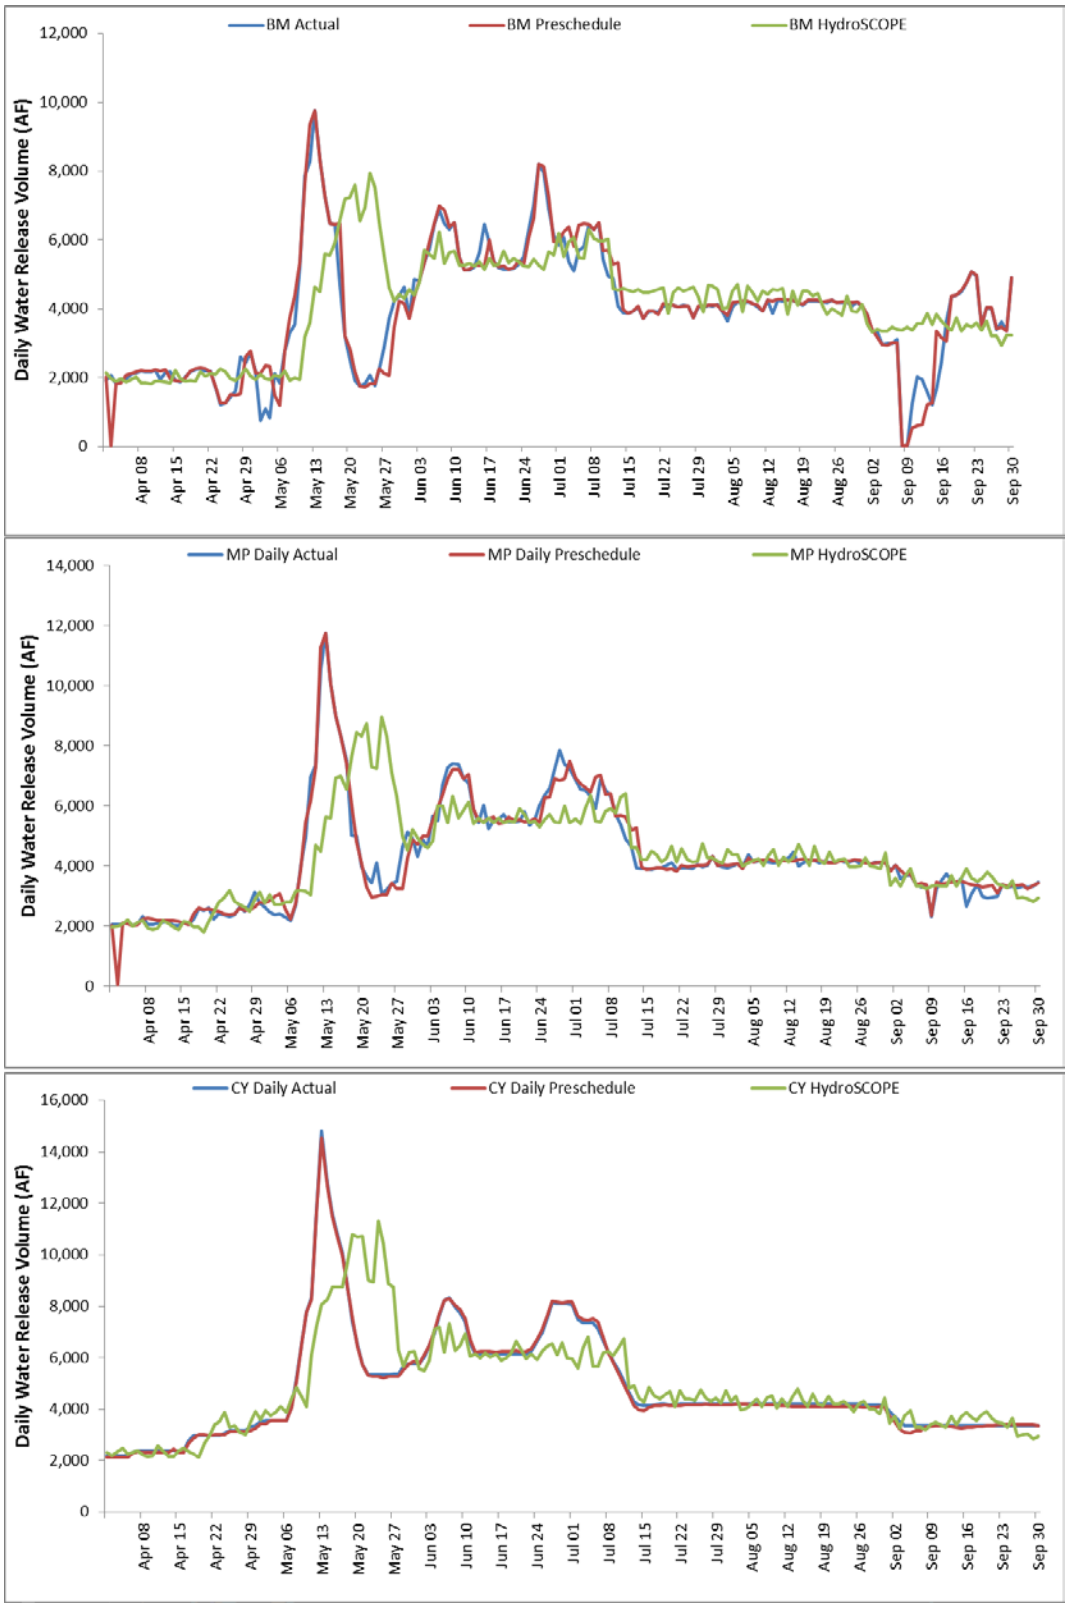

FIGURE 4-85 Comparison of Actual Daily Water Release with Day-Ahead Prescheduled Level and Optimized HydroSCOPE Amounts

TABLE 4-32 Comparision of Actual and HydroSCOPE Monthly Water Releases

	Actual Water Releases (AF)				HydroSCOPE/CHEERS Water Releases (AF)			
Month	BM	MP	CY	Total	BM	MP	CY	Total
April	61,216	69,564	80,164	210,944	56,586	67,975	80,207	204,768
May	120,285	155,220	207,949	483,454	134,925	167,344	220,393	522,662
June	175,513	184,467	206,664	566,645	156,264	165,634	187,979	509,877
July	144,319	148,672	158,420	451,411	149,992	150,579	157,882	458,453
August	128,422	128,815	130,031	387,269	129,196	129,869	131,599	390,664
September	93,535	100,453	102,187	296,175	101,139	101,373	103,394	305,906
Total	723,290	787,192	885,415	2,395,898	728,103	782,774	881,454	2,392,331

Table 4-33 shows a comparison of several power system metrics for actual operations, along with CHEERS results for both stand-alone and WUOT integrated model runs. Both runs assume a limit of one unit start/stop per day at a cost of $50 per start and stop. Compared to actual operations the CHEERS integrated run has a savings of about $662,000. The equivalent analysis in CHEERS stand-alone mode is $492,000. It should also be noted that the CHEERS integrated results have the highest economics of all stand-alone runs, despite slightly lower reservoir water releases and very restrictive start/stop cycle limits. The improved economics are primarily attributed to more favorable distribution of daily water release volumes over the study period. That is, more water is released when it has a higher value. In addition, a more accurate projection of inflows benefits schedulers because it allows them to better match day-ahead STF purchases with real-time scheduling needs.

4.4.1.6 CRSP Retrospective Summary

As with any retrospective analysis, it is very difficult to model what the past would have been under an altered reality and modeling improvements are almost always possible. Therefore, the absolute numbers that are produced in this study should be viewed in that context. However, results of this study suggest that there is sufficient indicative evidence that CHEERS may potentially improve CRSP operations. This involves a systemic approach of Aspinall operation that recognizes CRSP interactions and interdependencies over both time and space.

The comparison of actual operations with CHEERS model results in a stand-alone mode suggests the following possible areas of improved operations:

- Operating units at higher water-to-power conversion efficiency;
- Reduced number of unit starts and stops;
- Less ramping;

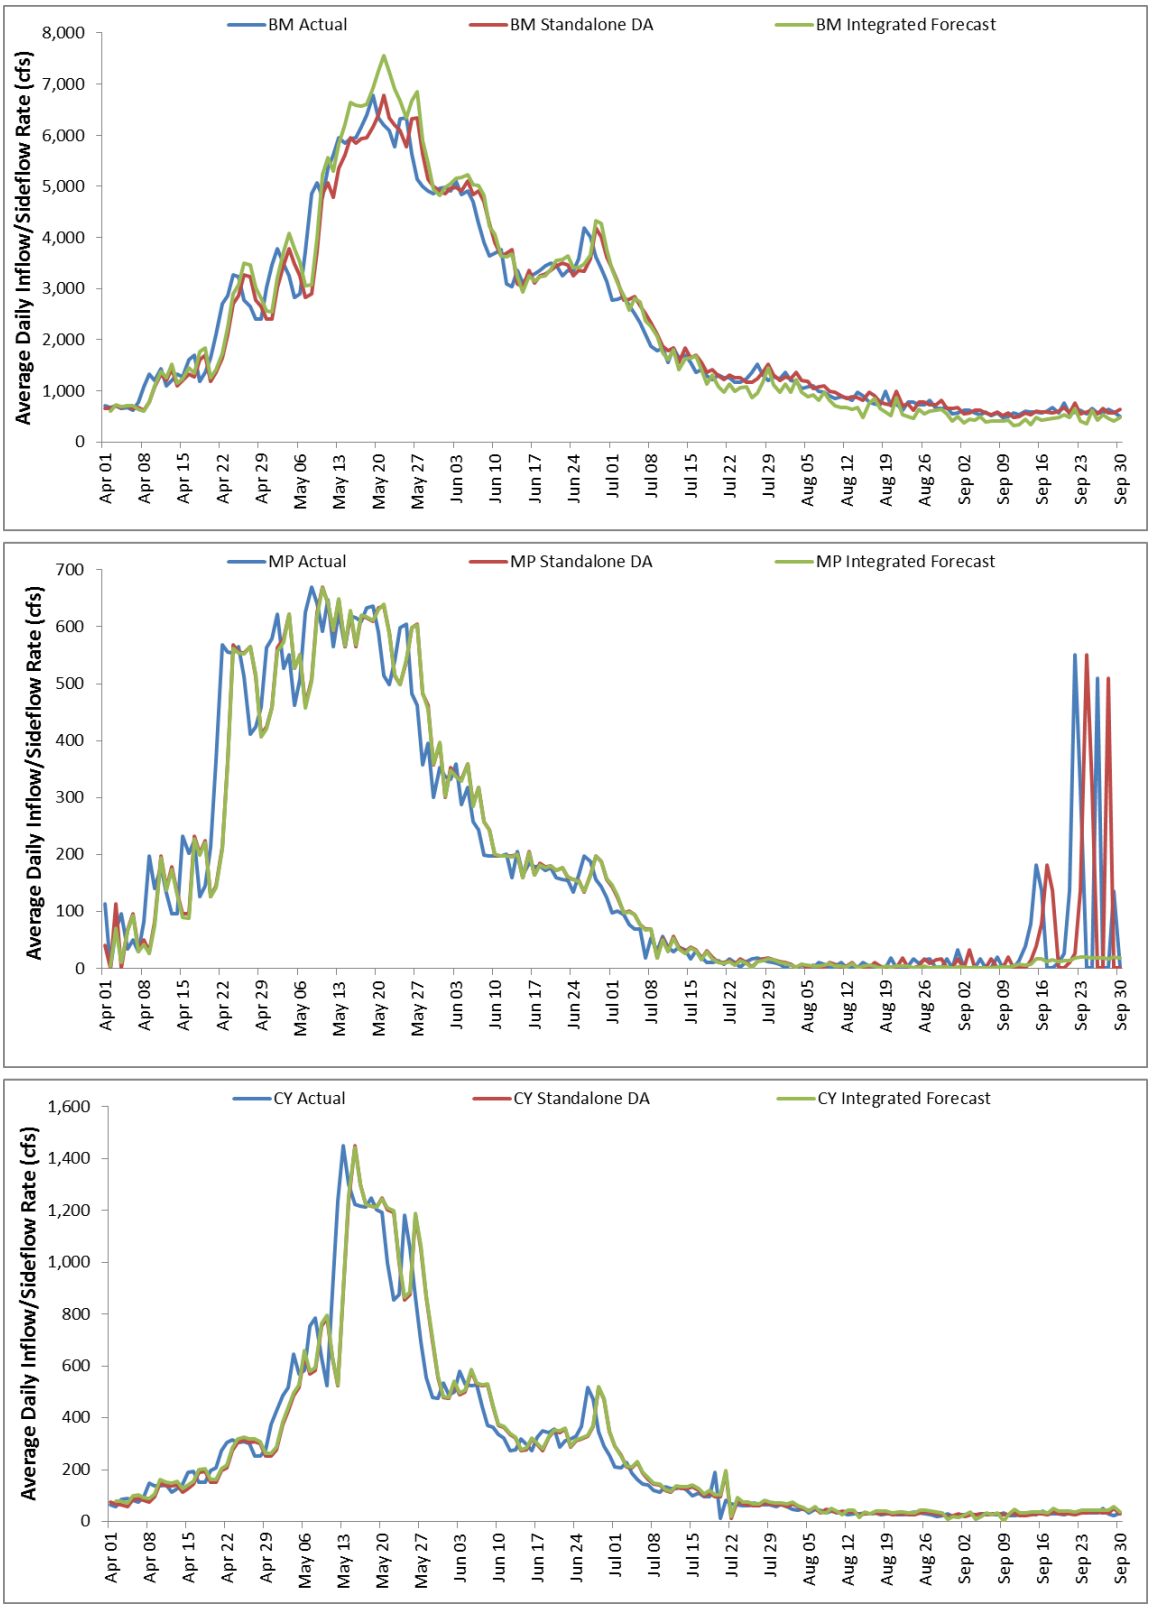

FIGURE 4-86 Comparison of Actual Inflow/Side Flows with Persistence and EHFS Forecasts

TABLE 4-33 Comparision of Power System Metric for Actual and CHEERS WUOT Integrated Model Results Assuming a One Unit Start/Stop Limit Per Day

Analysis Metric	Actual	Integrated WUOT CHEERS Run	Difference from Actual	CHEERS Standalone Run	Difference from Actual
System Cost ($1,000)	–2,592	–3,253	662	–3,085	494
Generation (MWh)	619.3	629.6	10	625.5	6
Number of Start/Stop Cycles	641	635	–7	618	–24
Power Ramping (MWh)	104,217	101,805	–2,412	99,849	–4,369

- Improved STF transactions that improve off- and on-peak block purchases and sales; and

- Reduction in balancing spot market transactions.

It was also shown that running CHEERS as an integrated component of WUOT may not only further improve power economic and operational efficiency, but also enhance the environment. The implementation of the entire WUOT package would require additional changes to current institutional arrangements and coordination between Reclamation and EMMO staff. If full implementation of the WUOT does not come to fruition, EMMO schedulers will still be able to take advantage of improved inflow forecast produced by the EHFS. Furthermore, as time permits and as more information is collected about EHFS inflow forecast accuracy, EMMO staff could potentially use the statistical method discussed in previous sections and in Appendix J to help manage operational risks and reduce reservoir elevation errors.

Argonne, EMMO schedulers, and Reclamation staff will continue to collaborate on CHEERS and WUOT CRSP related analyses. In the late summer of 2013, the most recent version of the model will replace the older one and automate connections between CHEERS and EMMO data sources. In addition, during this time Argonne and EMMO staff will jointly perform day-ahead and real-time model runs and compare the results to decisions that will be made using current business practices. The intent of this future exercise is to train EMMO staff in the use of CHEERS for future integration of CHEERS and potentially other WUOT components into Western EMMO and Reclamation business practices.

4.4.2 Oroville Complex

4.4.2.1 Hydroforecasting

The Upper Feather River Basin in California drains an 8,340 km^2 area upstream of Lake Oroville, California. The majority of the land in the basin is maintained by the United States Forest Service and is mostly unregulated until it reaches Lake Oroville. The climate is

Mediterranean, and most of the precipitation occurs during the cold season (November–March). The elevation ranges from 275 m to over 2900 m, with an average snow line at 1650 m; this is a mixed snow-/rain-dominated basin with a high daily flow variability due to a high sensitivity to daily temperature changes (Koczot et al. 2005). Lake Oroville is a major water storage area of the California State Water Project (CDWR 1998) and supports a large variety of water uses: flood control, irrigation water supply, hydropower (760-MW capacity), and water quality, among others. Snowmelt accounts for 40% of the long-term mean annual flow. Because of the mixed snowmelt/rain dominance, the accuracy of both initial conditions and climate forecasts are critical for medium-range and seasonal flow forecasts in the Feather River Basin (Kozcot et al. 2005).

The EHFS provides medium-range and seasonal forecasts at Oroville (Figure 4-87 and Table 4-34) to the WUOT's seasonal hydrosystem analysis, environmental performance, day-ahead scheduling, and real-time operation components. By conveying uncertainty, ensemble-based forecasts are well suited to risk-based water resource management. It is desirable that the ensemble range capture (or bracket) the potential range of flow over the forecast period.

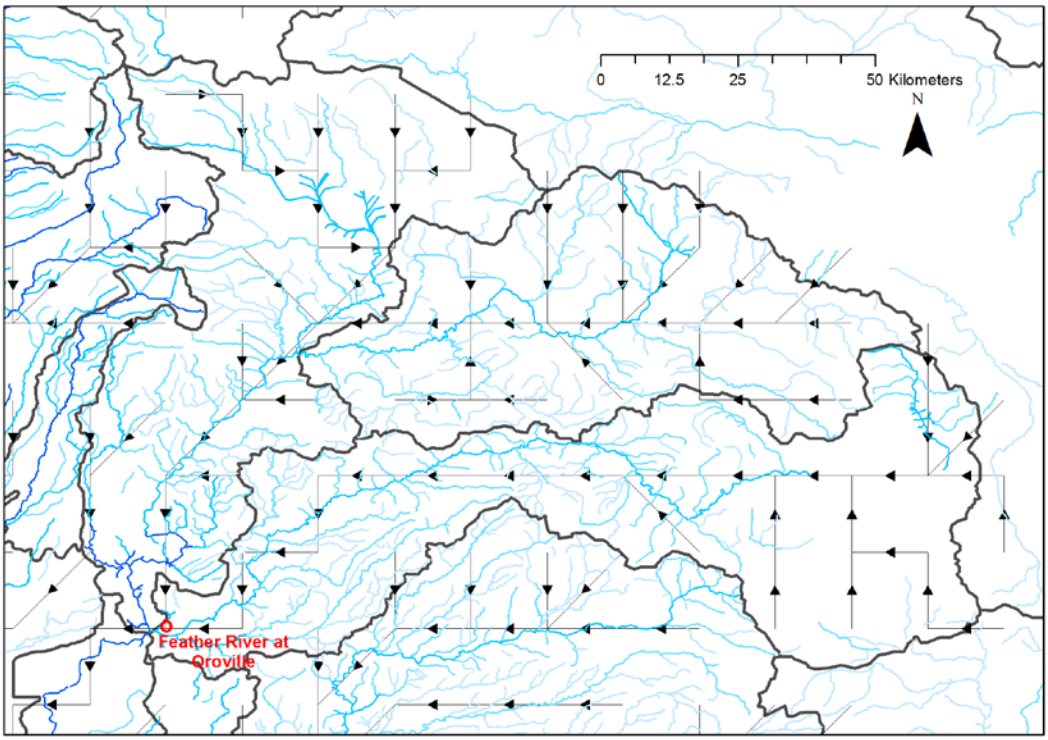

FIGURE 4-87 Feather River Basin, VIC Routing, and the Forecast Point at Oroville

TABLE 4-34 The Forecast Point Lookup is Manually Established to Provide a Relationship between the One-Eighth-Degree Routing Network, the Node in the Integrated Toolset, and Corresponding On-the-Ground Latitude/Longitude When Available

Forecast Point	Latitude	Longitude	Feature Description	EHFS Latitude	EHFS Longitude
OROVI	39.536266	–121.482576	Oroville Dam	39.5625	–121.4375

Calibration. Calibration of the EHFS VIC-routing model setup was performed over the Feather River Basin at the daily time step, at a one-eighth-degree spatial scale with respect to the observed naturalized daily flow obtained from the California Data Exchange Center. The calibration was performed using the Multi-Objective Complex Evolution of the University of Arizona (MOCOM-UA) method (Yapo et al. 1998), as applied in Voisin et al. (2011). Soil and routing parameters were both optimized at the daily time step. Four objective functions were used: the variance explained, the relative bias, the relative standard deviation difference, and the absolute value of annual mean volume error. The performance of the simulated flow is evaluated with respect to the naturalized daily flow at Oroville Dam (Table 4-34) (station code ORO; see http://cdec.water.ca.gov).

Table 4-35 presents standard measurements of the flow simulation performance for the calibration period (2000–2005) and the verification period (1990–2000) for the following components: correlation, NSE, and raw and relative RMSEs and MSEs. The performances for both periods (calibration and verification) are very similar, with slightly higher correlation and NSEs, but also RMSE during the validation period. This similarity in performance for those two periods is promising for stable performances for future simulations (i.e., in forecast mode).

Figure 4-88 presents the monthly mean flow climatology during both the calibration and the validation periods. It allows a visual check of the monthly performance. Overall, we can expect the summer flow to be slightly underestimated in forecast mode, while spring flow might be overestimated at times. Finally, Figure 4-89 shows daily flow hydrographs for the simulated flow and the naturalized daily flow at Oroville for water years 1994–2000. Overall, daily peak flows are underestimated in magnitude. The rising limb of the peak is timed correctly, but the peak might be about a day late. These issues are currently being explored.

Medium-Range Forecasts. The medium-range forecast provides distributed hourly meteorological forecasts over the basin and daily flow forecasts out to 13 days for 15 traces and an average trace for Oroville for 2011. As described previously, the GFS medium-range weather forecasts were downscaled to the grid of the hydrological model and calibrated to remove systematic biases and improve the ensemble forecast probabilistic skill. These data were used to

TABLE 4-35 Daily Streamflow Simulation Performance in the Feather River Basin at Oroville, California

Parameter	Calibration (2000–2005)	Validation (1990–2000)
Annual Relative Bias	0.00	0.00
Daily correlation	0.91	0.93
Daily NSE	0.81	0.84
Daily NSE (log flow)	0.84	0.79
RMSE/Obs Mean	0.46	0.67
MSE/Obs Mean	0.19	0.16

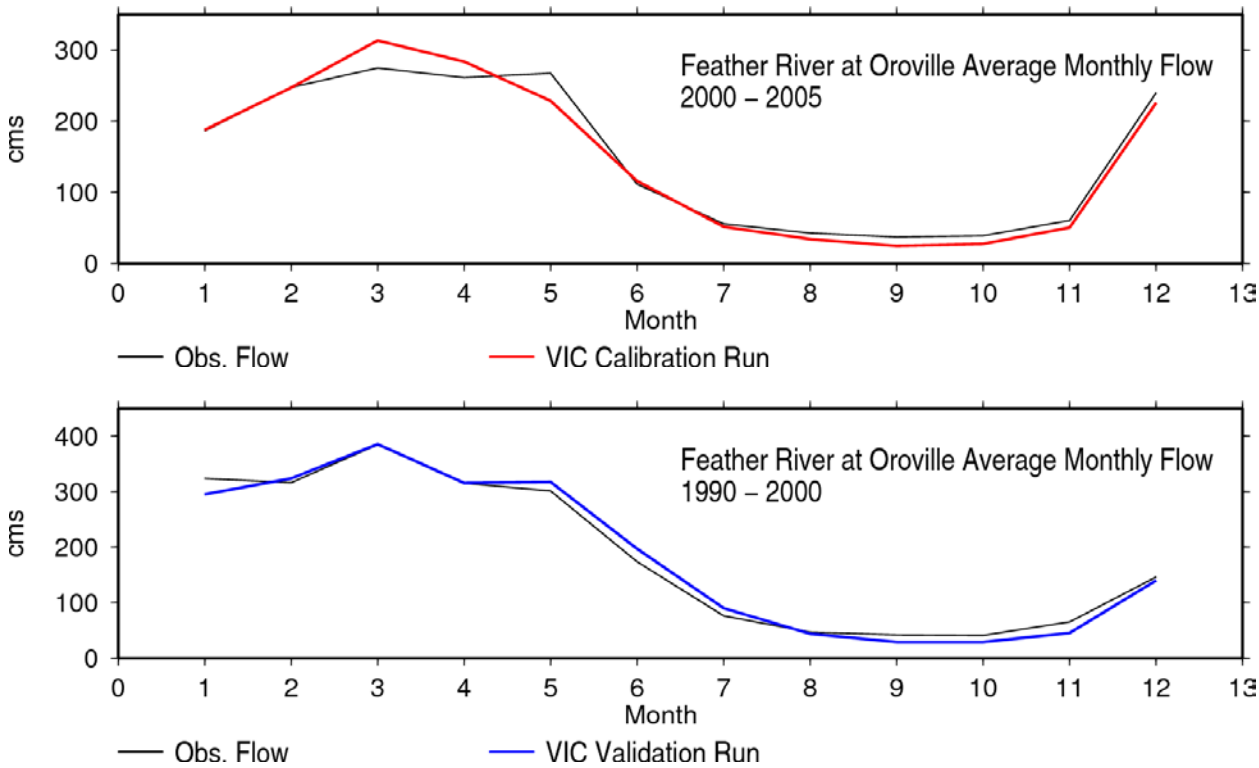

FIGURE 4-88 Monthly Flow at Oroville, California, during the Calibration (upper panel) and Validation (lower panel) Periods for Observed Naturalized Flows and Modeled Flows

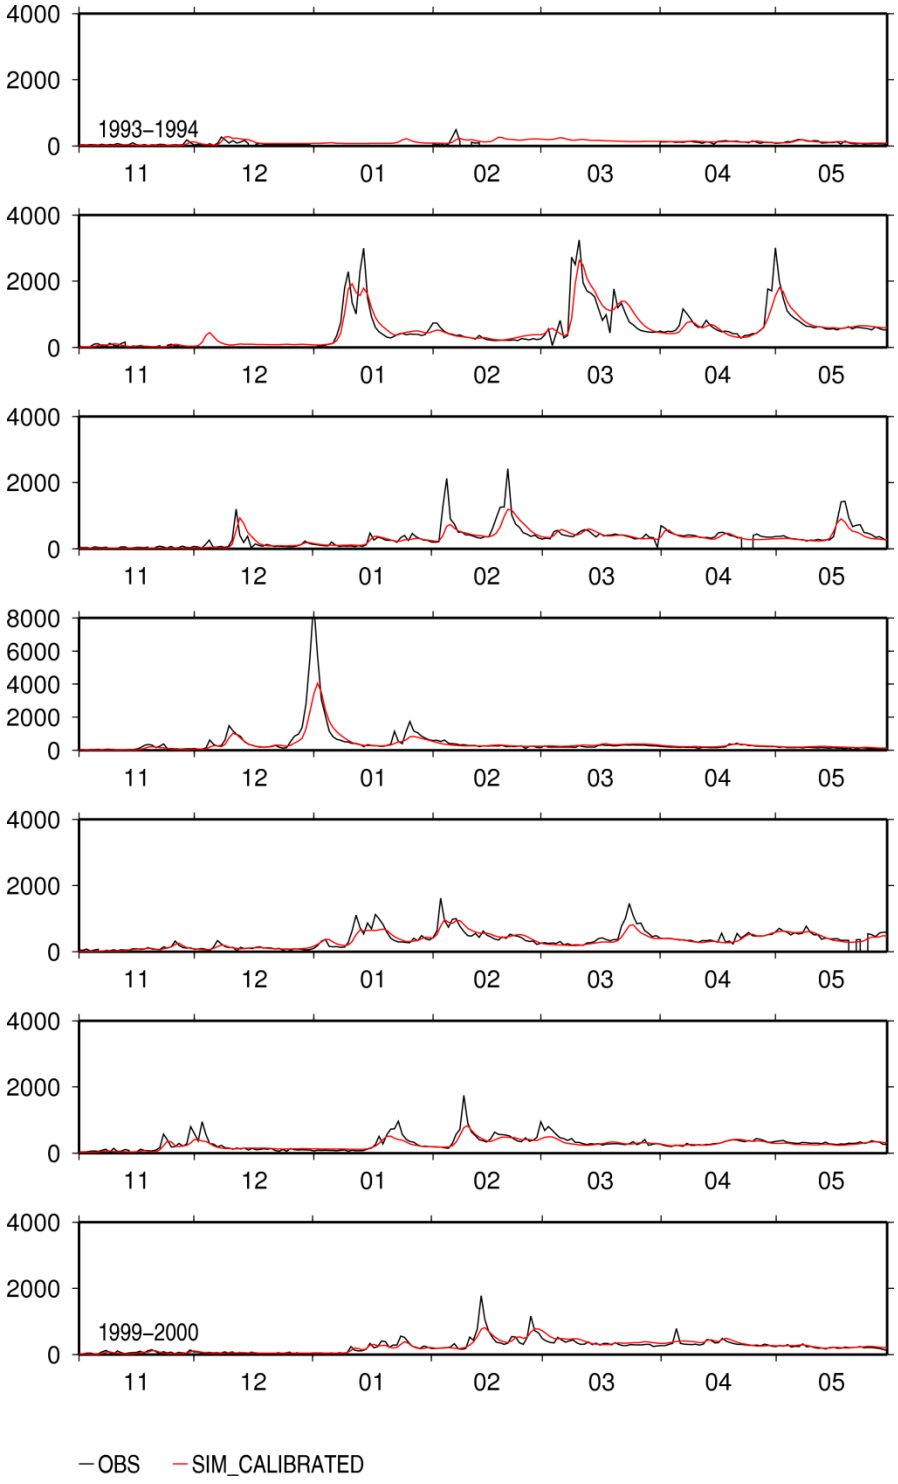

FIGURE 4-89 Daily Naturalized Observed (black) Flow (cms) and EHFS Simulated Flow (red) at Oroville, California (each panel shows a shortened water year for clarity, 1994–2000)

drive the hydrology model, which was previously spun-up using near-real-time observed meteorology. Figure 4-90 shows the observed naturalized flows and the ensemble average for the 1-, 2-, and 3-day lead forecasts for April–September, 2011. No streamflow assimilation was applied for the Feather River Basin because it is dominated by fast concentrations of flow rather than by the subsurface flow where the data assimilation process operates.

Seasonal Forecasts. Forecast results for the demonstration period are shown at Oroville in Figure 4-91, which presents the monthly average 12-month seasonal ensemble forecasts issued April 1, 2011, with forecast flows to March 30, 2012. The black line shows observed flow; the box and whisker plots show EHFS results based on ensemble meteorological forecasts. The boxes show the 25%, 50%, and 75% non-exceedence flows, while the whiskers represent the minimum and maximum ensemble members.

The close agreement between the ensemble and black lines and the center of the box plot indicates that the hydrologic model does a good job representing initial conditions and therefore the seasonal hydrologic conditions. The ensemble range also brackets flows in all months.

ESP flow forecasts' skill is in the prediction of the volumetric April–July flow. Figure 4-92 shows an evaluation of the April 1 EHFS April–July volumetric flow forecast at Oroville. Each ensemble member of the ensemble seasonal forecast is equally probable. Each quantile (ESPXX; x-axis) represent the chance that the flow threshold will not be exceeded. The minimum and maximum values indicate the full range of uncertainty. The ensemble mean volumetric forecast indicates a convergence of the traces and compares well with the observed value.

Products. The database was populated with streamflow and meteorological forcing data. For each forecast day (daily April 1, 2011, through September 30, 2011) an assemblage of the ensemble average of the medium-range forecast and the seasonal forecast issued before the medium-range forecast date was entered in the database. Each forecast day had 344 days with forecast flows in cubic feet per second. The first 13 days are from the medium-range forecast, and the remainder of the forecast period is from the daily data of the seasonal forecast.

The medium-range ensemble meteorological forecasts are derived from the pre-processed GFS medium-range weather forecasts. It includes one-eighth-degree gridded, 13-day, 15-member, hourly meteorological data, issued daily for the period of 1990–2010. It includes data parameters such as precipitation, air temperature, vapor pressure, atmospheric density, wind speed, specific humidity, incoming shortwave radiation, and incoming longwave radiation.

The seasonal forcing group includes a 49-member ensemble hourly meteorological forcing data at one-eighth-degree spatial grid resolution for a horizon of 366 days. The dataset includes data parameters such as precipitation, air temperature, vapor pressure, atmospheric density, wind speed, specific humidity, incoming shortwave radiation, and incoming longwave radiation. The dataset is static as each trace is a resampling from each of the 49 or 50 years (leave 1 year out when overlapping) in the 1960–2010 period.

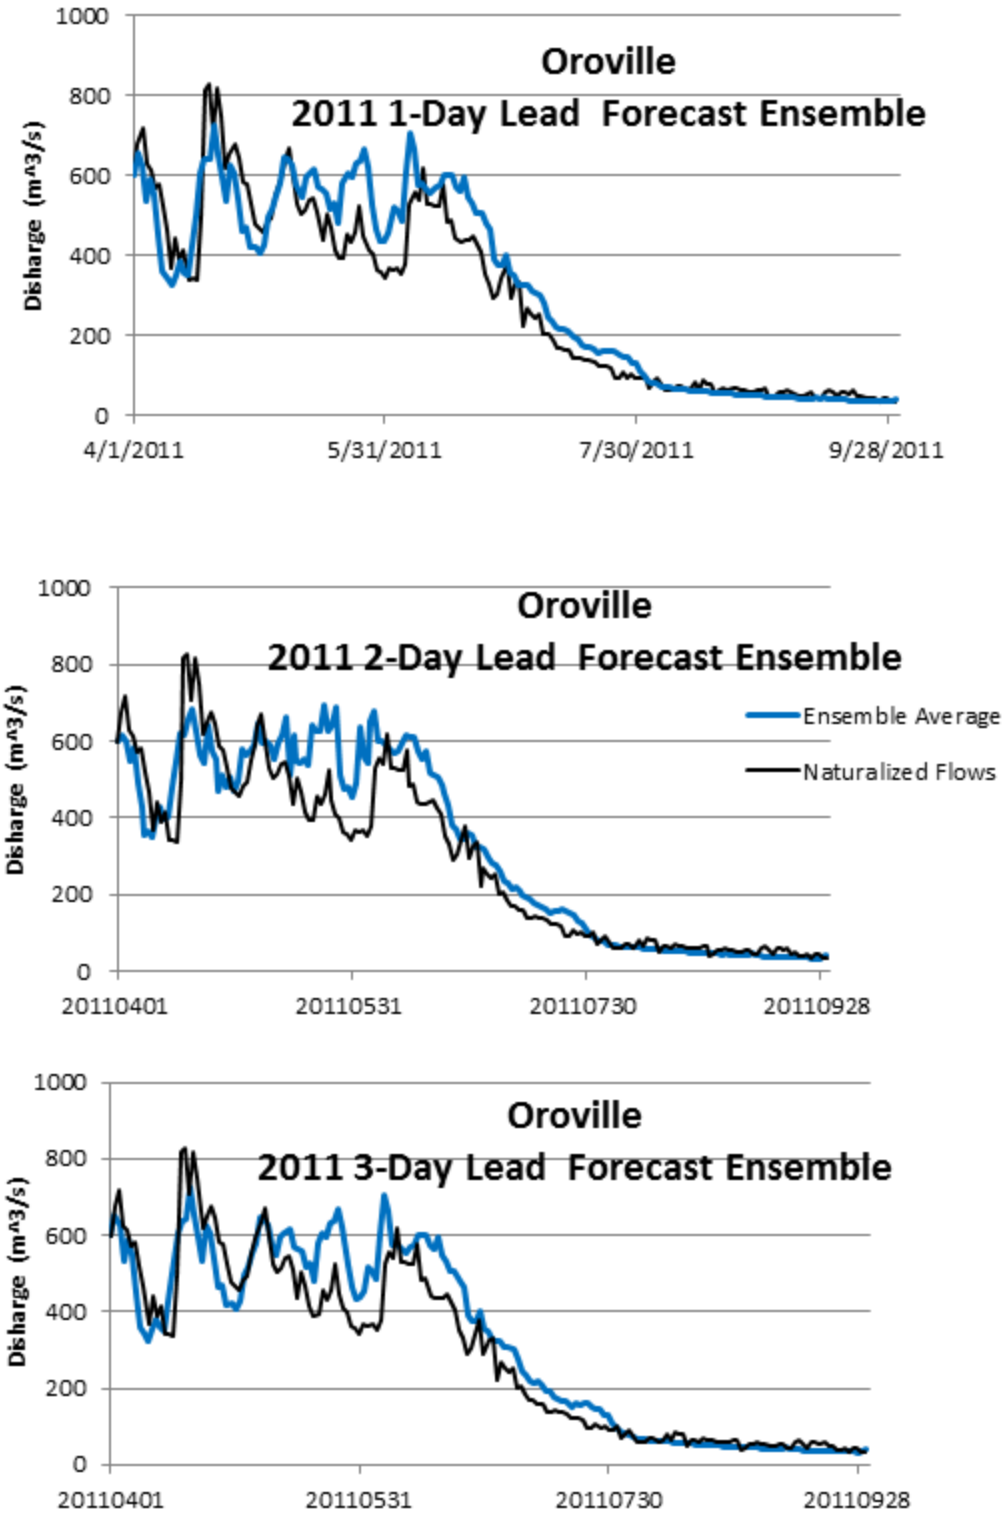

FIGURE 4-90 Observed Flow at Oroville for 2011 and EHFS 1-, 2-, and 3-day Lead Ensembles (Blue is the mean ensemble forecasts; the black line is the naturalized observed flows.)

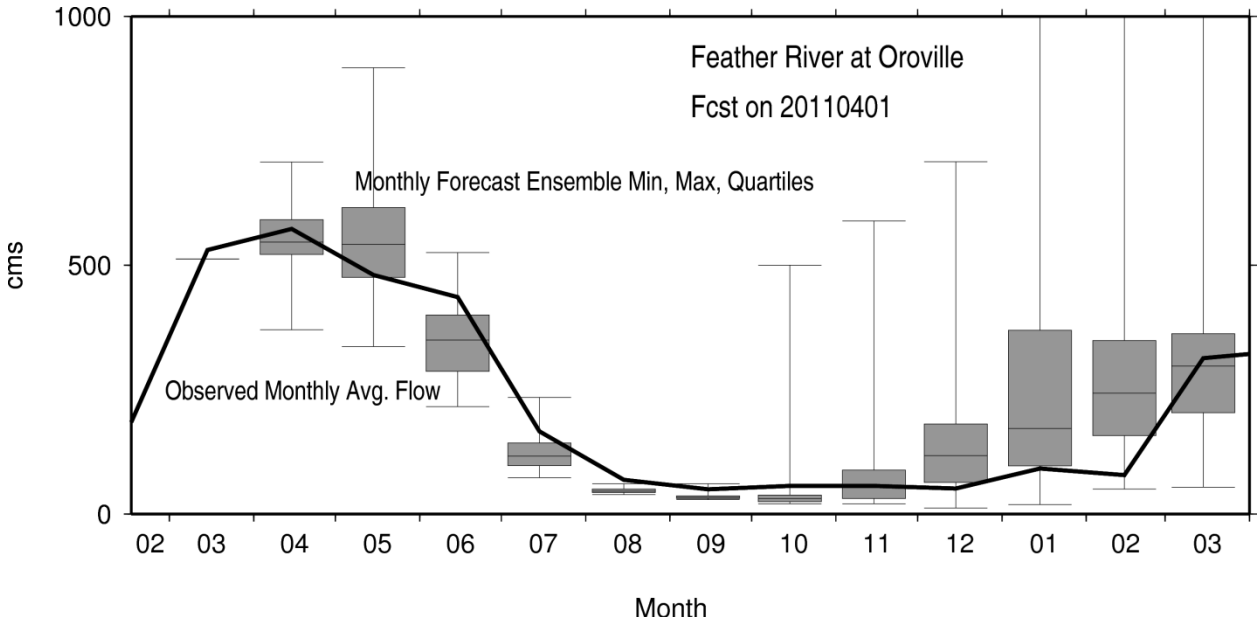

FIGURE 4-91 Seasonal Ensemble Forecast at Oroville, Issued April 1, 2011, with Forecast Flows to March 30, 2012 (The black line shows actual observed flow and the box and whisker plots show EHFS results based on ensemble streamflow forecasts. The boxes show the 25%, 50%, and 75% non-exceedence flows, while the whiskers represent the minimum and maximum ensemble members.)

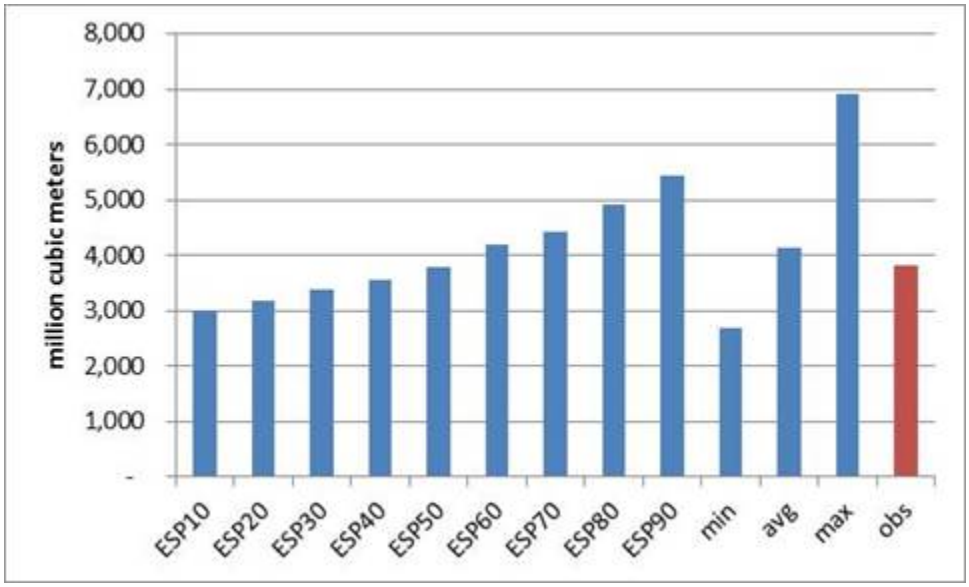

FIGURE 4-92 Evaluation of the April 1, 2011, EHFS April–July Volumetric Probabilistic Forecast for the Feather River at Oroville (The ESPXX shows the seasonal volume corresponding to 10, 20, … 90% nonexceedance probabilities. The ensemble forecast mean, minimum, and maximum are also shown to provide the entire range of the uncertainty.)

4.4.2.2 Seasonal Model

The Oroville Complex retrospective analysis was simulated over a 6-month period that ran from April 1, 2011 to September 30, 2011. For HydroSCOPE, a full DAKOTA optimization is run for each day of the simulation and like the Aspinall Cascade analysis, an adequate balance between simulation run time and accuracy was obtained with 500 simulations per optimization. Each day of the simulation consists of seven steps, as follows:

1. Obtain the 6-month hydrologic forecast from the EHFS.

2. Run the DAKOTA optimization to maximize revenue and the IRF score.

3. Select the optimal release schedule from the DAKOTA results and extract the predicted reservoir releases.

4. Perform a deterministic simulation that uses the predicted reservoir releases, historical inflows from data for the first day, and forecasted inflows for the balance of the simulation.

5. Write the results from Step 4 to the common database for use by the other tools.

6. Extract the system conditions at the start of the second day for use as the initial condition of the next day's simulation.

7. Return to Step 1 and repeat for each day running from April 1, 2009, through September 30, 2009.

For each day of the analysis, a single baseline simulation is run using the forecasted inflows that sets the baseline release schedule for each reservoir, which are the releases necessary to keep each reservoirs elevation on target. During the optimization, DAKOTA makes adjustments to the baseline release schedules to simultaneously maximize total revenue and the environmental score. For each timestep in a simulation, the operational logic calculates the maximum and minimum release rates from Oroville that meet the downstream flow targets and does not allow DAKOTA to make adjustments beyond those limits. A penalty function is used to adjust the revenue and environmental score downward based on the number of flow-target violations that occur during the simulation as compared to the 6-hr adjusted historical data. DAKOTA will terminate the optimization when no more improvements can be found or the 500 simulation limit has been reached. The latter case was used in the 183 optimization runs used in this analysis. The operational logic and penalty function is explained in more detail below.

Model Domain. Figure 4-93 shows a conceptualization of the Oroville Complex model as used by HydroSCOPE. The system consists of four reservoirs (Oroville, Diversion, Forebay, and Afterbay), and seven river reaches. The upper boundary of the model is the inflow point to Oroville and the lower boundary is the Gridley Gauge (California Data Exchange Center gauge ID GRL 09152500), which is a little over 16 mi downstream of Oroville at river mile 50.

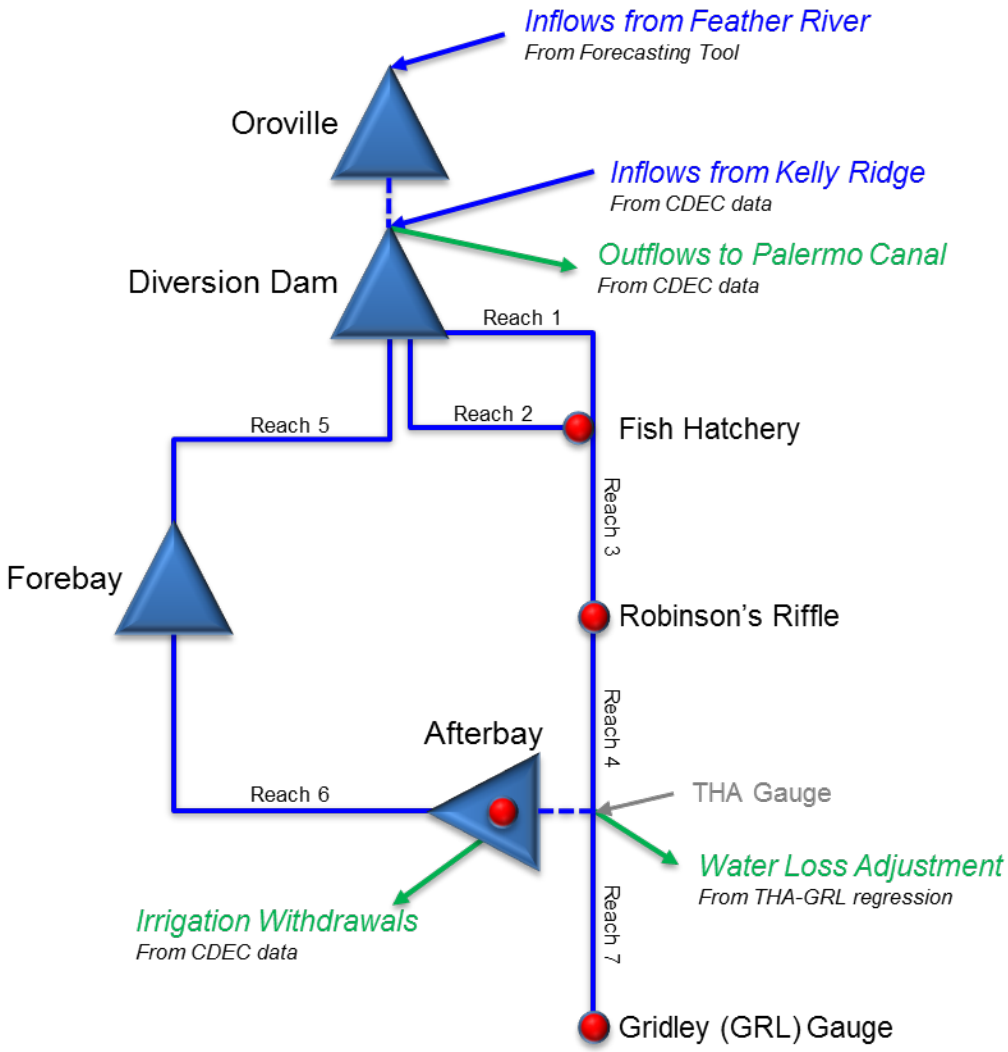

FIGURE 4-93 Conceptual Model for the Seasonal Tool (System inflows are in blue and system outflows are in green. The red dots are the environmental assessment points. Reaches 1 and 2 are treated as zero-length reaches in the model.)

Inflows to the model occur where the Feather River flows into Oroville and where the discharge canal from the Kelly Ridge Power Plant enters the diversion pool. The EHFS tool supplies Feather River inflows to Oroville while the Kelly Ridge discharge is based on historical values. Outflows are from Oroville to the Palermo Canal, from the Afterbay to four irrigation canals, from the bottom boundary at the Gridley gauge and via evaporation from all the reservoirs and reaches. Outflows to the Palermo Canal and the Irrigation canals are based on historical data from the CDEC website.

An outflow adjustment point was added to the model to correct for an apparent flow loss between the CDEC THA gauge in the Feather River immediately below the Afterbay and the GRL gauge (Figure 4-94). The correction is made using a linear relationship between the 6-hr average flows at THA and at GRL one timestep later (the model timestep is 6 hours) and is only applied when calculating the penalty function and the IRF score. According to the California Department of Water Resources staff, the reason for the discrepancy between the gauges is unknown.

Optimization Multipliers. HydroSCOPE begins the optimization process by first calculating a set of default releases from each reservoir that are based on historical operations. The default releases are calculated by the model, using the forecasted inflows and adjusting the reservoir releases to match the reservoir elevations and to honor any flow targets while staying as close to the historical releases as possible. Each optimization simulation modifies the default releases by individually adjusting

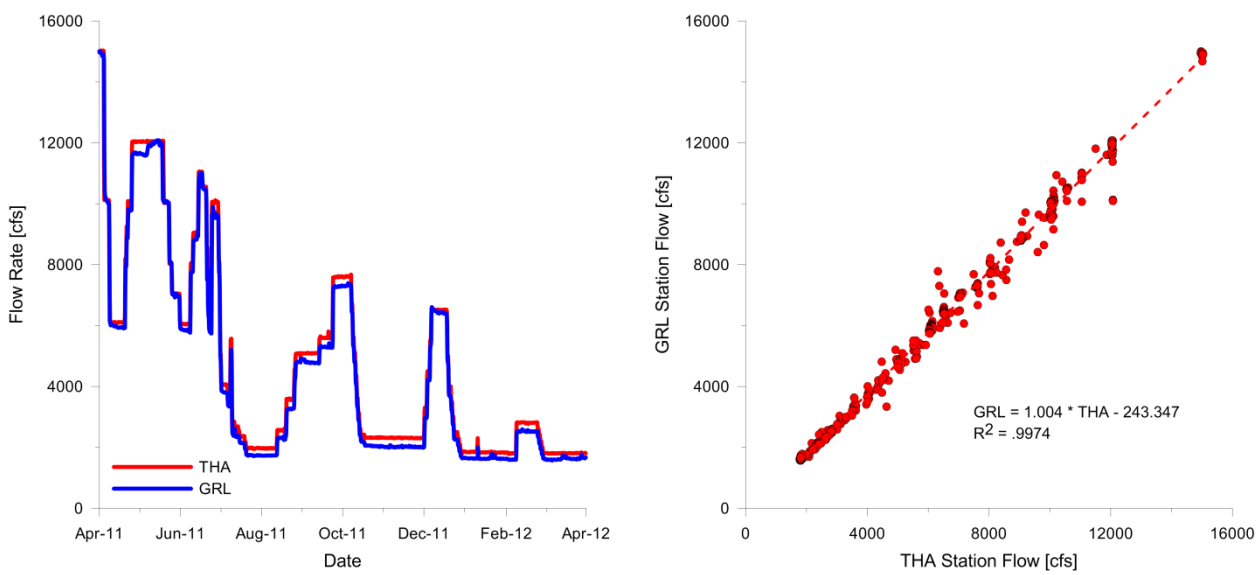

FIGURE 4-94 Relationship between Flows at the THA Gauge and the GRL Gauge (The THA gauge is in the Feather River immediately downstream of the Afterbay inflow point. The GRL gauge is approximately 7.5 mi further downstream.)

a set of multipliers that are applied to the default releases. The optimal solution is the one where the modified outflows maximize the IRF score and revenue. When conducting an optimization from the present day forward, the historical operations are determined by either matching forecasted conditions to similar historical conditions or by defining a set of operational rules that can accurately mimic on-the-ground operations. For the retrospective analysis, the historical operations are the actual operations during the analysis period. This was done to help reduce the computational overhead of conducting the analysis and to save time coding and debugging the complex operational logic.

The default releases are modified using 15 multipliers for each reservoir, with the first seven multipliers applied to the first 7 days of the simulation, the next three multipliers to the next 3 weeks of the simulation, and the final five multipliers applied to the remaining 5 months of the simulation. Multipliers are allowed to vary between 0.95 and 1.05 (±5%). For the Oroville analysis, the multipliers are applied to the midpoint of their respective time ranges and mapped to the 6-hour time step of the model using piecewise Hermite interpolation (Birkhoff et al. 1968). This avoids the discontinuities that would otherwise occur and creates more realistic transitions between each of the multiplier time ranges. A plot showing the multiplier function for the releases from Oroville for the April 1 simulation is shown in Figure 4-95.

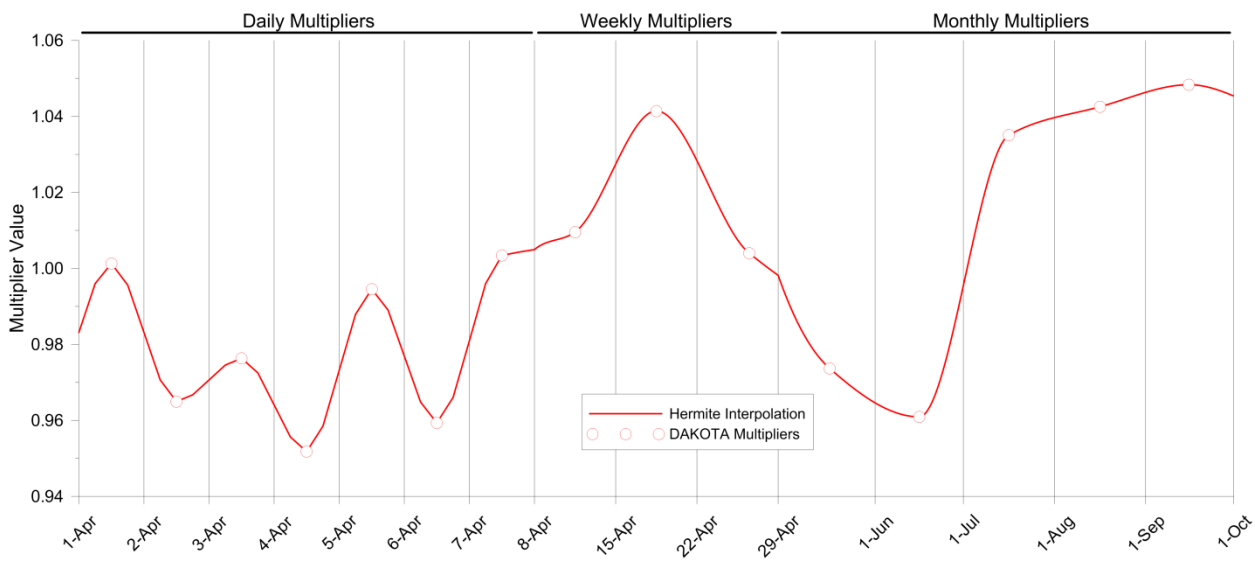

FIGURE 4-95 The Multiplier Function for the Releases from Oroville (DAKOTA is allowed to adjust each point ±5% as part of the optimization process. A Hermite interpolation scheme is used to map the multiplier points to the 6-hour time step of the model. Note that the x-axis scale changes.)

Data Processing. Because the model calculates the various flows using a conservation of mass approach, the historical data from which the default releases are based were pre-processed to insure the water balance is correct and to identify and fix any problems or inconsistencies in the underlying data. After cleaning the data to account for any data gaps, outliers, and so forth, the data are mapped to the 6-hour time step of the seasonal model and then used to calculate the reservoir elevations over time. For the water balance to be correct, the calculated reservoir elevations must match the historical elevations, which was not the case here so adjustments to the Oroville inflows, and the releases from the Thermalito Diversion Pool (TDP), Thermalito Forebay (TFB), and Thermalito Afterbay (TAB) were made using a set of assumptions as follows:

1. Outflow data from Oroville are assumed correct, including diversions to Palermo Canal;

2. Inflow data from Kelly Ridge Reservoir are assumed correct;

3. Elevation data are assumed correct;

4. Flow at in the Feather River below Thermalito (known as the THA gauge) is assumed correct; and

5. Flow in the power canal is one-way from the TDP to the TFB.

The assumption that the outflow data from Oroville are correct was made so that Oroville could be balanced in isolation from the other reservoirs and to keep the outflows consistent with the standalone analysis done by the day-ahead planning model, which requires only an independent water balance for each day rather than a full time-series balance.

To balance the water in Oroville, the inflows necessary to meet the reservoir elevation over time were back-calculated as a function of the reservoir outflow and storage using a simple water balance equation:

$$I^n = \frac{S^{n+1} - S^n}{\Delta t} + O^n \tag{4-10}$$

where I is the sum of all inflows, O is the sum of all outflows, S is the storage, and Δt is the time step (6-hours in this case). The superscripts, n and $n+1$, refer to time step indexing with n being the value at the start of the current time step and $n+1$ being the value at the end of the current time step. Over the 1-year timespan from April 1, 2011, through March 31, 2012 (recall that the last day of the analysis, September 30, 2011, needs a 6-month simulation to evaluate, which requires the data to be processed through March 31, 2012), this resulted in calculated inflows that were noisy, oscillated wildly, and ranged from –3780 cfs to 53000 cfs. To correct for this, the Tao (1999) methodology was used to calculate the inflows based on the relationship between the moving average of the inflows calculated with the water balance equation against the moving average of the inflows obtained from streamflow measurements. In this case, the naturalized "perfect" forecast from the EHFS was used as the streamflow measurement inflow. Figure 4-96

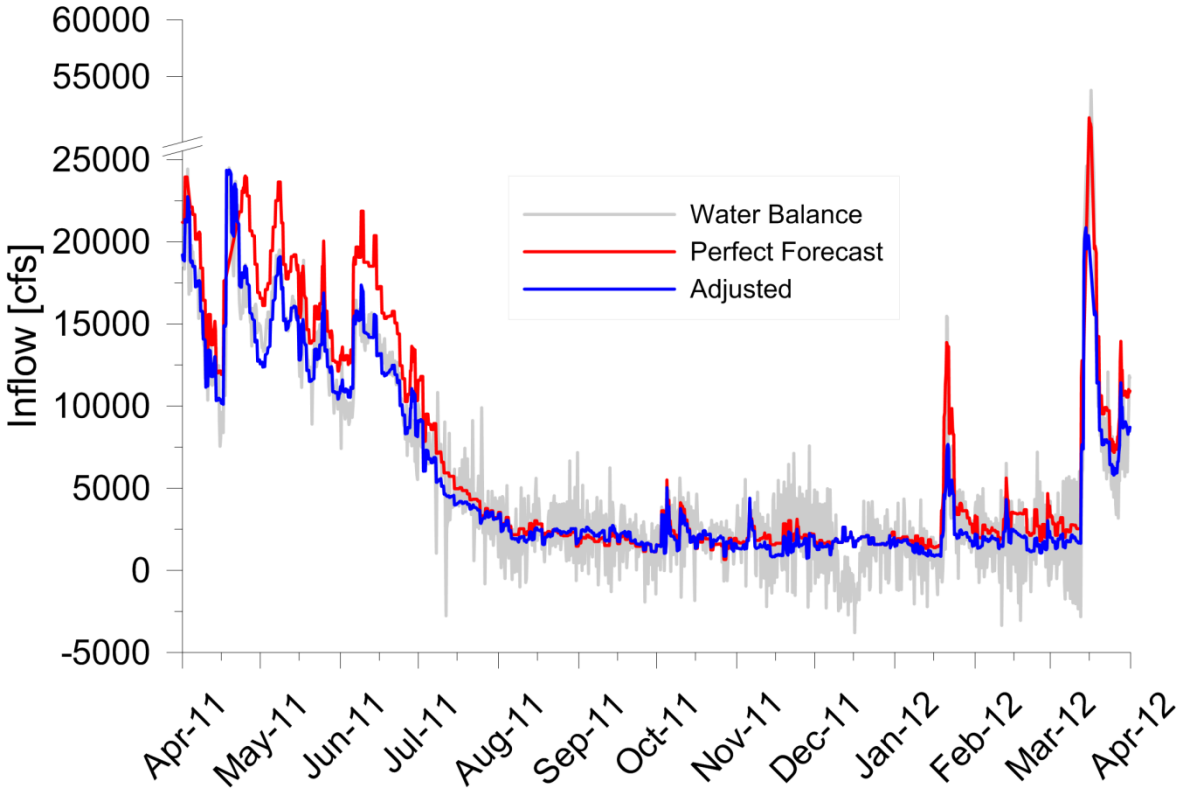

FIGURE 4-96 The Perfect Forecast Supplied by the EHFS (red line) Includes Naturalized Flows That Must Be Adjusted to Match the Water Balance in the Reservoir; a Linear Transformation Is Used to Make the Adjustment (blue line)

plots the water balance derived inflows, the EHFS perfect forecast, and the final adjusted inflows used in the model.

When compared to the EHFS perfect forecast, the adjusted inflows produce approximately 15.5% less inflow over the full 12-month period. The reason for the difference is that the EHFS perfect forecast provides an unregulated inflow estimate while the adjusted inflow implicitly accounts for upstream regulation, withdrawals, and the like. To insure that the daily forecasts supplied by the EHFS were not similarly biased, a linear transformation was applied to each of the forecasts based on the relationship between the adjusted inflows and the EHFS perfect forecast. Figure 4-97 shows a plot of the 183-day (6-month) inflow volume for each forecast date of the original forecasts and the adjusted forecasts. The differences during the summer months are slight due to the lower inflow volumes for those forecasts. The volume difference on April 1 is corresponds to an average flow difference of about 1000 cfs.

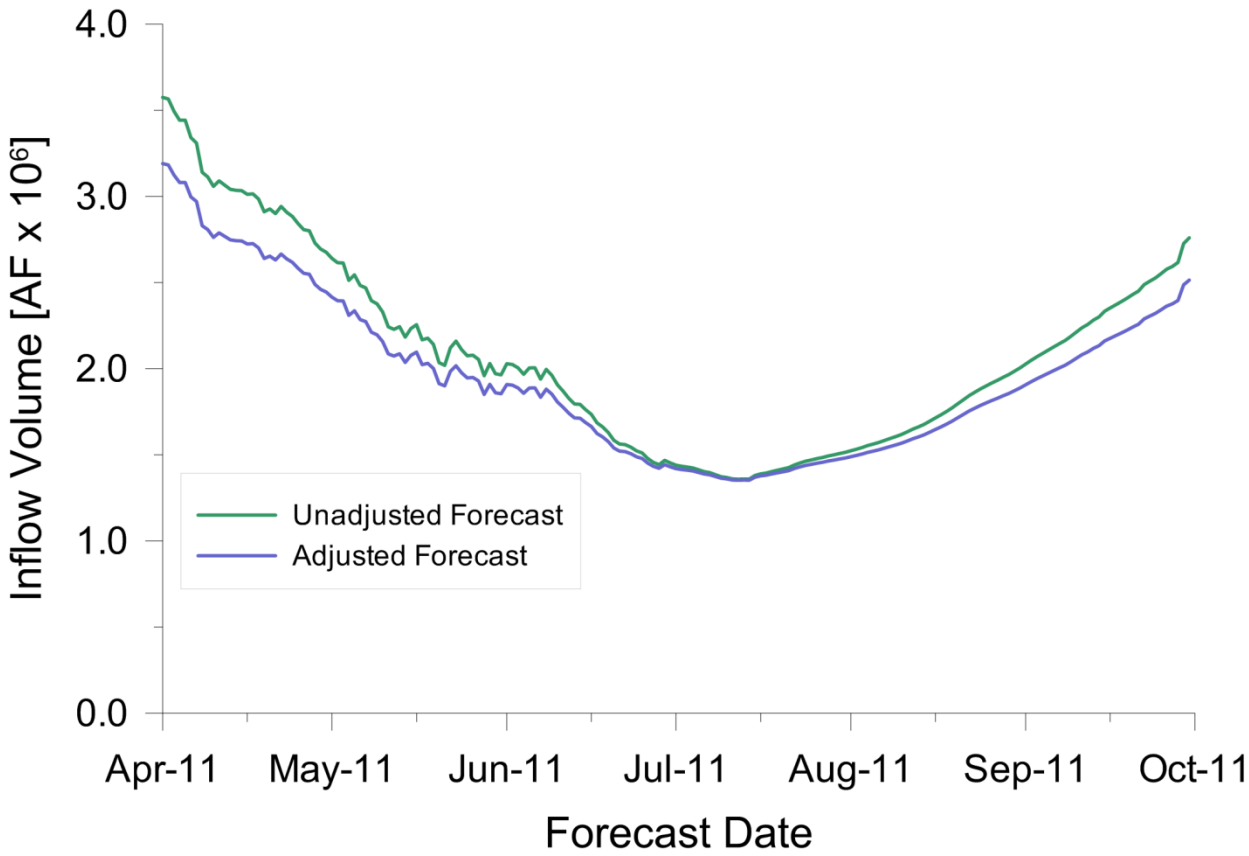

FIGURE 4-97 Difference in the Forecasted Inflow Volume for a 6-month Period from Each Date before and after the Adjustment Was Made (The adjusted inflows are used in the retrospective analysis.)

Calibration. The model was calibrated using DAKOTA with an objective of minimizing the root mean squared error (RMSE) between the field data and the simulation results for a full year simulation from January 1, 2005, through December 31, 2005. Like the Aspinall Cascade calibration, the decision variables (i.e., the calibration parameters) are the *a*, *b*, and *c* coefficients of the wind function (equation [3-8]), the solar extinction coefficients (equation [3-5]), the minimum and critical diffusion coefficients, and the A_{1-6} coefficients used to calculate the effective diffusion (equations [3-21]–[3-25]). Data were available for the vertical profile temperatures in Oroville, the outflow temperatures from each of the four reservoirs, and at the Fish Hatchery and the Fish Barrier Dam. To help reduce the noise in the data as well as the simulated temperatures, 6-hour averages were used for comparing the vertical temperature profiles and various flow temperatures. While the simulation was run for the entire year, the calibration only looked at April through mid-December to provide a better match to the important summer months and to avoid anomolies associated with some very high releases during the winter.

Penalty Function. The model uses a penalty function approach that reduces the value of the IRF score based on how well the model is able to comply with the regulatory constraints. Regulatory constraints for the purpose of the penalty function are the flow and elevation constraints that are defined in the operational regulations. The temperature constraints are evaluated by the IRF tool as part of determining the IRF score.

The penalty function is an integrated score from 0 to 1 that is calculated by comparing the model's compliance to the historical compliance across all constraints (historical compliance is based on historical data that have been processed to match the 6-hour time step of the model). If the model's performance matches or exceeds the historical performance, then the penalty function returns a 1; otherwise, the penalty function rapidly declines to 0 as a function of the percentage of time the system is out of compliance (Figure 4-98).

The IRF score is multiplied by the penalty function and thus provides feedback during the optimization process for DAKOTA to learn what release schedules perform the best. Without the penalty function, the binary nature of the constraints (met vs. not met) provides little structure for DAKOTA to work with. The constraints considered in the penalty function are as follows:

1. High Flow Channel

    a. If the average flow in any 1-hour period exceeds 2500 cfs during the time between October 15 and November 30, then minimum flow from October 1 through March 31 is the highest 1-hour flow during the time between October 15 and November 30, minus 500 cfs.

2. Down Ramp Rates (DRR) at Robinson's Riffle (RR)

    a. 600 cfs < flow at RR ≤ 2500 cfs → DRR ≤ 300 cfs/day

    b. 2500 cfs < flow at RR ≤ 3500 cfs → DRR ≤ 500 cfs/day

    c. 3500 cfs < flow at RR ≤ 5000 cfs → DRR ≤ 1000 cfs/day

    d. 5000 cfs < flow at RR → DRR ≤ 5000 cfs/2 hours

3. Afterbay Elevation Criteria (Figure 4-99)

    a. Between March 15 and May 31, the water surface elevation must pass through 133.5 ft at least once every 9 days.

    b. Between July 1 and August 15, the water surface elevation must be between 132.0 ft and 135.0 ft.

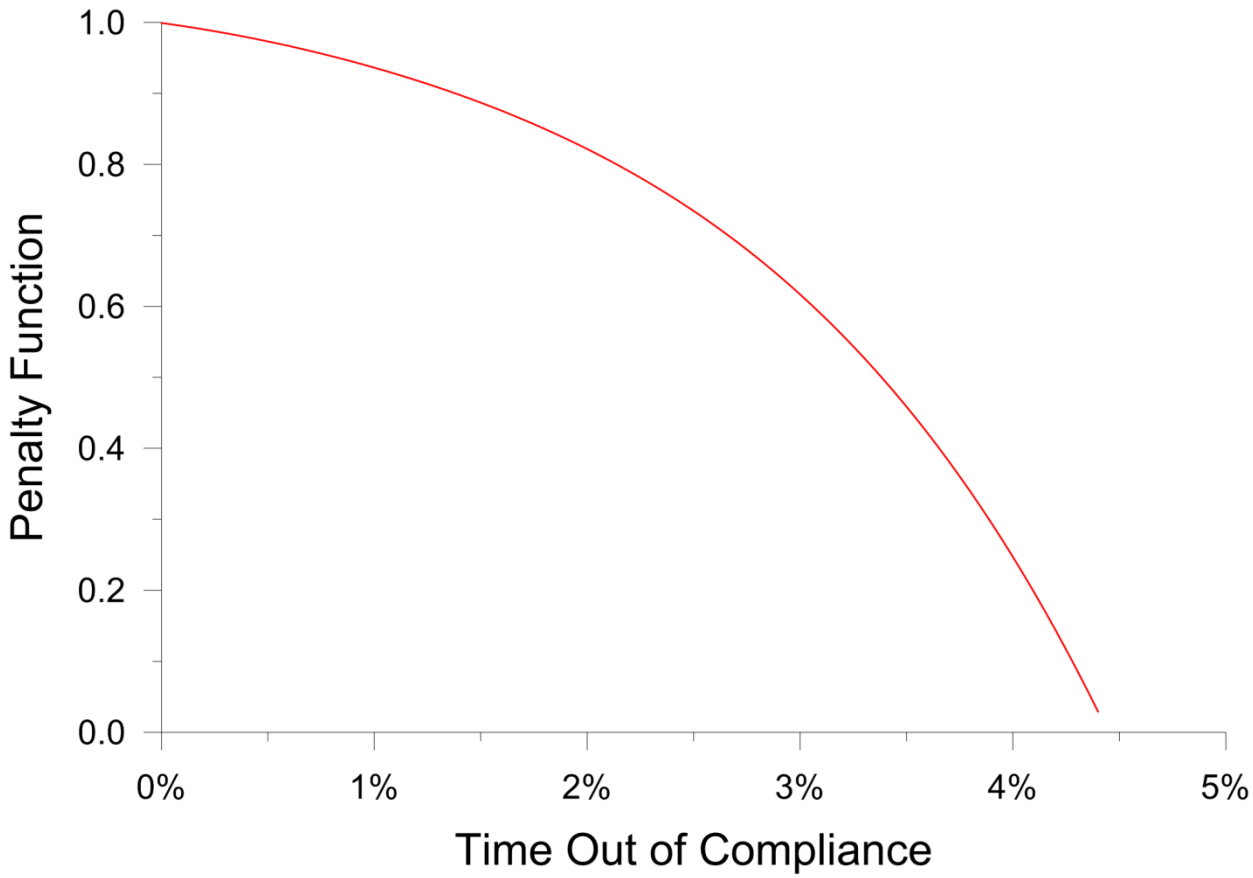

FIGURE 4-98 The Penalty Multiplier is a Function of the Cumulative Time That All the Objectives Are out of Compliance (The time out of compliance is relative to the time out of compliance for the 6-hour averaged historical data [if the model matches the historical performance, the time out of compliance is 0%].)

c. Between August 16 and September 15, the water surface elevation must be between 131.0 ft and 136.0 ft.

d. Between May 1 and September 30, the water surface elevation must be maintained for a continuous 12-hour period at least once a month (30 days was used in the model).

Three additional constraints, the minimum flow rates at Robinson's Riffle (600 cfs), the minimum flow rates in the high-flow channel (1700 cfs or 1000 cfs as a function of previous and current water years, water elevation in Oroville, and the time of the year), and the diversions to the fish hatchery (based on historical data), are built into the operational logic of the model and thus are never violated. The clearly defined nature of these constraints provides the operational structure for the model to develop a default release schedule that can then be optimized. Including them in the penalty function instead of as part of the operational logic provides too much freedom to the optimization process, which then requires many more simulations to complete.

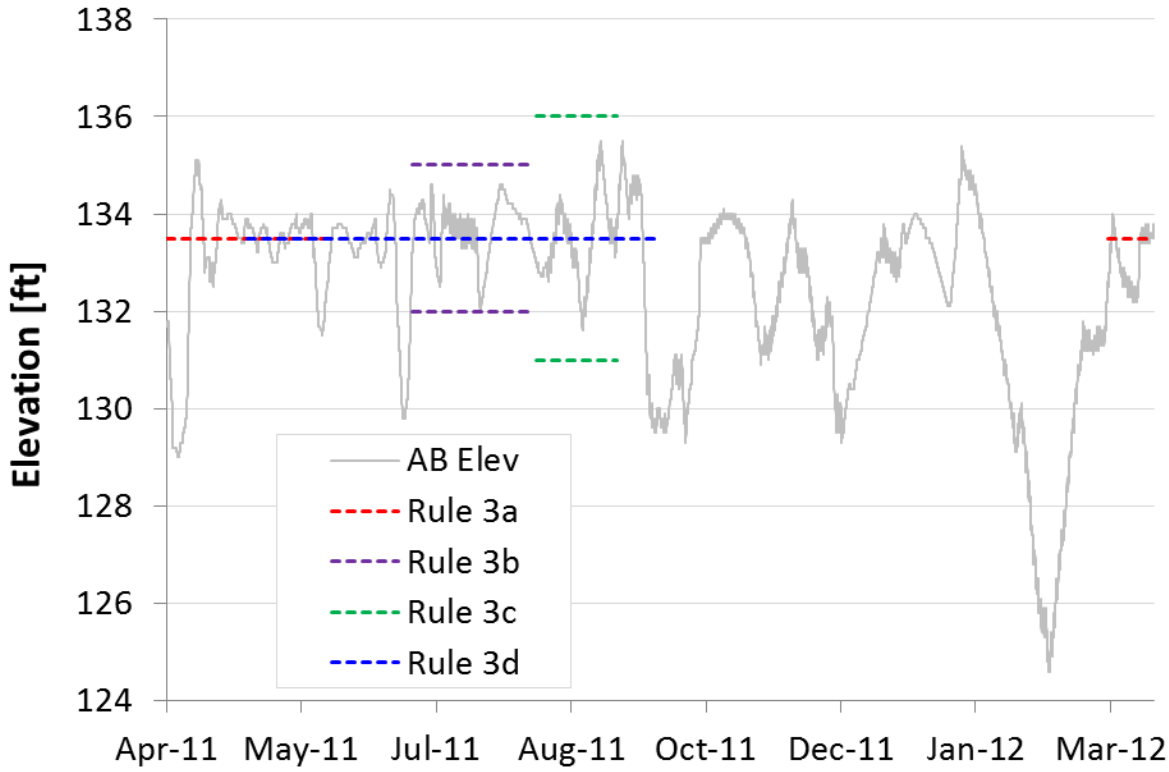

FIGURE 4-99 Elevation Objectives for the Afterbay during the Retrospective Analysis Simulation Period

Operations. At the start of each optimization run, HydroSCOPE runs once using the forecasted inflows to determine the default outflows for the next 6 months, which are then subsequently modified by DAKOTA to find the optimal solution. The default outflows are determined by setting target elevations for the four reservoirs and then adjusting the outflows as a function of the inflows to match those targets. Delay factors are included that allow corrective releases to become larger as a function of the difference between the current elevation and the target elevation, but that limits the correction in any one time step to a percentage of the difference. This attenuates the programmatic (i.e., unrealistic) oscillations that would otherwise occur with the model fully correcting and re-correcting at each time step.

Starting with the elevation target, an initial release value is calculated to correct any differences and then scaled with the delay factor. It is then compared to the maximum and minimum releases possible that would still meet the regulatory flow constraints and that are within any physical limits of the system (e.g., maximum flow rate through an outlet structure). If the initial release value is between those values, then it is the used as the default release for that time step. If it is above or below those values, it is given the appropriate maximum or minimum value. During optimization, DAKOTA is allowed to adjust the flow rate between the maximum and minimum values calculated at each time step. Note that because the inflows to the TDP, TFB, TAB, and LFC may change during optimization, the maximum and minimum allowable

releases may be different than the default case and thus are recomputed at each time step. Operational criteria specific to each reservoir are described below.

Oroville. Releases from Lake Oroville are used to control the flows throughout the system. The minimum flow targets in the HFC are set as a function of the current and previous water year's inflow volumes, and the September 30 projected water surface elevation. The CDWR rule set describes four scenarios for determining the minimum flows, High Runoff, High Runoff/Low Storage, Low Runoff, and Low Runoff/Low Storage. For the 2011 simulation year, the scenarios for all forecasts were predetermined to be High Runoff, which sets the minimum flow rate in the HFC at 1700 cfs from October 1 through March 31 and at 1000 cfs otherwise.

The Oroville elevation targets in the model are based on the required available storage curves for the low-flow forecast condition (Figure 4-100). Without confirming with the California Department of Water Resources (CDWR), it appears that the historic operations were operated to match the low-flow curve until about mid-April, when an adjustment was made to match the high-flow curve, until approximately October 1, when operations reverted back to honoring the low-flow curve. Since the low-flow curve encompasses the majority of the historical elevation curve, it is used as the target for calculating the default outflows.

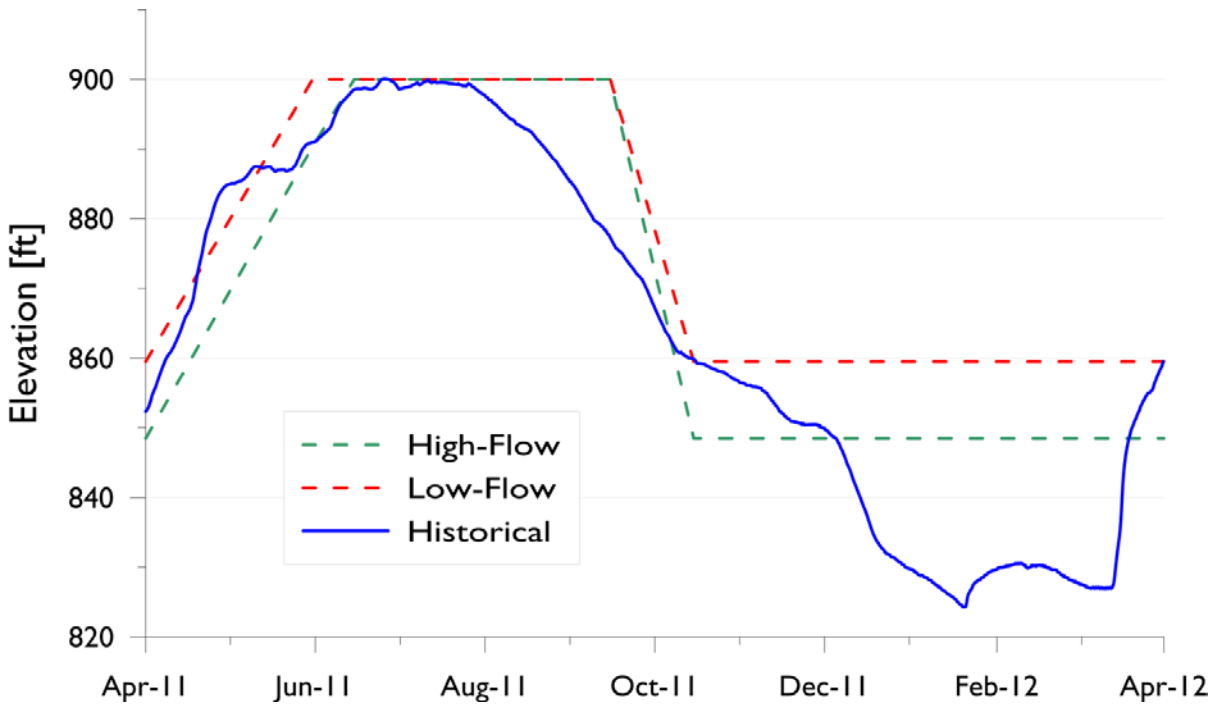

FIGURE 4-100 Historical Elevation for Oroville as Compared to the High-Flow and Low-Flow Storage Requirements (The low-flow requirement is used as the target elevation in the retrospective analysis.)

The outflow structure for Oroville is unique in that it uses a shutter system that consists of two screened inflow chutes (called intakes) that are opened from top to bottom by successively removing a set of shutters that cover them. Shutter operations require that the number of shutters removed for each intake be within one of each other. For that reason, and to help simplify the modeling process, HydroSCOPE simulates a single intake with a total capacity equal to the sum of both intakes. Because the simulated elevations are not drastically different from historical elevations, the historical positions of the shutters were used in the model as opposed to creating logic that would remove/replace the shutters as a function of the system state. The historical shutter positions were supplied by the CDWR (Bui 2013). A schematic of a shutter system is shown in Figure 4-101.

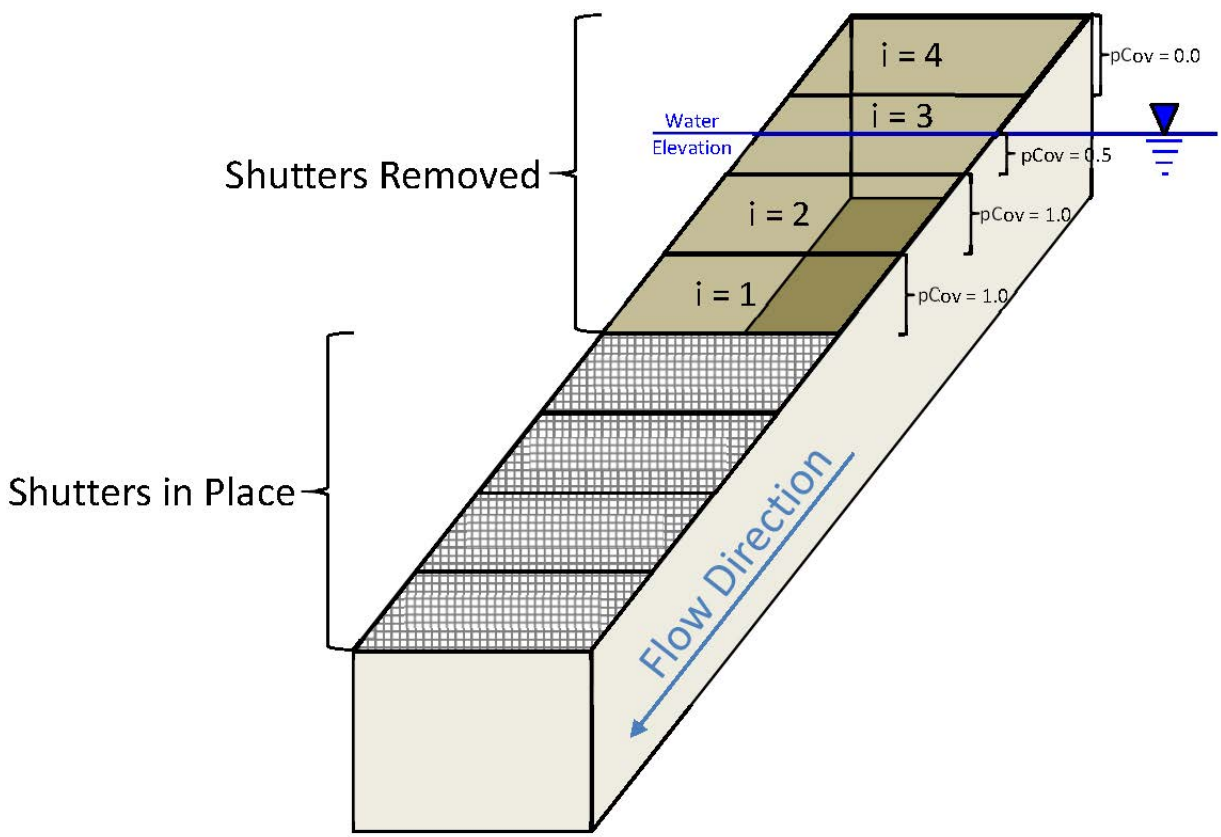

FIGURE 4-101 Schematic of an Eight-Shutter Intake System with Four Shutters Still in Place (Oroville has two intake structures with 15 shutters, each ranging in elevation from 612 ft to 891 ft. The example shows a hypothetical case with the water surface elevation at the mid-elevation of the third open shutter [i = 3] and the values of the *pCov* variable used in equations [4-11] and [4-12].)

The unique nature of the intake system required that the method be modified to calculate the withdrawal profile (i.e., how much water is released from each model layer) in the reservoir. The goal of the model is to honor the behavior and dynamics of the intake structures, but not necessarily the physics. Using the results from the U.S. Bureau of Reclamation hydraulic model studies of the Oroville intake structures (BOR 1965), an exponential relationship was developed that relates the flow through a shutter interval to the depth of the interval in the water column. The relationship preserves the empirical results from the BOR (1965) study and shows that approximately 90% of the total flow is withdrawn from the lowest 10% of the total depth (Figure 4-102). Total depth is defined as the water surface elevation to the top elevation of the highest installed shutter. In reality, the curve shown in Figure 4-102 is also a function of the flow rate, but since adequate calibration was achieved using the simpler configuration, including the additional complexity was deemed unnecessary.

To calculate the withdrawal from each layer, HydroSCOPE assumes a point withdrawal structure at the mid-elevation of each open shutter with the flow through each structure calculated using the following equations (these equations produced the results in Figure 4-102):

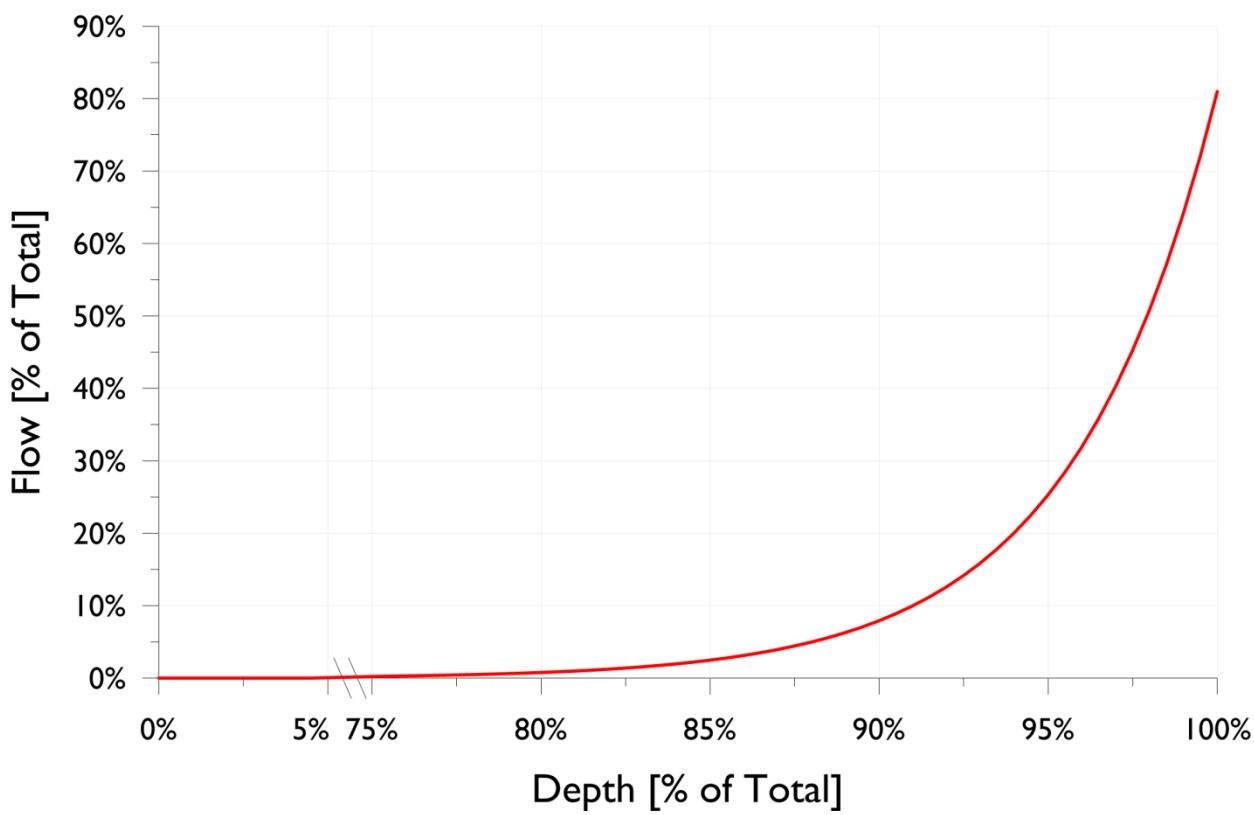

FIGURE 4-102 Outflow Rates are Calculated for the Mid-elevation of Each Open Shutter Using an Empirical Representation of the BOR (1965) Experimental Results (The function produces a cone of withdrawal that is skewed towards the bottom two or three open shutters.)

$$Q_o(i) = \begin{cases} Q_{tot}/Q_d & for\ i = 1 \\ Q_o(i-1) * p * pCov_i & for\ i > 1 \end{cases} \quad (4\text{-}11)$$

where $Q_o(i)$ is the total flow through shutter interval i, Q_{tot} is the total outflow from Oroville, p is the exponential decline rate (0.1901 is used here), $pCov_i$ is the percent submergence of each shutter (a value of 1 indicates the entire shutter interval is submerged), and Q_d is a flow divisor calculated as:

$$Q_d = \sum_{i=1}^{numOS} p^{i-1} pCov_i \quad (4\text{-}12)$$

where *numOS* is the number of open shutters. The index variable *i* assigns the value of 1 to the lowest, open shutter and counts upward from there. The intake structures in Oroville have 15 shutters each ranging in elevation from 612 ft to 891 ft. The cone of withdrawal for every outlet structure (i.e., for every open shutter) is calculated using the method described above and then summed to determine the complete outflow pattern in the reservoir.

Power generation is calculated using a relationship that describes the total flow through the power plant, the water surface elevation, and the power generated. The relationship is developed using historical data supplied by the CDWR for 2003 through 2009 (not every year was represented in the data set) and is defined as:

$$P = AQ^2 + BQh + Ch^2 + DQ + Eh + F \quad (4\text{-}13)$$

where *P* is power in megawatts (MW), *h* is the water surface elevation in feet, *Q* is the total flow through the power plant in cubic feet per second, and the coefficients A through F are listed in Table 4-36. The resulting function for the range of flows and elevations encountered during the analysis is shown in Figure 4-103.

TABLE 4-36 Power Calculation Coefficients for the Hyatt Power Plant

Coefficient	Value
A	2.4057E-21
B	9.0526E-05
C	5.1394E-04
D	–0.0305
E	–0.8752
F	372.1676

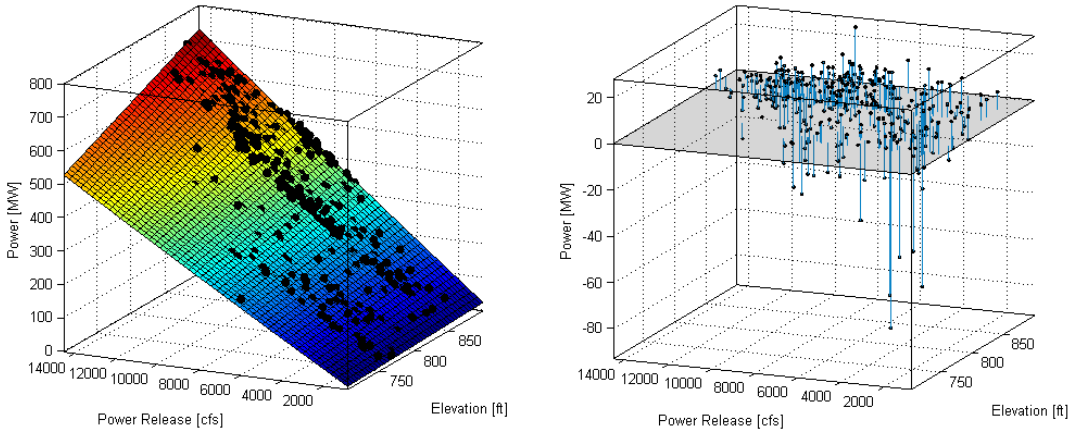

FIGURE 4-103 Plot of the Power Generation Curve for the Hyatt Power Plant (The plot on the right shows the residuals between the historical data and the curve fit. With the exception of cases with low flow and low elevations, the fits are within ±5% of the data.)

Thermalito Diversion Pool. Operations for the TDP are relatively simple as compared to Oroville. The TDP is operated to divert water to the fish hatchery, to the low-flow channel via the fish barrier dam, and to the power canal, which feeds into the TFB. Assuming that historical elevations for the TDP were a result of the operating logic and system needs at that time, the target elevation for the TDP is set as the historical elevations. From the CDWR California Data Exchange Center (CDEC) Web site, historical elevations for the TDP are only available for the beginning of each month, so linear interpolation was used to fill in the values during each month. Because the total storage of the TDP is relatively small, the error introduced from assuming a linear elevation change from month to month is small. Using the historical elevations keeps the relationship between the TDP and TFB elevations, which controls the flow rate through the power canal. Power generation from the TDP is assumed to be constant at a rate of 2.75 MW.

Thermalito Forebay. The TFB uses its historical elevation as the operational target. Because the operational range for 2011 is less than 3 ft, the outflows from the TFB match the inflows from the power canal very closely. Like Oroville, power is calculated using a relationship of total flow through the power plant and the water surface elevation that was developed from historical data of the form of:

$$P = Ah + BQ + c \tag{4-14}$$

where P, h, and Q are defined above and the coefficients A, B, and C are listed in Table 4-37. A plot of the power function is shown in Figure 4-104.

TABLE 4-37 Power Calculation Coefficients for the Thermalito Forebay

Coefficient	Value
A	6.4356E-03
B	–0.3420
C	75.1282

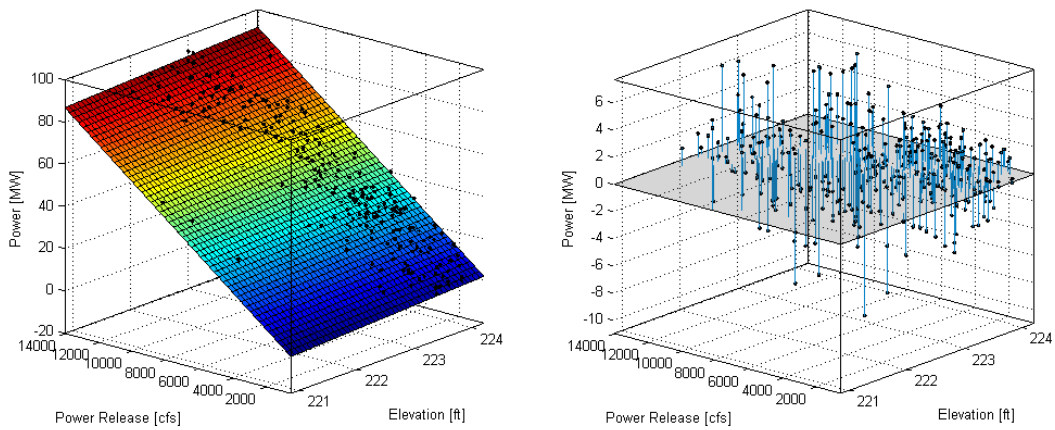

FIGURE 4-104 Plot of the Power Generation Curve for the Thermalito Forebay (The plot on the right shows the residuals between the historical data and the curve fit. The average error is <5%.)

Relationship of TDP and TFB Elevations. In contrast to HydroSCOPE, the CHEERS model fixes the elevations of both the TDP and the TFB to be equal. As mentioned above, the error introduced by this assumption is minimal due to the small storage capacities of the reservoirs in relation to the flow volumes. However, to maintain consistency between the HydroSCOPE results and the CHEERS results, the final HydroSCOPE optimized result is post-processed to equalize the TDP and TFB elevations by changing the flow rate in the power canal. The resulting elevations and flow rates are then written to the database for use by CHEERS.

Thermalito Afterbay. Historical elevations are also used as the target elevations for the TAB but with additional logic to attenuate the corrections toward the target and to ensure flow requirements in the Feather River are met. As mentioned above, the diversions for irrigation are fixed based on historical data. The release logic for the TAB uses a deficit tracking approach that accumulates at each time step the volume difference of water that was actually released and the amount of water that should have been released to keep the water elevation on target. The deficit only accumulates when the reservoir elevation is below the target and when the needed release to meet the target is less than the minimum release needed for downstream requirements. If the reservoir goes above the target, additional releases as a function of the accumulated deficit are

added to the baseline releases. In this manner, the reservoir elevation is more aggressively adjusted downward as the deficit increases. The Thermalito Afterbay does not generate power.

Power Score. The WUOT is intended for implementation on a daily basis by downloading a new forecast from the PNNL server, running and optimizing the seasonal model, and running and optimizing the day-ahead planning model. Both the seasonal model and the day-ahead planning model interact with the IRF tool and use estimates of power generation and/or revenue for their objective functions. For the retrospective analysis on the Aspinall Cascade, average monthly prices were calculated for on-peak and off-peak times, multiplied by the generated power at each time step, and summed at the end of the simulation to provide a revenue objective for optimizing.

However, estimating electricity sale prices for 6 months into the future is difficult and full of uncertainty; thus for the Oroville complex a power score approach was adopted that does not rely on predicting prices. To calculate the power score, the power generation in megawatts is calculated at each time step and then multiplied by a scaling factor of 1 if it is during an off-peak time period, or 1.4 if it is during an on-peak time period. Matching the time step of the seasonal model, the 6-hour time step from midnight to 6:00 am is defined as off-peak while the other three time steps are defined as on-peak. The 1.4 multiplier is based on a weighted average difference between the off-peak and on-peak prices for April through September 2011 used in the day-ahead planning model but skewed downward to account for weekends and holidays, which the seasonal model does not do. Because the shortest timespan given to the optimization multipliers is daily, DAKOTA has little control on optimizing between on-peak and off-peak flows, so the value of the on-peak multiplier is of little significance.

If one were to divide the power score by the number of time steps in a simulation (732), the result can be loosely interpreted to be the average, system-wide rate of power generation in megawatts. Because the rate of power generation is a function of both head and flow as discussed above, the power score can be optimized with respect to operating efficiency while avoiding the uncertainties of trying to predict future sale prices.

Results. The nature of the retrospective analysis makes interpreting the results of the seasonal tool somewhat difficult in that when used in this retrospective capacity, the real value of the seasonal model lies in how it influences the day-ahead planning model and the improvements it brings there (those improvements are discussed in the day-ahead results section below). Thus, the results presented here focus on the ability of the optimization process to improve upon the default scenario with respect to the IRF and power scores.

Figure 4-105 is a set of plots showing the range of variability of the total releases from Oroville that are sampled by DAKOTA during the optimization process. Each plot represents a statistical summary of the 500 simulations conducted for each optimization with the green bands defining the maximum and minimum difference in the releases as compared to the default release, the blue band showing the 25th to 75th percentiles, the purple line showing the mean, and the red line showing the optimized release. In this context, the default release schedule

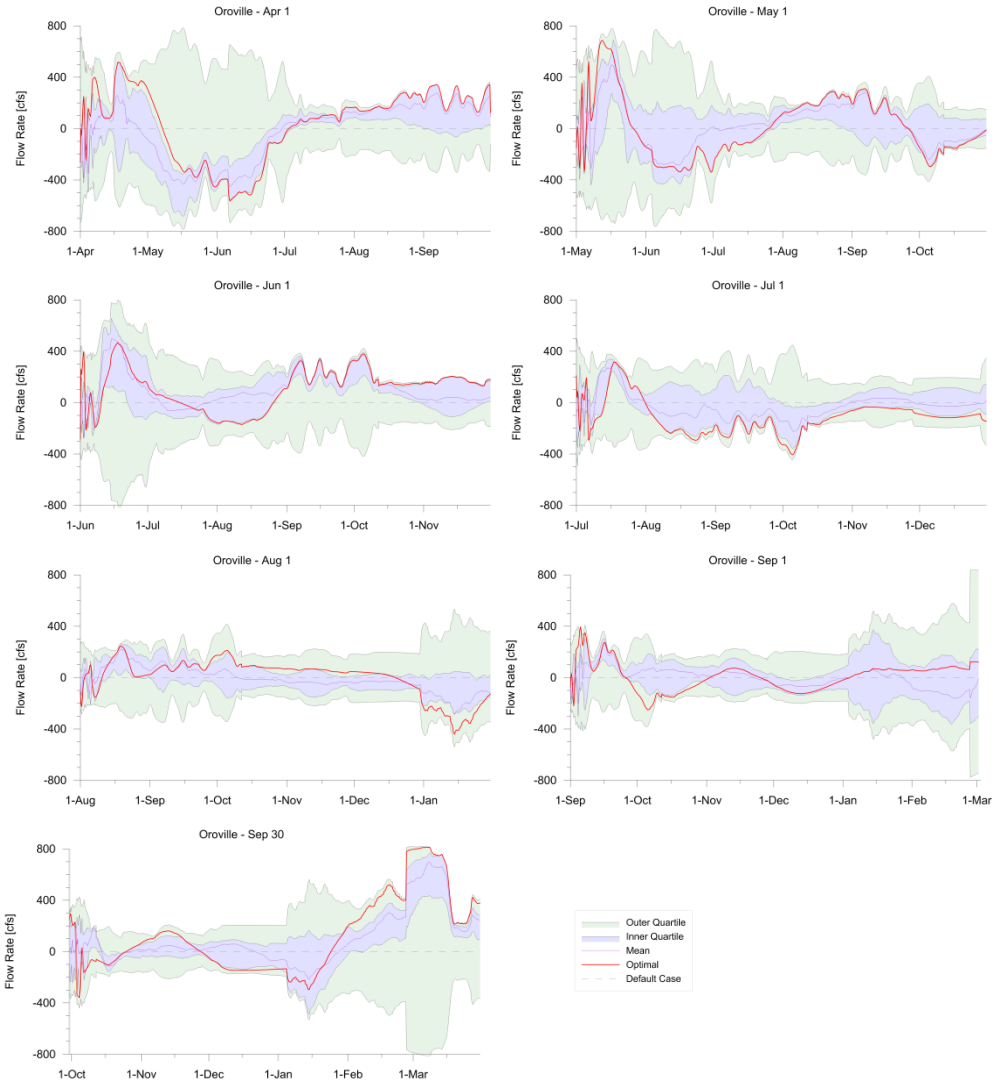

FIGURE 4-105 Plots Showing the Range of Variability of the Total Releases from Oroville Sampled by DAKOTA during the Optimization Process (The Flow Rate on the y-axis represents the difference between the particular statistic [maximum, minimum, mean, optimal, etc.] and the default discharge. Note how the optimal release varies between the maximum and minimum over time as well as for different start dates.)

would be a horizontal line at the 0 cfs mark on the y-axis. Note that the date for each plot is the start date of the optimization, which is conducted over the next 6 months from the start date.

In all cases, DAKOTA is sampling values above and below the default release schedule, as determined by the green bands, but ends up simulating the majority of runs above or below the default line. Of interest is the fact that the optimal release schedules vary between the maximum and minimum values within the individual plots. In addition, the optimum release as compared to the default value for a particular time of year can be different depending on the start date of the optimization. For example, the optimal release on September 1 as determined by the April 1, May 1, and June 1 optimizations is approximately 400 cfs above their respective default cases but then falls to 200 cfs below its default case for the July 1 optimization. The reason for this is mainly due to the differences in the forecasted inflows for each optimization date.

Figure 4-106 shows the tradeoff between the power score and the IRF score for selected dates of the analysis. The yellow diamonds show the scores of the optimal solution as determined through the lowest common rank approach discussed above. The green baseline point is the value of the default release schedule, which is the release schedule calculated by the model for the forecasted inflows and is the starting point of the optimization process. The purple diamonds represent the points on the Pareto front, which defines the line where improvement in one metric is only possible with a decline in the other metric. Note that for the July 1 optimization that the power score for the default case is higher than the optimal case. This can occur when the default case is on or close to the Pareto front and when improvement in one metric is greater than the corresponding decline in the other metric. In this case, the 0.4% decline in the power score is offset by a 6.9% increase in the IRF score. In practice, one would look closely at the causes behind those changes and decide whether the calculated optimal solution is really optimal or not. Table 4-38 lists the numerical values and percent change in the optimal and baseline solutions.

Figure 4-107 is a plot of the 7-day moving average of the difference between the historical temperature at the Robinson's Riffle gauge and that simulated by the model. The 7-day moving average smooths out the time step–by–time step variability, which makes the difference easier to see. On average, the simulated temperatures are approximately 0.05°F less than the historical data, which is well within the margin of error for the model and thus should not be considered a significant or accurate difference. The average temperature difference is slight, due to the IRF score calculation, which reduces the score only for simulations where the time out of compliance exceeds the historical case. The fact that the temperature difference between the historical and optimized simulations is small may be an indicator that the temperature performance of the system is near optimal. More work would be needed to determine if in fact that is the case; this would include optimizing the shutter positions to better meet the downstream temperature and/or increase the systems margin of error.

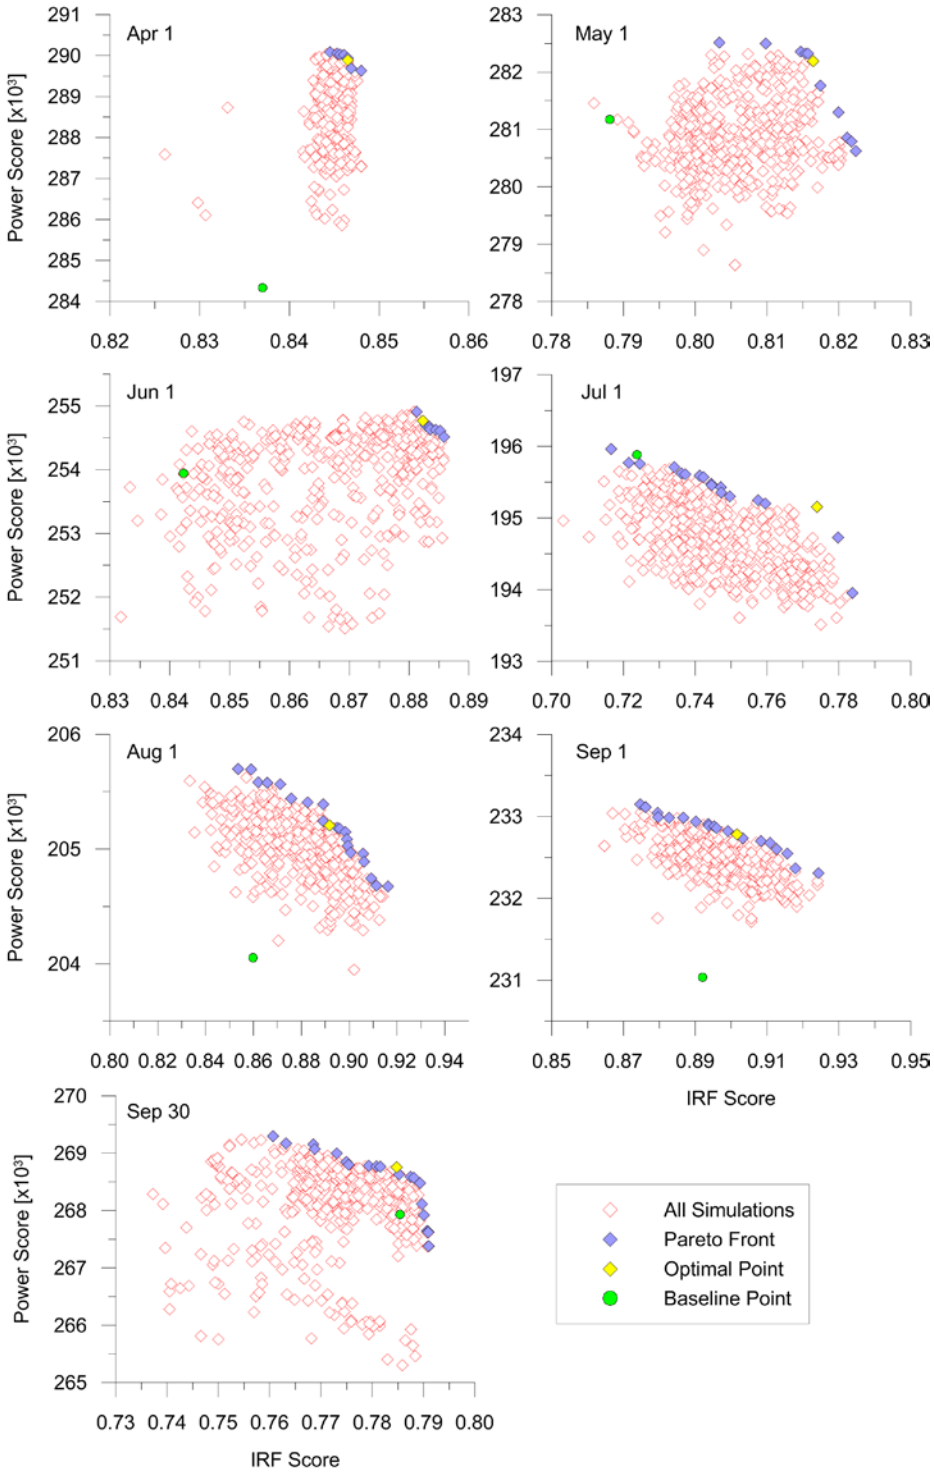

FIGURE 4-106 Plots of the Optimization Results for Selected Dates of the Analysis (The Pareto front defines the line where improvement in one metric results in a decline in the other metric. The yellow diamond represents the score of the optimum solution as determined with a lowest-common rank approach. The baseline point represents the relative scores of the default release schedule.)

TABLE 4-38 Power and IRF Scores and the Percent Improvement for Selected Dates of the Analysis Period[a]

	Power Score		IRF Score		Percent Improvement	
Date	Optimized	Baseline	Optimized	Baseline	Power Score	IRF Score
4/1/2011	289890.59	284326.00	0.846	0.837	1.96%	1.13%
5/1/2011	282191.96	281173.55	0.816	0.788	0.36%	3.60%
6/1/2011	254766.53	253941.65	0.882	0.842	0.32%	4.75%
7/1/2011	195158.55	195886.01	0.774	0.724	–0.37%	6.93%
8/1/2011	205204.53	204051.57	0.892	0.860	0.57%	3.73%
9/1/2011	232779.54	231033.49	0.902	0.892	0.76%	1.07%
9/30/2011	268751.89	267930.87	0.785	0.785	0.31%	–0.08%
Average	246963.37	245477.59	0.842	0.818	0.61%	2.95%

[a] Note that the IRF score and total revenue are for the 6-month period following and including the listed date. The average is for the selected dates only.

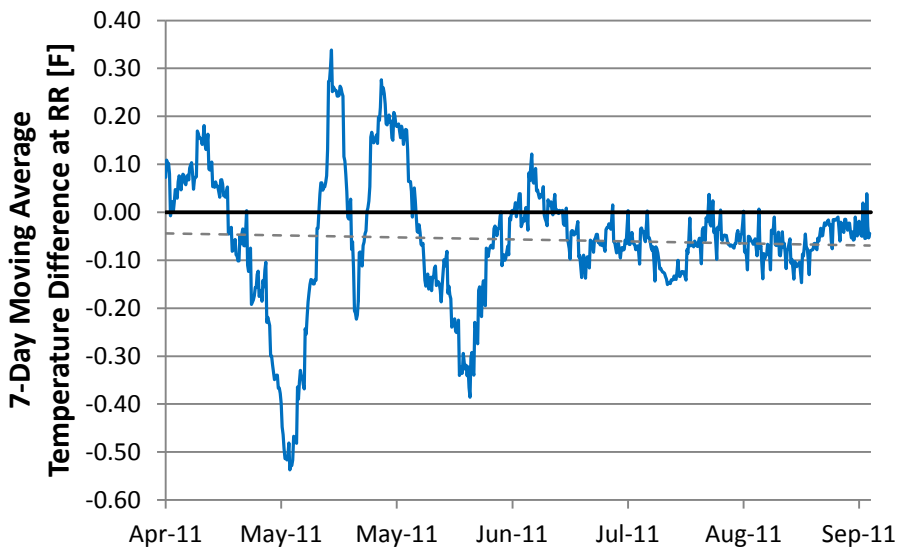

FIGURE 4-107 The Difference between the Optimized Temperature and the Historical Temperature at Robinson's Riffle (The dotted line is the average difference over time. The fact that the differences are slight may be an indication that the temperature performance of the system is near optimal. More work and study is needed to determine if that is actually the case.)

To examine how the day-by-day optimizations improved system performance, the first day of the 183 optimizations are concatenated to form a time series of the reservoir outflows for the April 1 through September 30 analysis period. The resulting release schedule is used as input to the model, which is then run a single time to calculate the power and IRF scores. Those results, which are called the daily results, are compared to another simulation that uses the historical release schedule (called the WB Results; WB stands for water balance) in calculating the scores. Both simulations use the historical, realized inflows to the system. The daily results release schedule is the release schedule that would have been realized had the seasonal tool been used on a daily basis to optimize the outflows and the power, and IRF scores are metrics that describe the integrated performance over the 6-month time period. While the daily release schedule will be further refined with the day-ahead planning model, the comparison provides insight as to the type of improvement that may be possible if only a seasonal analysis is used to optimize the system.

Figure 4-108 shows a plot of the Oroville releases and elevations for the daily and WB results. The daily results deviate from the WB results around mid-April when the operations logic increases the releases to keep the water surface elevation below the low-flow elevation target (see Figure 4-100). The daily results releases become less than the WB releases from late-April to mid-May as the operations logic tries to bring the water surface elevation up toward the elevation target, which it never can do. The result is that the reservoir never becomes completely full, topping out at 897.36 ft on June 26 (capacity is at 900 ft). From mid-May onward, the daily release flows closely mimic the WB releases due to the need to meet the various downstream flow targets. This indicates that during the summer months, there is minimal ability to optimize the system due to the lack of flexibility in the flow target regulations. One possible solution to this issue might be to optimize the system during the winter and spring to maximize the system flexibility during the summer and then optimize during the summer to maximize the power and IRF scores. A comparison of the daily release volumes for Oroville is shown in Figure 4-109, which clearly shows the differences in the timing of the releases between the daily results and the WB releases.

As a final comparison, Table 4-39 lists the power and IRF scores for the daily results and WB results simulations. The IRF scores are virtually identical, with the daily results release schedule underperforming the WB release schedule by just 0.02%. However, the power score for the daily release schedule is 3.67% higher than the WB schedule, indicating that an upside may be available in optimizing for power generation. As for the temperature comparisons, more work is needed to determine this more clearly.

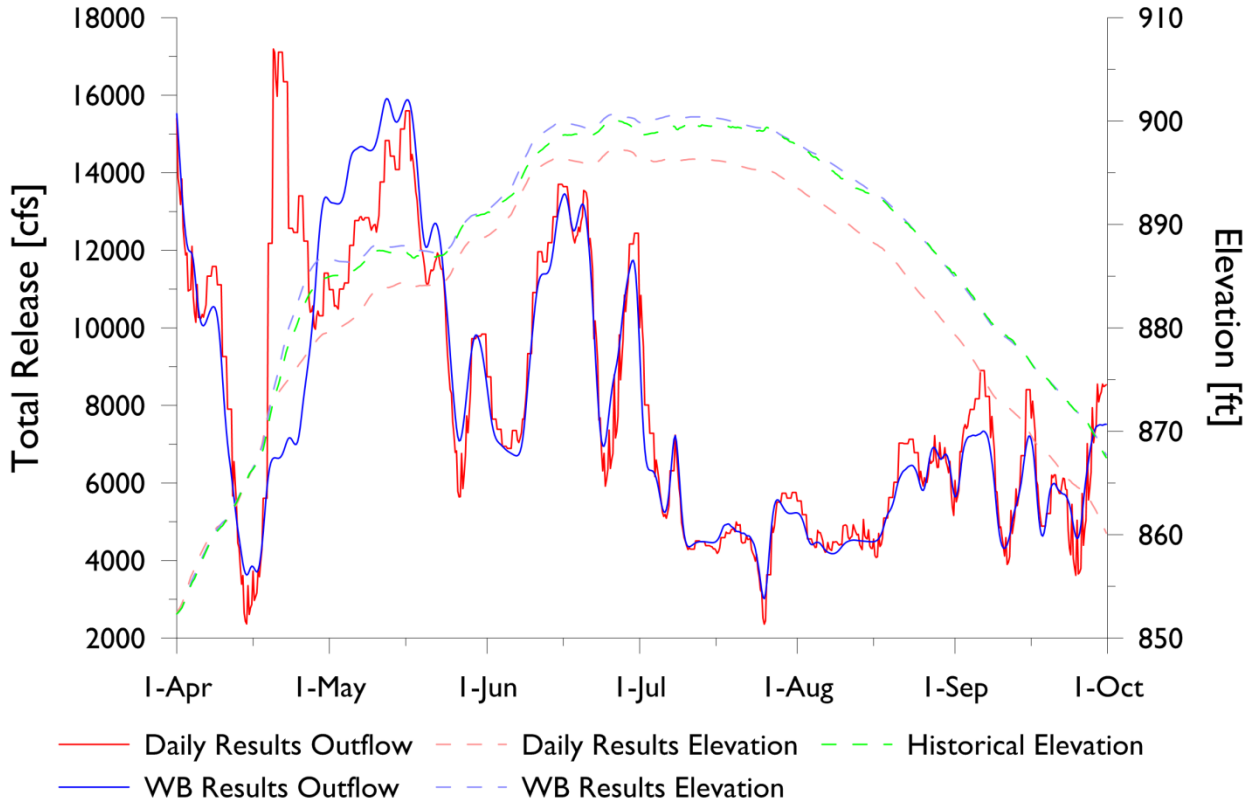

FIGURE 4-108 Plot of the Daily Result Releases, the Water Balance (WB) Releases, and the Corresponding Water Surface Elevations for Oroville (The model logic increases the daily result releases over the historical rates around mid-April to keep the water surface elevation below the low-flow elevation target.)

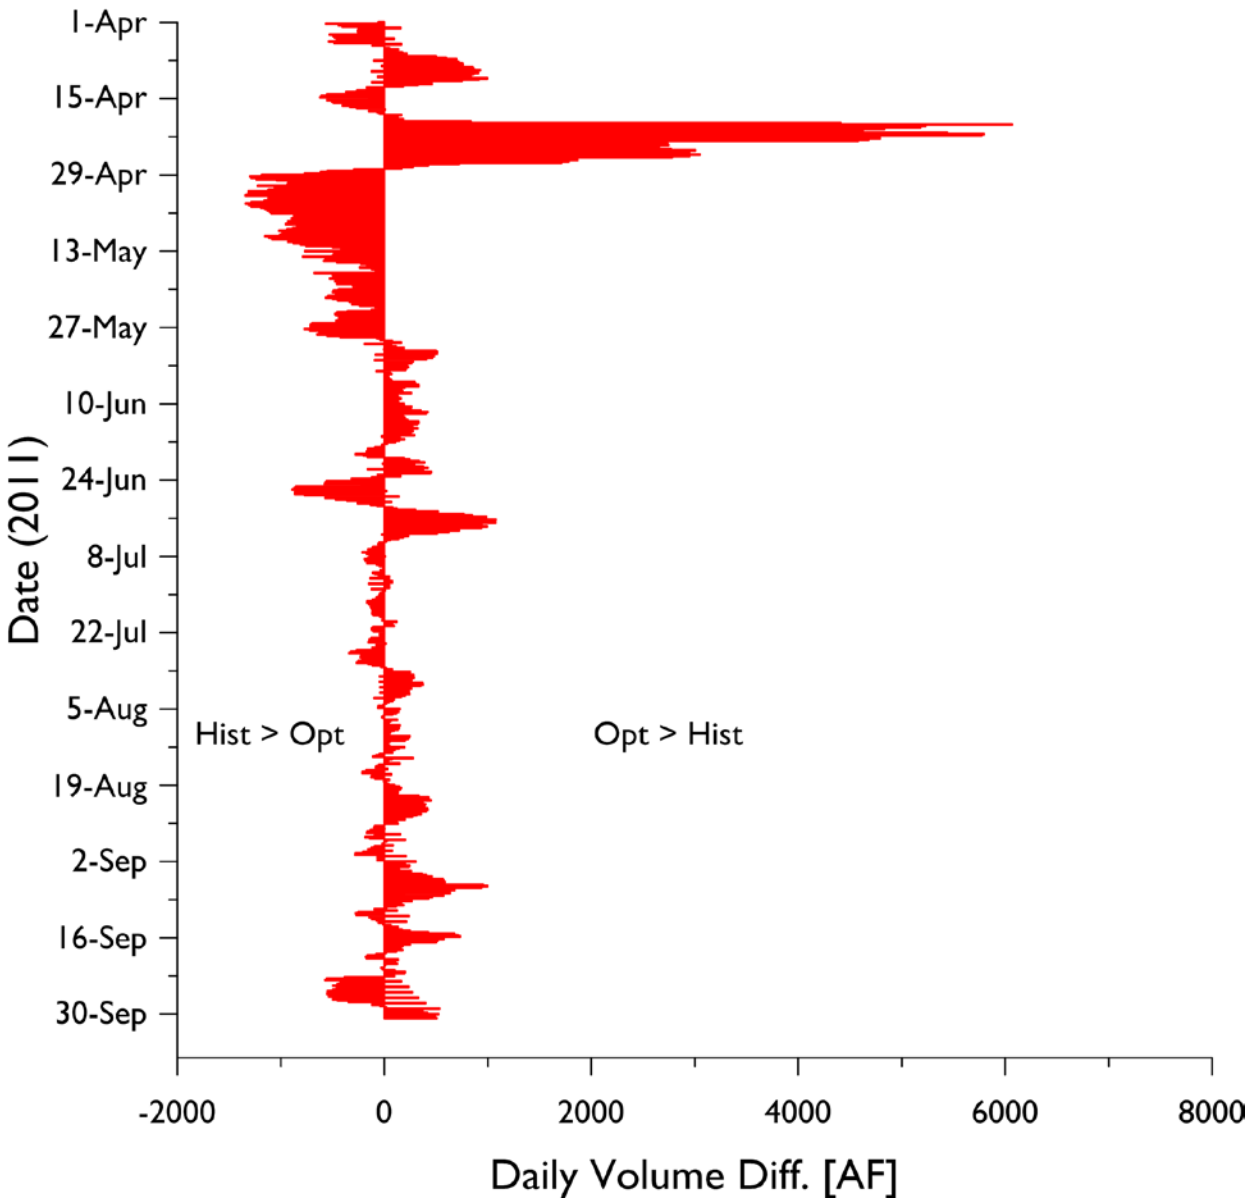

FIGURE 4-109 Plot of the Daily Volume Difference for Releases from Oroville of the Daily Results (Opt) and the WB Results (Hist) (The pattern clearly shows the timing difference in the releases between the two scenarios.)

TABLE 4-39 Power and IRF Scores of the Daily Results and Water Balance Release Schedules

Power Score		IRF Score		Percent Improvement	
Daily Results	WB Results	Daily Results	WB Results	Power Score	IRF Score
291901.16	281560.37	0.849	0.849	3.67%	–0.02%

Summary. To examine how the WUOT might perform in actual operations, a retrospective analysis for the Oroville complex was conducted by successively optimizing daily operations from April 1, 2011, through September 30, 2011. For the seasonal model, this involves performing a 6-month optimization for each day and then supplying the daily release volumes for the first day of each optimization to the day-ahead planning model, which then optimizes those releases over the next 24 hours. When used in this capacity, the seasonal model does not directly influence system operations, which makes comparing the seasonally optimized system to the historical performance difficult.

However, as shown above, results from the seasonal model do provide qualitative insight into the systems' operations that could help shape regulatory change and/or operational logic in order to improve the flexibility and resiliency of the system. In summary, those insights are as follows:

1. Direct comparison of the daily optimized releases to the historical releases show no improvement in the IRF score and about a 3.7% improvement in the power score.

2. Little to no improvement was made to the temperature performance of the system, indicating that changes in the releases within the parameters of the optimization have little influence on temperature. Improvements in temperature performance may only be possible by optimizing both the reservoir releases and the Oroville intake shutter positions.

3. Regulatory constraints in the form of downstream flow targets limit the ability to improve performance during the summer months. Optimizing the system during the winter and spring to maximize the systems flexibility during the summer months and then optimizing the summer months to maximize the power and IRF scores may improve this.

Future analysis and work for the seasonal model should focus on expanding the optimization objective function to include metrics such as system resiliency and minimizing "regret" (regret, as used here, is the difference between the projected performance and the realized performance). Improvements in the model's operational logic to possibly include decision logic for determining the shutter positions should also be explored.

4.4.2.3 Environmental Performance

Description of Environmental Objectives for the Feather River. Objectives were defined for the following two locations on the Feather River:

1. Gridley, California (meant to represent the high flow channel [HFC] reach between the outlet of Thermalito Afterbay and the Honcut Creek).

2. Robinson Riffle, located below the Thermalito Diversion Dam (meant to represent river reach in the low-flow channel [LFC]).

Apart from these two locations, environmental rules in the form of facility-mandated hard constraints were also applied for two other locations:

1. Thermalito Afterbay Reservoir (TAR) (meant to represent the shoreline habitat around the reservoir). Environmental performance evaluation at this location is to comply with reservoir elevation operation rules (rather than discharge objectives).

2. Oroville Fish Hatchery (OFH) (meant to represent the water temperature requirements for the fish hatchery). Environmental performance evaluation at this location is to facilitate respecting the water temperature agreement between Oroville Facility and the California Department of Fish and Game (CDWR 1983).

The IRF tool was used in Planning mode for the integrated run with the other toolset components of WUOT to optimize water use during the retrospective demonstration. The integrated retrospective runs were conducted for the second half of water-year 2011 (April–September 2011), and included consideration of environmental objectives and/or rules for all four locations mentioned above. It should also be noted that while for the TAR and OFH sites only environmental rules (constraints) are identified, there are also some hard environmental rules for the LFC and HFC sites. Besides the objectives and constraints mentioned below for each of the four locations, there are additional flow restrictions or operational constraints related to flooding at the confluence with the Yuba and Sacramento Rivers; however, those are not defined here because the HFC site represents the downstream boundary of our model network.

Environmental performance analysis using the standalone IRF-Historic mode was done only for the HFC site, due to unavailability of long-term (>25 years) historic data for the other three sites. As a result, objective and IRF scores in the IRF Historic mode are indicative of the overall and objective performance exclusively for the HFC site. The flow and thermal regimes of the LFC site have been significantly modified compared to their pre-regulated state due to operations of the Oroville-Thermalito Complex. Since most of the water that would have naturally flowed through the LFC site is bypassed through a complex channel of canals and reservoirs, comparing historic environmental performance to present performance would make little practical sense for this 8-mi reach. The TAR and OFH sites are new features on the Feather River that were nonexistent in the pre-regulated period. As a result, temporal comparison of performance from a historical perspective also makes little sense for these sites. Moreover, the TAR and OFH sites contain only rules or hard constraints rather than objectives having continuous functions, with the possibility of incremental improvement in the condition of objectives. While discharge, water temperature, and reservoir elevation data were provided by the California Department of Water Resources Oroville Facility, river stage was determined from rating curves developed in this study. Most of the raw data for discharge stage and volume elevation curves were also provided by the facility staff.

Environmental Objectives for the Feather River Location 1—High Flow Reach.

Objective 1.1: Facilitate Chinook Salmon Spawning
Relevant Period: October 15–November 30

Condition: The purpose of this objective is to provide optimal spawning conditions for spring- and fall-run chinook salmon by regulating the depth in the main channel. For this purpose, a generalized rating curve was used that describes the relationship between discharge and stage (depth of water column) for the HFC segment (see next section). Based on published studies conducted for this salmonid species, optimum depth was determined to range between 1.6 and 2.6 ft; the minimum and maximum acceptable depths were 0.4 ft and 4 ft, respectively (Sommer et al. 2001). In order to meet this objective, a minimum of 1000 points would be required, where a point (or a fraction of a point) can be accrued at each instance (data recorded at the 15-minute timestep) based on the function described in Figure 4-110, for all instances during the relevant period. However, a total of 2000 points (4512 instances are available during the mentioned period) would accomplish the objective to its maximum potential. Maintenance of water level within a 1-ft band to obtain the highest possible point also provides stable incubation conditions and minimizes chances of redd exposure. The recommended timing is based on peak spawning activity and considers annual variations (Sommer et al. 2001; CDWR 2003). In general, the optimal depth range for spawning should not be interpreted as merely an on-site necessity for spawning; this depth is also necessary for facilitating migration up to the natal

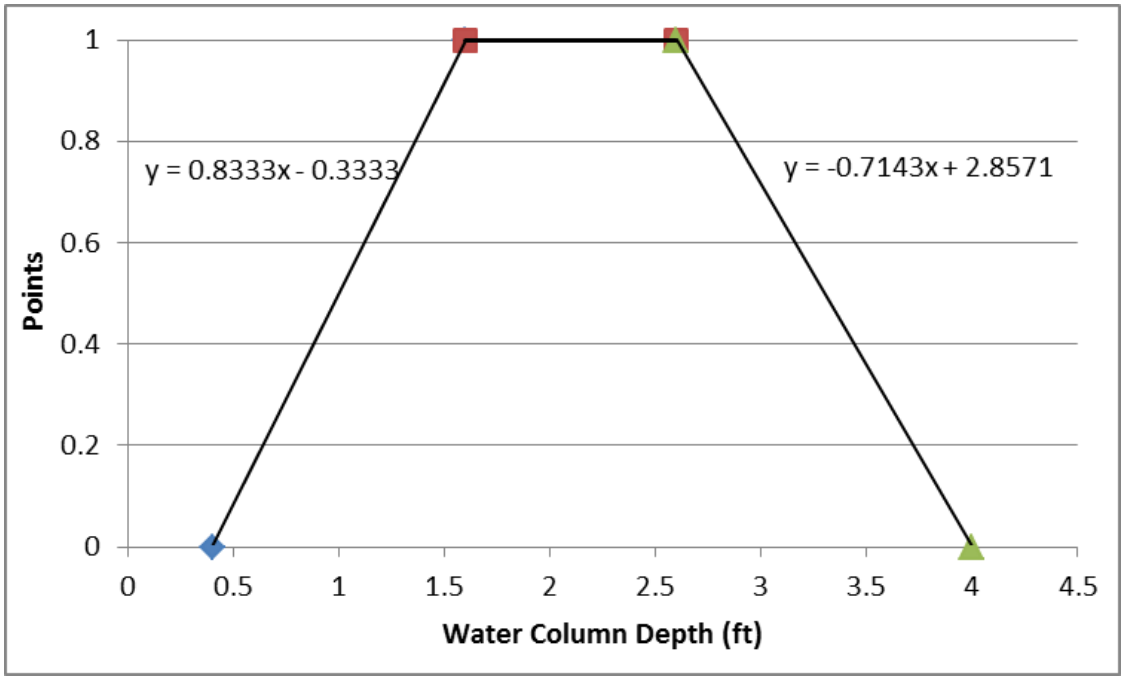

FIGURE 4-110 Relationship between Instantaneous Depth of Water Column and Suitability Points for Chinook Salmon Spawning in the High Flow Channel Section of the Feather River

creeks and tributaries. As defined, this objective is desired to be met in 9 out of 10 years on a long-term average, but after 2 consecutive years of failure to address it, its score would go to zero.

Objective 1.2: Facilitate Steelhead Trout Spawning
Relevant Period: December 1–February 28

Condition: This objective is intended to provide optimal depth for spawning steelhead trout by regulating the depth in the main channel. The rating curve for the HFC segment that was used in the previous example was also used here. Based on published studies conducted for this salmonid species, a suitable depth range was determined to be between 0.4 and 2.3 ft, although the optimum depth was deemed to be around 1.2 ft (CDWR 2003, 2004b). A minimum of 2000 points would be required for this objective where a point (or a fraction of a point) can be accrued at each instance (data recorded at the 15-minute time step) based on Figure 4-111 for all instances during the relevant period. However, a total of 4000 points (5664 instances are available during this period) would accomplish the objective to its maximum potential. A somewhat stable water level is also expected to provide ideal incubation conditions and minimizes chances of redd desiccation. As defined, this objective is desired to be met in 9 out of 10 years on a long-term average, but after 2 consecutive years of failure to address it, its score would go to zero.

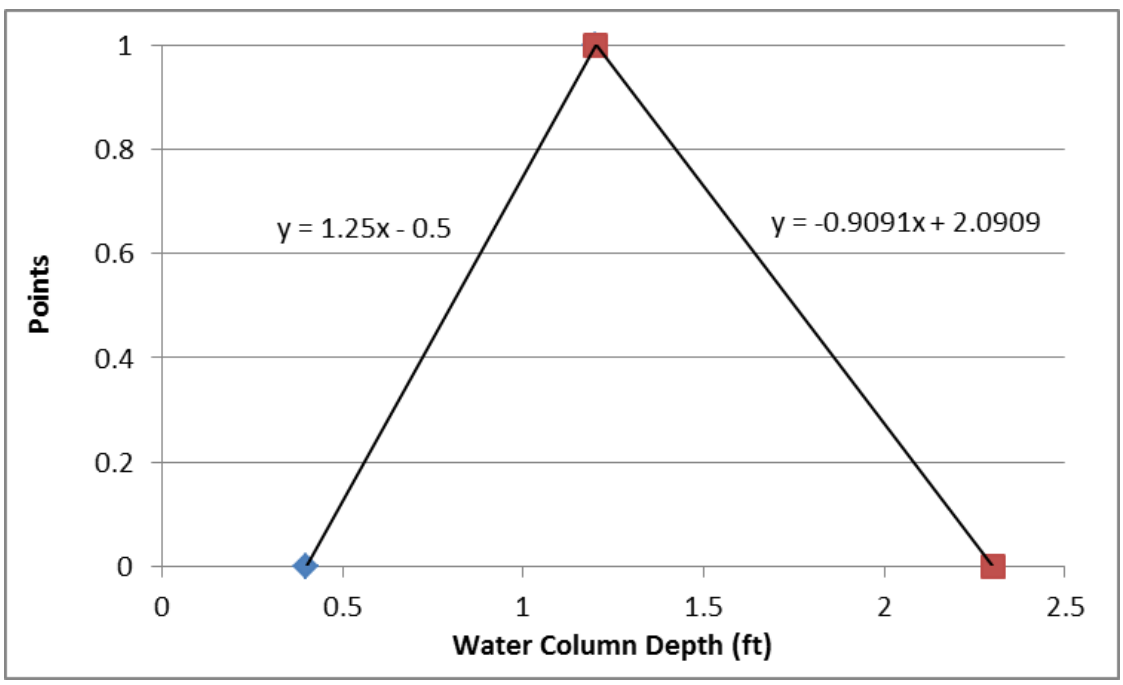

FIGURE 4-111 Relationship between Instantaneous Depth of Water Column and Suitability Points for Steelhead Trout Spawning in the High Flow Channel Section of the Feather River

Objective 1.3: Spring Pulse Flow for Out-Migrating Chinook Salmon for Below-Normal Years
Relevant Period: March 15–April 30

Condition: This objective is to ensure that out-migration of chinook salmon juveniles, primarily the sub-yearlings, occurs to some extent, except during very dry years. Outmigration in this salmonid species usually occurs over a 60–90 day duration (Sommer et al. 2001; CDWR 2003, 2004a). The spring pulse flow enables out-migration of salmon parrs downstream, and translocates developing juveniles to the estuary, where physiological (primarily osmoregulatory) adaptations associated with smoltification occur. Moreover, the sudden pulse of freshwater also flushes silt deposits from interstitial spaces in gravel bed, as well as triggering an inland spawning migration of coastal and estuarine adult populations of the spring-run chinook salmon. To achieve this objective to its minimum degree, maximum daily discharge should be at least 8000 cfs for 30 days within the relevant period; 45 days at or above the threshold accomplishes the objective to its full extent. This spring pulse flow magnitude and duration was recommended in a feasibility study related to reoperation of Oroville Reservoir to improve environmental flows in the Feather River, which indicated the threshold for hydrologically below-normal years (CDWR and BOR 2012). As defined, this objective is desired to be met in 8 out of 10 years on a long-term average, but after 4 consecutive years of failure to address it, its score would go to zero.

Objective 1.4: Spring Pulse Flow for Out-migrating Chinook Salmon for Above-normal Years
Relevant Period: March 15–April 30

Condition: This objective supplements the previous objective (objective 3.3) for hydrologically below-normal years, and is defined to further enhance out-migration of chinook salmon juveniles, especially during the wetter years when abundant inflow enters into the Oroville-Thermalito Complex. As a result, the definition of this objective is same as for objective 3.3, except for a higher threshold and lower frequency. To achieve this objective to its minimum degree, maximum daily discharge should be at least 10,000 cfs for 30 days within the relevant period; 45 days at or above this threshold accomplishes the objective to its full extent. This spring pulse flow magnitude was recommended to improve environmental flows in the Feather River, which indicated the threshold for hydrologically above-normal years (CDWR and BOR 2012). As defined, this objective is desired to be met in 1 out of 4 years on a long-term average; the score for this objective decreases as the observed frequency falls further behind the desired frequency.

Objective 1.5: Bed Mobilization to Maintain Channel Integrity
Relevant Period: February 20–March 20

Condition: This objective is achieved by exceeding the bank-full flow during peak spring flow period and is intended to maintain channel complexity in the HFC reach of the Feather River, especially in the lower reaches of the Oroville Wildlife Area. The resulting spike in flow is intended to perform many geomorphic functions that allow reshaping of the channel, such as meander migration, scouring sediment deposits, preventing vegetation establishment and

encroachment on gravel bars, and reconnection of the floodplain to the main channel. This objective would be accomplished by maintaining a minimum hourly flow of 35,000 cfs for a consecutive 12-hour period anytime during the relevant period. This objective is desired to be met in 1 out of 4 years on a long-term average; the score gets smaller as the objective's observed frequency falls further behind the desired frequency. The threshold bank-full discharge magnitude, timing, and duration and the recommended frequency are based on the feasibility study related to the reoperation of Oroville Reservoir to improve environmental flows in the Feather River (CDWR and BOR 2012).

Objective 1.6: Fremont Cottonwood Riparian Seedbed Preparation
Relevant Period: April 15–May 15

Condition: This objective is intended to provide flows of a certain magnitude to maintain groundwater level in the adjacent riparian floodplain for a duration sufficient for Fremont cottonwood seed germination. This action also results in creating point-bars suitable for cottonwood seedbed preparation (Trush et al. 2000). It has been recommended that seedbed preparation be followed by bed mobilization that creates unvegetated nursery sites through erosional and depositional processes. Barren nursery sites without herbaceous vegetation allow abundant sunshine needed for the low-endosperm-containing seeds and provide new seedbeds for this pioneer riparian species (Fenner et al. 1984; Mahoney and Rood 1998). This objective would be accomplished by maintaining a minimum daily flow of 12,500 cfs for a minimum of one consecutive 4-day period anytime during the relevant period; a consecutive 7-day flow at the mentioned threshold fully accomplishes the objective. The relatively larger magnitude of discharge would keep the soil moist in significant portions of the floodplain during the peak seed-release period, which lasts approximately 3 weeks. Moreover, the specified magnitude maintains the desired adjacent groundwater level, which is critical for the establishment of seedlings in the new seedbed. This objective is desired to be met in 1 out of 5 years on a long-term average; the score decreases as the objective's observed frequency falls behind the desired frequency. The duration of 4–7 days is based on sprouting time and the development of an initial root system, and the timing is based on the peak seed release and dispersal, which usually occurs between April and June, and their viability periods (CDWR and BOR 2012). The desired frequency is based on historic pre-regulated flow observations, which indicated that this objective was met on average once every 5–10 years, during moderate flood events (McBain and Trush 2002; Trush et al. 2000).

Objective 1.7: Fremont Cottonwood Stand Establishment and Recruitment
Relevant Period: May 1–June 30

Condition: This objective is intended to continue from the previous seedbed development objective by maintaining ramp rate such that the newly initiated cottonwood seedlings may become established as mature stands and allow future recruitment in this species. This objective is accomplished by maintaining a maximum within-day elevation change of 0.083 ft (1 in.) for 60 consecutive days during the relevant period. Cottonwood recruitment models recommend that post-germination decline in river stage not to exceed 1 in./day, because that allows the growing roots to maintain contact with the capillary fringe of the receding water table in the riparian substrate (Busch et al. 1992; Mahoney and Rood 1998). The consecutive days prevent any

interruption in capillary contact for the developing root system and allow the seedlings to grow at the same rate as the water-level recedence rate during the stage at which they are most vulnerable to water stress. The 60-day duration is based on the plants' overall growth rate, as recommended in the feasibility study related to reoperation of Oroville Reservoir to improve environmental flows in the Feather River (CDWR and BOR 2012). This objective has the same desired frequency as the seedbed objective (i.e., 1 out of 5 years on a long-term average), and should be targeted in conjunction with and sequenced in the mentioned order. By the end of first year, cottonwoods generally develop a well-established root and shoot system to sustain them through seasonal conditions of drought and flooding, and beyond 2 years plants are somewhat less susceptible to drought stress as long as groundwater levels stay within 10–20 ft of the wooded riparian surface (CDWR and BOR 2012).

Objective 1.8: Pulse Flow for Out-migrating Fall-Run Chinook Salmon
Relevant Period: February 1–February 28

Condition: This objective is to allow out-migration of chinook salmon juveniles, primarily the early hatchlings from spawning of the fall run. The timing is based on the peak emigration activity observed for this run-type of Chinook (Sommer et al. 2001; CDWR 2003, 2004a). The shorter-duration pulse flow enables out-migration of these ocean-type individuals, and thus reduces completion over limited resources for the next wave of hatchlings from the spring run, which are usually more abundant and are considered river-type. Since ocean-type Chinooks live in the river only for 3–4 months, unlike their ocean-type counterparts that may live in the river for 12–15 months, spatial and temporal separation in habitat possibly allows for better growth rate and overall stronger year-class. For achieving this objective to its minimum degree, maximum daily discharge should be at least 8000 cfs for 5 days within the relevant period; 7 days at or above this threshold accomplishes the objective to its full extent. A shorter duration compared to the spring emigration is based on the lower numbers of fall-run salmon. The pulse flow magnitude is based on the spring-pulse threshold for hydrologically below-normal years (CDWR and BOR 2012). This objective is desired to be met in 8 out of 10 years on a long-term average.

Environmental Objectives for the Feather River Location 2—Low Flow Reach.

Objective 2.1: Facilitate Chinook Salmon Spawning
Relevant Period: October 15–November 30

Condition: The intention of this objective is to provide optimal spawning conditions for the spring- and fall-run Chinook salmon by regulating the depth in the LFC. Since operation of the Oroville Dam began, a greater portion of Chinook spawning is known to occur in this section of the river, compared to the HFC section (Sommer et al. 2001). However, further migration upstream through the LFC is blocked by the fish-diversion dam, and the fish are rerouted through the hatchery. For the LFC section, we used the same optimal depth for Chinook salmon as in the HFC section described in the earlier section (i.e., acceptable depths between 0.4 and 4 ft, and preferred depths between 1.6 and 2.6 ft). It should be noted that change in river cross-sections between the HFC and LFC locations translates to differences in discharge levels at the two sites,

even if the specified magnitude of depth is same for the two sites. A generalized rating curve describing the relationship between discharge and stage (depth of water column) for the LFC segment was used (see next section). In order to meet this objective, a minimum of 1000 points would be required, where a point (or a fraction of a point) can be accrued at each instance (data recorded at the 15-minute time step) based on the function described in Figure 4-110, for all instances during the relevant period. However, a total of 2000 points (4512 instances are available during this period) would accomplish the objective to its maximum potential. Maintenance of a somewhat stable water level during the period of peak spawning minimizes chances of redd scour and exposure. As defined, this objective is desired to be met in 9 out of 10 years on a long-term average, but after 2 consecutive years of failure to address it, its score would go to zero.

Objective 2.2: Facilitate Steelhead Trout Spawning
Relevant Period: December 1–February 28

Condition: This objective is meant to provide optimal depth for spawning steelhead trout by regulating the depth in the LFC. The rating curve for the LFC segment that was used in the above example was also used here. For this LFC section, we used the same optimal depth for steelhead trout as in the HFC section described previously (i.e., acceptable depths between 0.4 and 2.3 ft, and preferred depth of 1.2 ft). A minimum of 2000 points would be required for this objective, where a point (or a fraction of a point) can be accrued at each instance (data recorded at the 15-minute time step) based on Figure 4-111 for all instances during the relevant period. However, a total of 4000 points (5664 instances are available during this period) would accomplish the objective to its maximum potential. Water level stability would provide ideal incubation conditions and minimize chances of redd scour and desiccation. As defined, this objective is desired to be met in 9 out of 10 years on a long-term average, but after 2 consecutive years of failure to address it, its score would go to zero.

Objective 2.3: Maintain Necessary Thermal Conditions for Early Life-Stages of Chinook Salmon
Relevant Period: September 1–December 31

Condition: The intention of this objective is to provide the necessary water temperature conditions for spawning, egg incubation, and alevin development of Chinook salmon by regulating the temperature of the water released through the LFC. Managing thermal regime in the LFC section of the river is important because the majority (up to three-quarters) of Chinook spawning is known to occur in this section of the Feather River (Sommer et al. 2001). Water temperature in the LFC is primarily affected by the intake strata from Lake Oroville, the volumes involved, and the ambient condition and residence time in the Thermalito Diversion Pool. Water temperature issues can sometimes occur in early fall, especially during periods of high ambient temperature (>80°F) in the discharges from the shallow diversion pool. Water temperature in the LFC can thus be regulated to some extent by managing the hypolimnetic releases from the Oroville Dam and its residence time in the diversion pool. Recommended ideal water temperature of 60°F is the upper limit for egg development, fry hatching, and alevin growth and emergence in this salmonid species. In order to meet this objective, a minimum of 90 points would be required, where a point can be accrued daily during the relevant period for which the

mean-daily water temperature is less than or equal to 60°F. However, a total of 120 points would accomplish the objective to its maximum potential. This objective is defined only using the upper thermal limit because the critical lower thermal limit of 38°F is usually not encountered and thus not a concern during typical climate and operation conditions. As defined, this objective is desired to be met in 9 out of 10 years on a long-term average, but after 2 consecutive years of failure to address it, its score would go to zero.

In addition to the objectives defined for the two locations, the following operation rules or constraints were placed:

Constraints for the Feather River HFC site:

1. Minimum release rule from the Oroville-Thermalito facility:

Scenario 1: If the total April through July unimpaired inflow into Oroville Reservoir during the most recently concluded April 1 to July 31 time period is more than 1,068,100 AF, and if Oroville Reservoir's predicted elevation is more than 733 ft by the upcoming September 30, then the following rule applies:

Time Period	Minimum Flow Requirements (cfs)
October through February	1700
March	1700
April through September	1000

Scenario 2: If the total April through July unimpaired inflow into Oroville Reservoir during the most recently concluded April 1 to July 31 time period is more than 1,068,100 AF, and if Oroville Reservoir's predicted elevation is less than 733 ft by the upcoming September 30, then the following rule applies:

Time Period	Minimum Flow Requirements (cfs)
October through February	1275
March	1275
April through September	750

Scenario 3: If the total April through July unimpaired inflow into Oroville Reservoir during the most recently concluded April 1 to July 31 time period is less than 1,068,100 AF, and if Oroville Reservoir's predicted elevation is more than 733 ft by the upcoming September 30, then the following rule applies:

Time Period	Minimum Flow Requirements (cfs)
October through February	1200
March through September	1000

Scenario 4: If the total April through July unimpaired inflow into Oroville Reservoir during the most recently concluded April 1 to July 31 time period is less than 1,068,100

AF, and if Oroville Reservoir's predicted elevation is less than 733 ft by the upcoming September 30, then the following rule applies:

Time Period	**Minimum Flow Requirements (cfs)**
October through February	900
March through September	750

Minimum flow requirements are based on minimum instantaneous values recorded at Gridley. The unimpaired April through July 2010 flow was 2,030,706 AF, and the observed April through July 2011 inflow was 3,405,831 AF; the predicted reservoir elevation for both years (i.e., 2011 and 2012) was above 733 ft. Therefore, scenario 1 was adopted for the modeling period of April–September 2011, and the 6-month advancing period of the tool (i.e., October 2011–March 2012).

2. Maximum release rule from the Oroville-Thermalito facility:

   If the total combined discharge from the facility (i.e., LFC + HFC) is more than 2500 cfs anytime between October 15 and November 30, then the minimum flow through March 31 needs to be maintained at 500 cfs below the maximum observed discharge from October 15 to November 30. This constraint was put in place to ensure that if salmon spawning occurred at levels exceeding 2500 cfs, appropriate conditions are maintained for the redds to prevent exposure and provide suitable conditions for developing juveniles, at least long enough for them to adapt to fluctuations in flow.
   Discharge did not exceed 2500 cfs between October 15 and November 30 of the 2011 or 2012 water-years.

3. Maximum down-ramp rule for Thermalito Afterbay releases into the HFC:

   When release from Thermalito Afterbay outlet is less than 2500 cfs in the Feather River HFC, a maximum down-ramp of 200 cfs is applicable on a rolling 24-hour basis for the entire period that the discharge is below the threshold.

Constraints for the Feather River LFC site:

1. Minimum flow rule for Robinson Riffle (LFC):

   A minimum instantaneous discharge of 600 cfs is applicable all year round.

2. Maximum down-ramp rule for releases through the LFC:

   Scenario 1: When discharge is less than 2500 cfs at the LFC, a maximum down-ramp of 300 cfs is applicable on a rolling 24-hour basis for the entire period during which the discharge is below the threshold.

Scenario 2: When discharge is between 2501 and 3500 cfs at the LFC, a maximum down-ramp of 500 cfs is applicable on a rolling 24-hour basis for the entire period during which the discharge is below the threshold.

Scenario 3: When discharge is between 3501 and 5000 cfs at the LFC, a maximum down-ramp of 1000 cfs is applicable on a rolling 24-hour basis for the entire period during which the discharge is below the threshold.

Constraints for the Thermalito Afterbay site:

1. Minimum reservoir elevation rule for preventing water-fowl nesting along the shoreline:

   A minimum elevation of 133.5 ft must be reached once every 9 days from March 15 through May 31, without exceeding 134 ft.

2. Stable reservoir elevation rule for allowing grebe nesting along the shoreline:

   Elevations must be within a range of 132 to 135 ft from July 1 through August 15 and a range of 131 to 136 ft from August 16 through September 15.

3. Minimum reservoir elevation rule for allowing Giant Garter Snake foraging along the shoreline:

   A minimum elevation of 133.5 ft must be maintained for a continuous 12-hour period at least once a month from May through September.

Constraints for the Oroville Fish Hatchery site:

Maximum water temperature limits at the Oroville Fish Hatchery site during specific periods of the year are listed in Table 4-40.

TABLE 4-40 Constraints Related to Water Temperature for the Fish Hatchery Site

Period	Maximum Temperature °F
January 1–May 15	55
May 16–May 31	59
June 1–June 15	60
June 16–August 15	64
August 16–August 31	62
September 1–September 30	56
October 1–December 31	55

Developing Rating Curve for the Low and High Flow Channels. In order to develop discharge-stage relationship for the low and high flow channels of the Feather River, data from two transects were used for each reach. In the high flow channel, discharge-elevation data at River Mile (RM) 50.4 and 54.95 were used, while for the low flow channel, discharge-elevation data at RM 64.996 and 65.002 were used. Transect RM 50.4 represents Gridley Gauging Station, RM 54.95 represents southern end of the Oroville Wildlife Management Area, and RM 64.996 and 65.002 represent channel transects at the downstream and upstream sides of the Highway 70 bridge that traverses the Low-Flow Channel. The raw data containing discharge and the corresponding elevation values for three of the four transects were obtained from a Hydrologic Engineering Centers River Analysis System dataset provided by CDWR (Tables 4-41 and 4-42). For the transect RM 50.4, discharge and elevation data were obtained from California Data Exchange Center reported values for Gridley Gauge. In some instances at RM 50.4, elevation values for specific discharge levels were obtained from linear interpolation between two observed values (Table 4-42).

Initially, the bed elevation for each transect was approximated by fitting a third-order polynomial to the discharge(x)-elevation(y) data and using the y-intercept value (Table 4-43). Stage values were then computed from elevation by subtracting the constant from all elevation values for a particular transect (Tables 4-41 and 4-42). Because of variation in channel profile between transects, a generalized channel profile for a particular section was developed by calculating the mean stage values from the transects representing that section (Tables 4-41 and 4-42). The mean stage at specific discharge levels for the low and high flow channels were used to develop the final discharge stage relationship (Figures 4-112 and 4-113). Rather than fitting a curve, the final discharge stage relationship is composed of three-segment piecewise linear regression. A piecewise regression routine that would be available as part of the WUOT graphic user interface (GUI) was used to fit regression lines by iterations to minimize sum square of errors based on user-specified number of breakpoints. While linearization allows faster computation of a stage from a particular discharge value during optimization that involves hundreds of iterations for each time step, the multiple pieces maintain the integrity of the generalized channel profile of the particular section of the river (Figures 4-112 and 4-113).

Historic IRF Scores for the Feather River. IRF scores for the Feather River HFC site (described in Section 4.4.2.3) were calculated using the IRF historic mode, and the objective scores were calculated similarly using objective definitions mentioned in that section (see Section 4.4.2.3). Scores in the historical mode were based on the observed annual hydrograph reported from January 1984 to March 2011 at the Gridley gauging station (Figure 4-114). The overall annual historic IRF scores for the HFC Site of the Feather River, along with annual scores for individual objectives are also shown in Figure 4-115. Long-term discharge data was unavailable for the LFC-site and therefore objective and site scores using the IRF Historic Tool, which works in standalone mode, could not be calculated for this site.

TABLE 4-41 Discharge, Elevation, and Stage (computed) Values for Two Transects in the Low Flow Channel of Feather River

	RM 65.002		RM 64.996		
Discharge (cfs)	Elev. (ft)	Stage (ft)	Elev. (ft)	Stage (ft)	Mean Stage (ft)
600	124.82	0.78	124.82	0.79	0.78
800	125.09	1.05	125.08	1.05	1.05
1000	125.33	1.29	125.33	1.3	1.29
1500	125.88	1.84	125.88	1.85	1.84
2000	126.38	2.34	126.37	2.34	2.34
2500	126.84	2.8	126.83	2.8	2.80
3000	127.29	3.25	127.29	3.26	3.26
3500	127.68	3.64	127.67	3.64	3.64
4000	128.04	4	128.03	4	4.00
4500	128.39	4.35	128.38	4.35	4.35
5000	128.71	4.67	128.7	4.67	4.67
5500	129.03	4.99	129.01	4.98	4.98
6000	129.33	5.29	129.32	5.29	5.29

TABLE 4-42 Discharge, Elevation, and Stage (computed) for Two Transects in the High Flow Channel of Feather River[a]

	RM 54.95		RM 50.4		
Discharge (cfs)	Elev. (ft)	Stage (ft)	Elev. (ft)	Stage (ft)	Mean Stage (ft)
1000	84.2	0.9	73.01	0.16	0.55
2000	84.7	1.4	73.98	1.13	1.26
4000	85.3	2.0	75.19	2.34	2.18
6000	87.0	3.7	76.17[a]	3.32	3.49
8000	88.1	4.8	77.05[a]	4.2	4.52
10,000	89.1	5.9	77.87[a]	5.02	5.44
12,000	90.0	6.7	78.61	5.76	6.25
20,000	93.1	9.8	81.29[a]	8.44	9.13
40,000	97.5	14.2	86.61[a]	13.76	13.99
60,000	100.2	16.9	88.54[a]	15.69	16.29

[a] Elevation values are obtained from linear interpolation.

TABLE 4-43 Computed Bed Elevations at Four Cross-Sections in the High and Low Flow Sections of the Feather River

Transect ID	Bed Elevation (ft)	R-square from Third-Order Polynomial
RM 50.4	72.85	0.9983
RM 54.95	83.29	0.9986
RM 64.996	124.03	1
RM 65.002	124.04	1

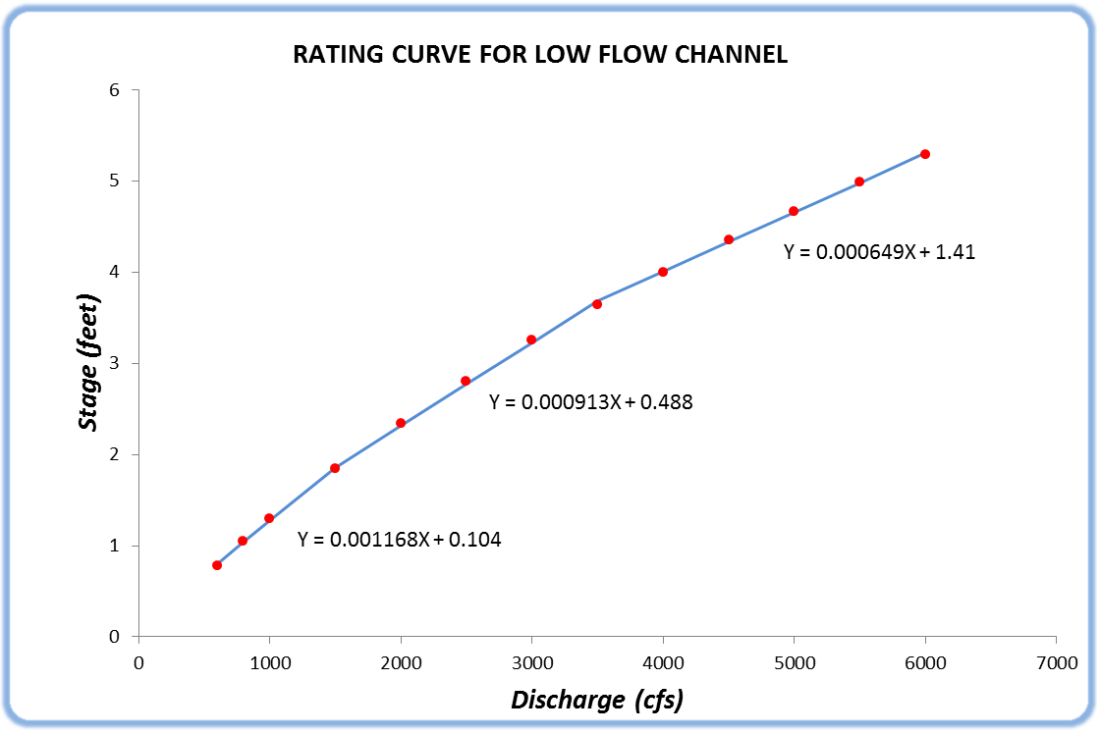

FIGURE 4-112 Discharge Stage Relationship Developed for the Low Flow Channel of Feather River Using Three-Segment Piecewise Linear Regression

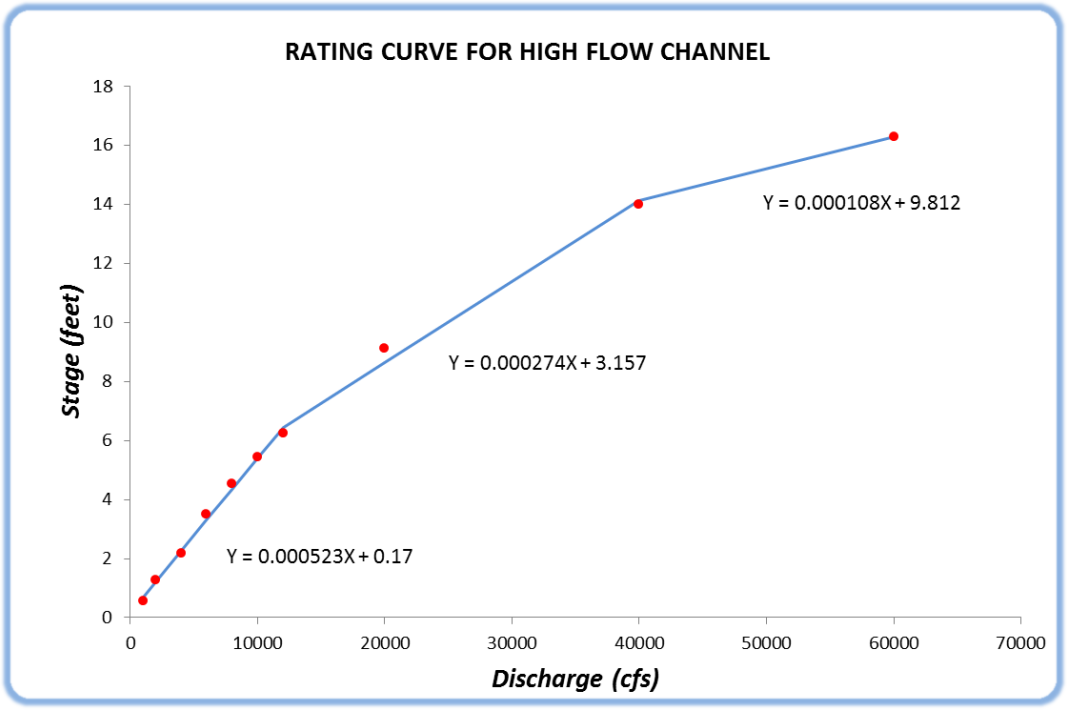

FIGURE 4-113 Discharge Stage Relationship Developed for the High Flow Channel of Feather River Using Three-Segment Piecewise Linear Regression

The annual environmental performance score (IRF score) for the HFC site varied from 0 to 0.21 during the 28-year period from 1984 to 2011. Although scores for most individual objectives fluctuated during the analysis period, annual site IRF scores were zero during the first 11 years as one or more objectives did not get any score during that period; however, from 1995 onward, objective scores were more indicative of performance based on defined relationships (Figure 4-115). This is because there are comparatively more fluctuations in cumulative frequency during the initial period with fewer records. On an individual basis, even though none of the objectives were accomplished up to their mentioned desired frequency, objective 1.5, which was related to the bed-mobilization, had the highest performance relative to other objectives for this site. Scores for several of the objectives increased in 2006 and again in 2011(Figure 4-115), as a result of the high flows with appropriate flow-regime features observed in those years (Figure 4-114). As of the most recent record (i.e., 2011), performance among all objectives varied from 0.06 to 0.71, from a long-term historical perspective. However, objective 1.7, which was related to hardwood stand development on the riparian areas, was not accomplished in any of the years during the period of record. In order to prevent the site score from becoming zero for all years as a result of non-occurrence of that objective, we chose a feature on the IRF tool that exercised the assumption that the objective was accomplished before the period of record. In addition, objective 1.2, which is related to spawning in steelhead trout, was accomplished to the desired level only in the last year of analysis (i.e., 2011); in all earlier years its score was zero. As a result, the feature of the IRF tool, which exercised the assumption that an objective was accomplished before the period of record, was also used for this objective.

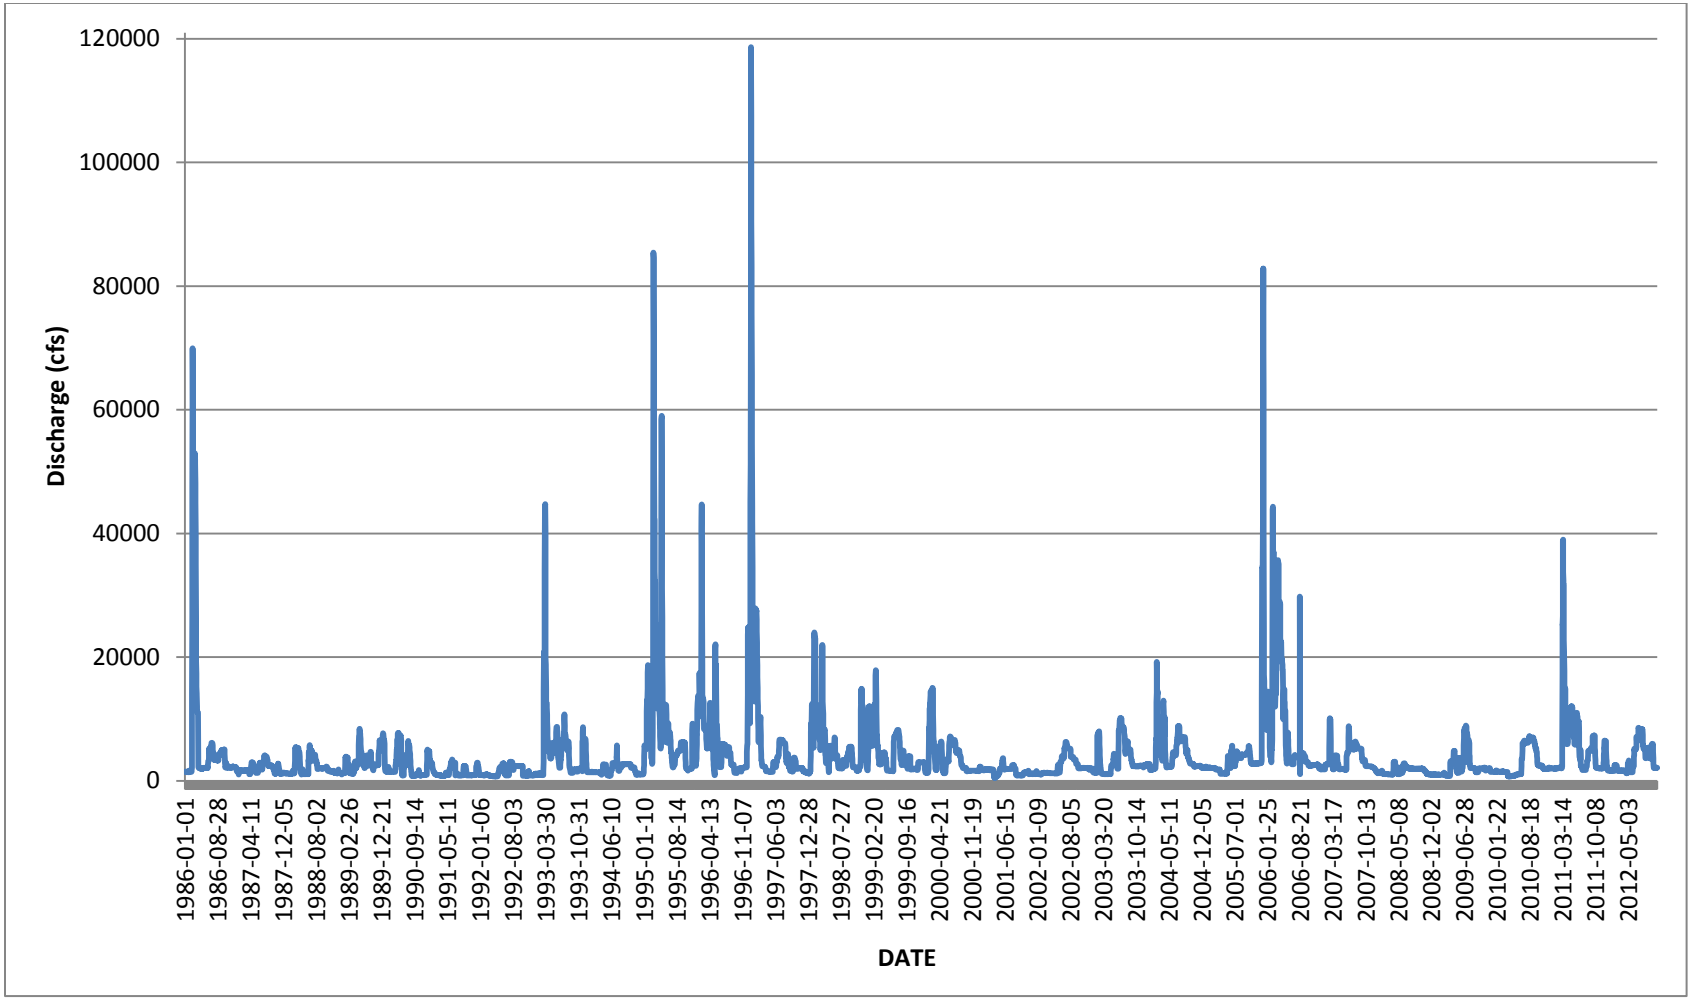

FIGURE 4-114 Actual Observed Hydrographs for the HFC Site in the Feather River, California, from 1986 to 2012 (Flow data for the HFC site was recorded from CDWR gauge located near Gridley, California.)

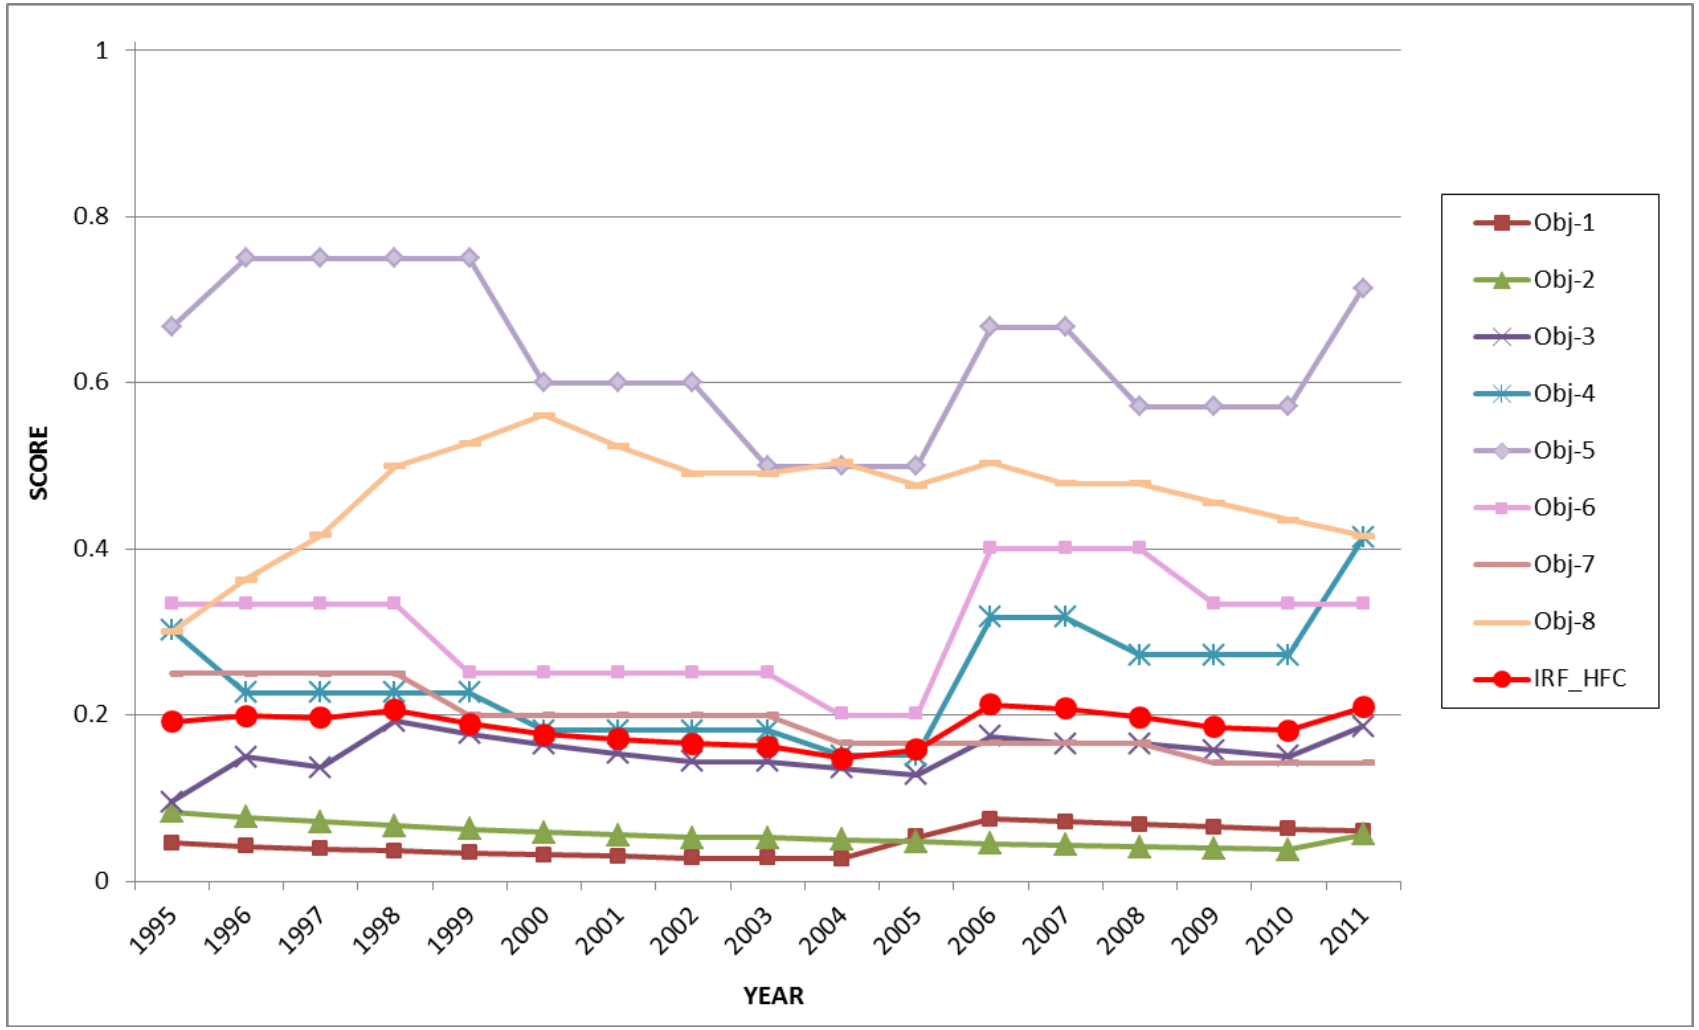

FIGURE 4-115 Annual Site IRF Scores (shown in red) and Objective Performance Scores for Objectives 1.1–1.8 for a 17-year Period (1995–2011) at the Feather River HFC Site Using the IRF Historic Tool (Data from 1984–2011 was used in the analysis; scores prior to 1995 are not shown because several of the objectives did not obtain any score in the first 10 years of record.)

IRF Scores for the Feather River Using IRF-Planning Tool Integrated within the WUOT. The planning mode of the IRF tool was used to calculate overall IRF scores at every 6-hour time step in order to develop the optimized hydrograph for Feather River between April 1, 2011, and September 30, 2011 (Figure 4-116). In this case, optimization between the IRF tool and the HydroSCOPE Tool took place as if real time was March 31, 2011; it advanced to the next time step to generate the most optimized hydrograph in terms of environmental and economic performance for that time based on inflow and power-price forecasts, while respecting applicable constraints. With subsequent advancement to the next time step, the process repeated itself until the entire modeling period had been optimized.

Integration of IRF-Tool with the Seasonal Hydrosystems Analysis Tool of the WUOT improved the overall environmental performance resulting from operation of the Oroville-Thermalito Complex on the Feather River. Overall environmental performances as indicated by the IRF scores were better for almost all the days in the 6-month modeling period compared to the observed facility operation during the same period (Figure 4-116). Site-specific scores for the HFC and LFC sites during the 6-month modeling period are also depicted in Figure 4-116. Even though improvement in the overall environmental performance in the optimized operation was somewhat moderate, it is important to note that enhancement in environmental condition was not accomplished by compromising on the revenue from power-generation, while meeting the power- and water-delivery commitments of the facility. The application of IRF component guided the HydroSCOPE Tool from an environmental perspective to identify co-optimal solutions for jointly maximizing environmental performance and power revenue. Moreover, these improvements in environmental conditions are above and beyond the environmental constraints that the facility had to abide by during actual operations.

Results from the HydroSCOPE optimization indicated that environmental performance improved at the HFC site (Figure 4-116). Figure 4-117 shows daily IRF scores under the optimized operation for all objectives that were defined for this site. For some of the objectives in the optimized run, IRF scores improved over the actual observed facility run; however, for other objectives, full accomplishments were noticed for both the observed and optimized runs (objectives 1.1, 1.2, and 1.5). The latter aspect of reaching the environmental ceiling for those objectives also explains why there was little relative difference between the environmental performance of observed and optimized runs for this site (see Figure 4-117), as any further improvement in performance beyond full potential is not possible. From an individual objective perspective, the greatest improvement in environmental performance was seen in objective 1.4, which pertains to accomplishing Chinook salmon outmigration by releasing a spring pulse flow (Figure 4-118). It should also be noted that for one of the worst-performing objectives (objective 1.7), which is related to stand development for riparian hardwoods, we chose a feature on the IRF tool that allowed us to state that the objective was accomplished before the period of record. Without this action, the score for this objective would have gone to zero (which also would have resulted in the site and overall IRF scores for the facility becoming zero) because this objective was never accomplished during the period of record, even though indirect evidence from riparian tree lines indicates this objective was accomplished during the pre-regulated period.

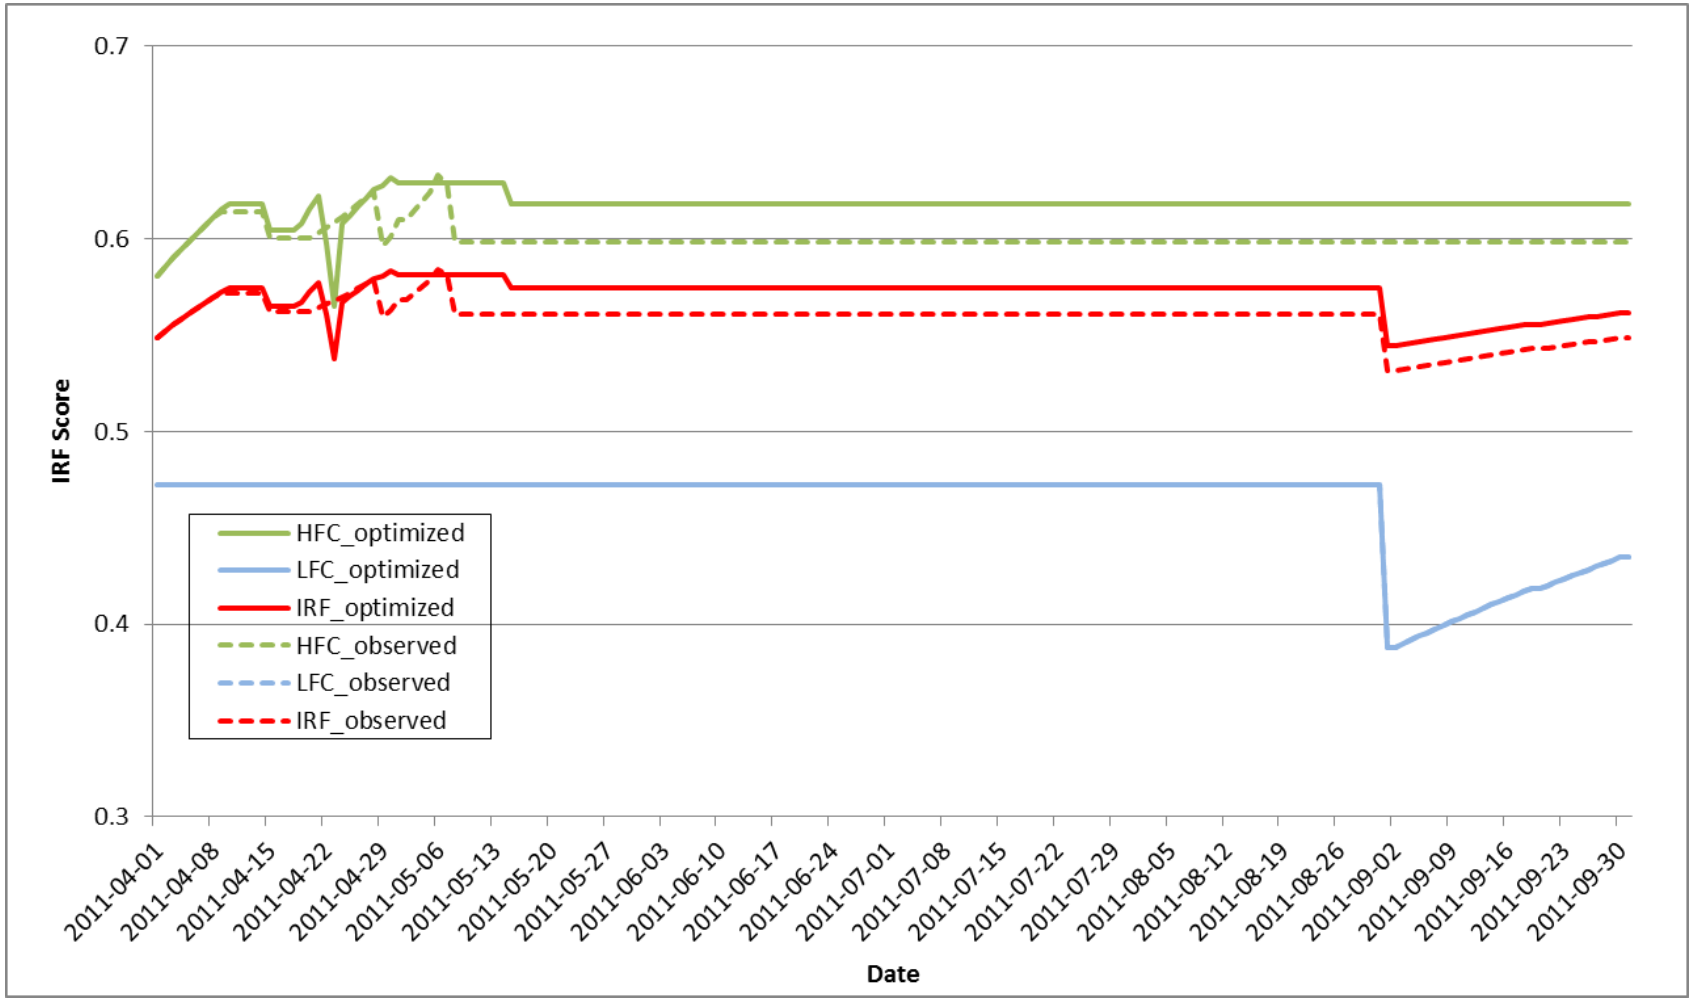

FIGURE 4-116 Daily IRF and Site Environmental Performance Scores Based on Observed Facility Operation (dotted line) and IRF Integrated WUOT Optimized Simulation (solid line) for the Feather River Using the IRF Planning Tool for April 1, 2011, to September 30, 2011 (The observed and optimized scores for the LFC site were identical.)

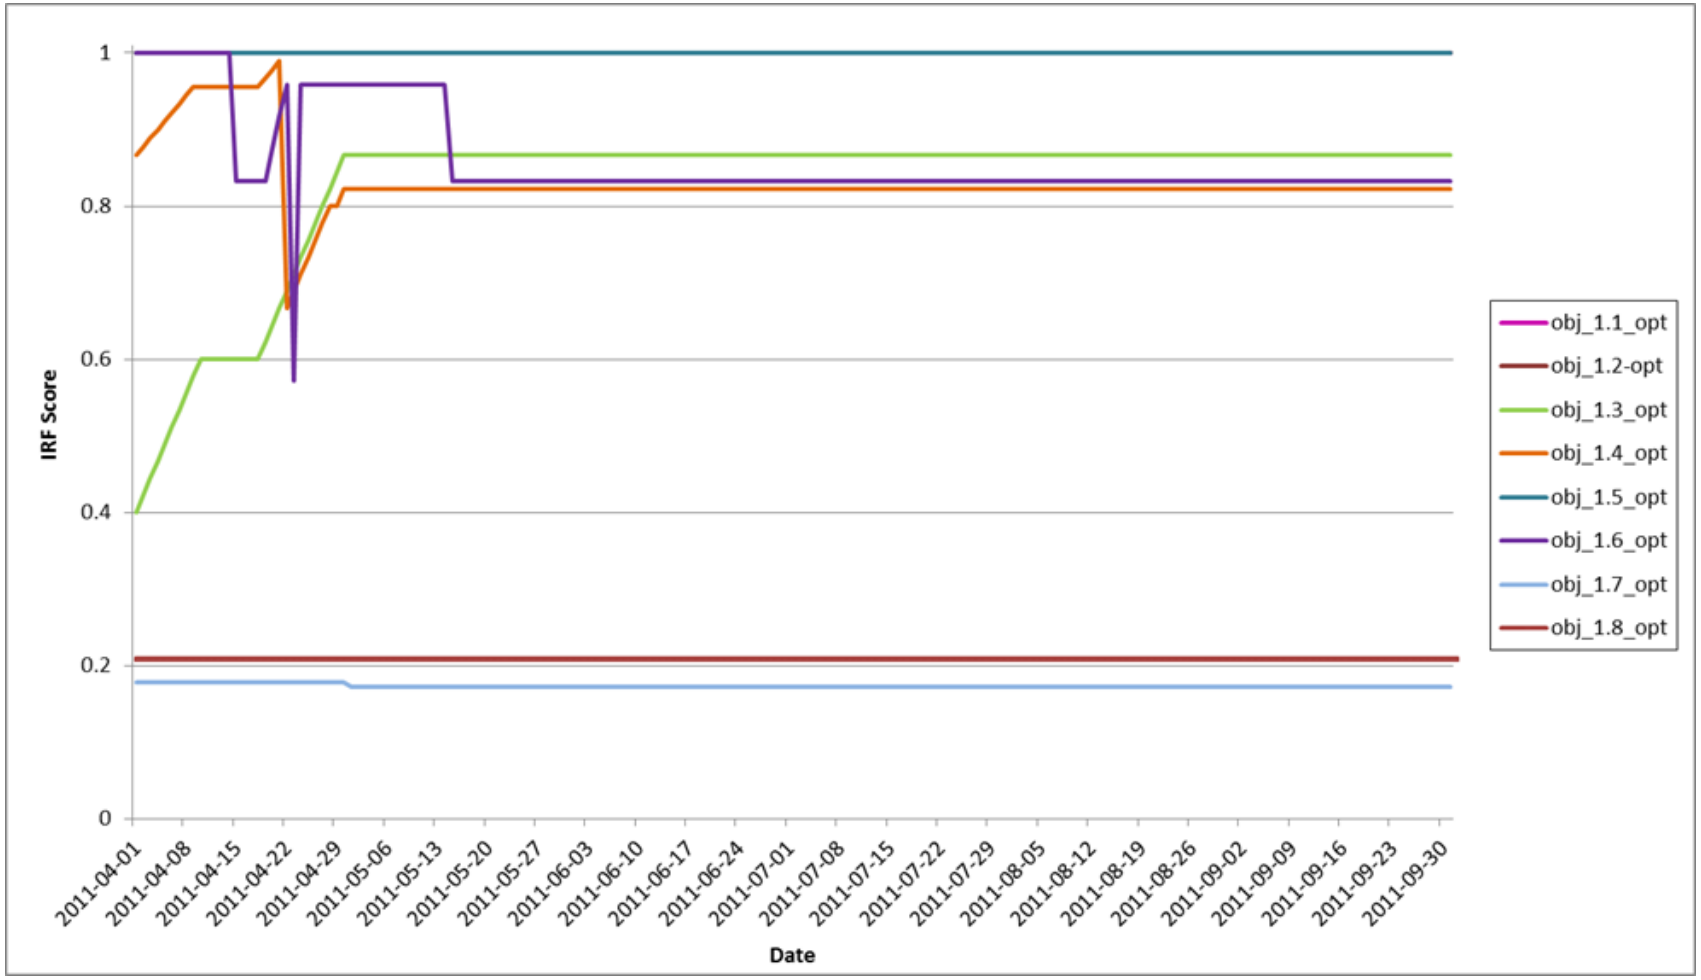

FIGURE 4-117 Daily Objective Performance Score for Objectives 1.1–1.8 Based on Optimized WUOT Simulation for the Feather River Using the IRF Planning Tool for April 1, 2011, to September 30, 2011, at the HFC Site (Objectives 1.1–1.8 are described in detail earlier in the result section. Objectives 1.1 and 1.2 are masked in the plot because of overlapping scores, but were accomplished to the maximum level; their IRF scores were 1 during the modeling period.)

Results from the HydroSCOPE optimization further indicated that improvement in environmental performance was limited only to the HFC site; however, for the LFC site, environmental performance of the optimized run was almost identical to that of the actual facility run. This is also evident from the optimized and observed hydrographs from the two sites (Figure 4-119). While modifications in reservoir releases from the Thermalito Afterbay Reservoir affected the hydrograph at the HFC site in the optimized run, releases from the Thermalito Diversion Dam into the LFC site in the optimized scenario were identical to those during actual facility operation. Hydrographs at the HFC site indicate that peak flow during the modeling period occurred around April 25, 2011, in the optimized operation. This magnitude of peak flow was approximately 50% larger, albeit for a shorter duration, than the observed facility operation. Another reason for only moderate improvement in the overall environmental performance is that even though the 2011 water year was comparatively a hydrologically wetter year, most of the high flow occurred before the modeling period (i.e., before April 1, 2011). As a result, flexibility was not a possibility in reshaping the high flow periods of the hydrograph to enhance environmental and economic value. In all probability, the IRF scores would have been higher had the modeling period started 2 months earlier.

One of the primary reasons the LFC site did not contribute toward improving the overall environmental performance was that only three objectives were defined for that site, compared to eight for the HFC site (Figure 4-120). Since most of the water that would have naturally flowed through the LFC site is bypassed through a complex channel of canals and reservoirs for power generation, only a few environmental objectives were relevant at this 8-mile reach. Because the overall IRF score is a geometric mean of all environmental objectives, a site with fewer objectives (in this case the LFC site) has less influence on the overall environmental performance for the facility. Even though, from an environmental standpoint, significant spawning in salmonids is known to occur at the LFC site, the upstream end of this segment acts as a “dead-end” as far as any natural upstream fish migration is concerned. The basic plumbing design of the facility offers limited flexibility for this site under the existing infrastructure and operation.

We also observed that, for some of the time steps, hard constraints related to reservoir elevation at the Thermalito Afterbay and water temperature limits at the Oroville Fish Hatchery were violated in the observed facility run. Not knowing the extenuating circumstances that led to those violations, we also relaxed the constraints in the optimized run to the same magnitude as observed in the observed facility run for those time steps. This also allowed fair comparison between the observed and the optimized operations. In some instances the violations were not avoidable because of initial condition (i.e., March 31, 2011). For example, in early April, when the optimization started, the reservoir was low to the extent whereby even with attempts to raise the reservoir elevation, it could not be done soon enough to prevent the violation. Conditions on or before the modeling period limits the optimization model as to how soon certain thresholds can be attained, so in this regard of constraint violation, the optimized runs were no worse than the observed facility operation during the same period.

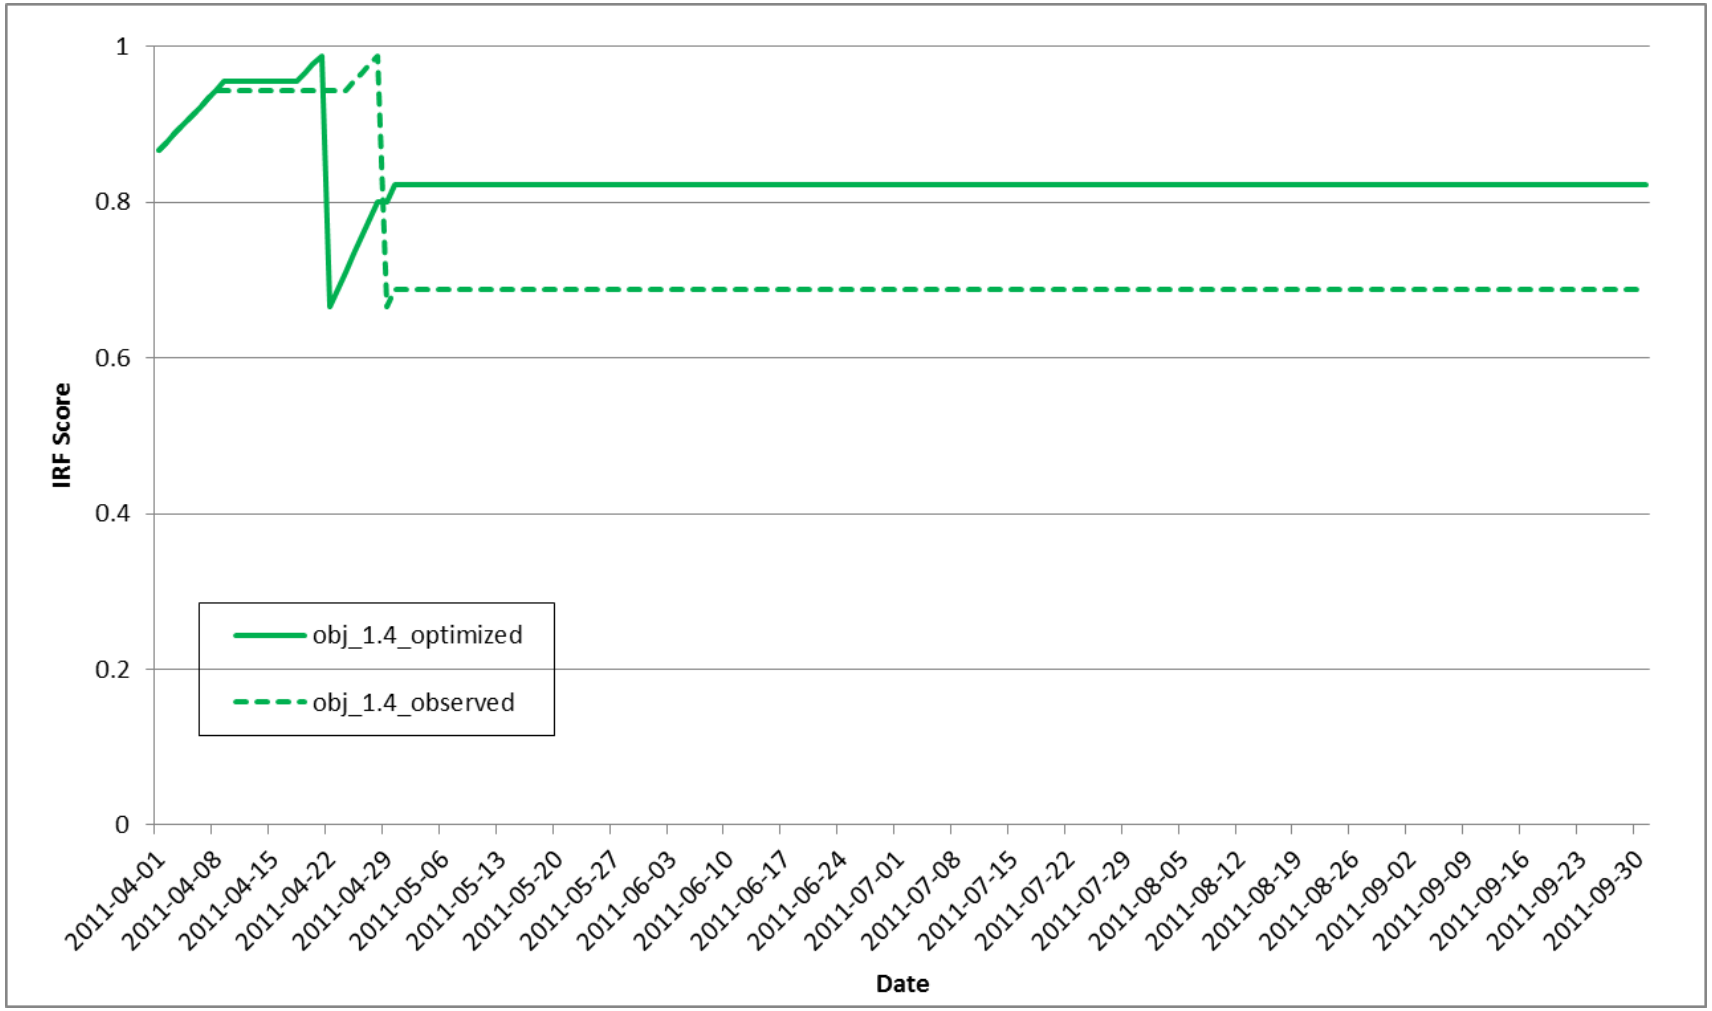

FIGURE 4-118 Comparison of Daily Objective Performance Score for Objective 1.4 between Observed Facility Operation (dotted line) and Optimized WUOT Simulation (solid line) for the Feather River Using the IRF Planning Tool for April 1, 2011, to September 30, 2011, at the HFC Site (Objective 1.4 pertains to accomplishing Chinook salmon outmigration by spring pulse flow, and is described earlier in the result section.)

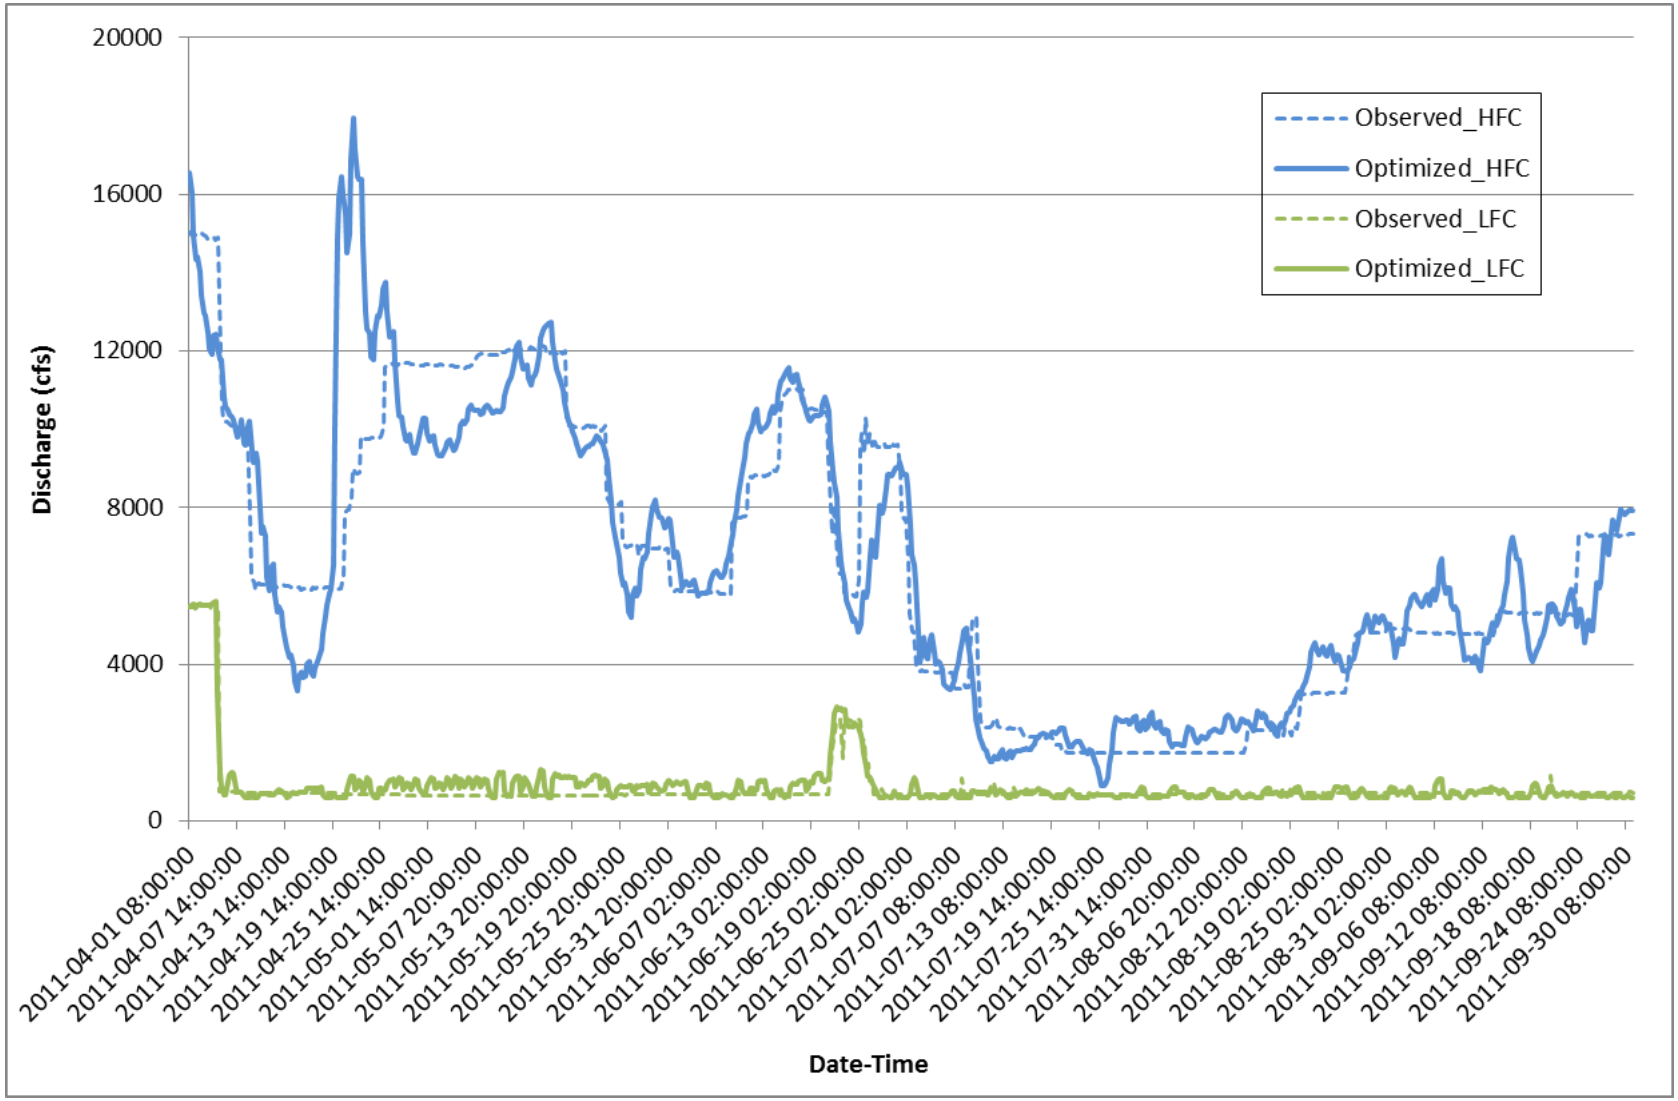

FIGURE 4-119 Actual Observed (dotted line) and WUOT Optimized (solid line) Hydrographs for the Two Sites in the Feather River from April 1, 2011, to September 30, 2011

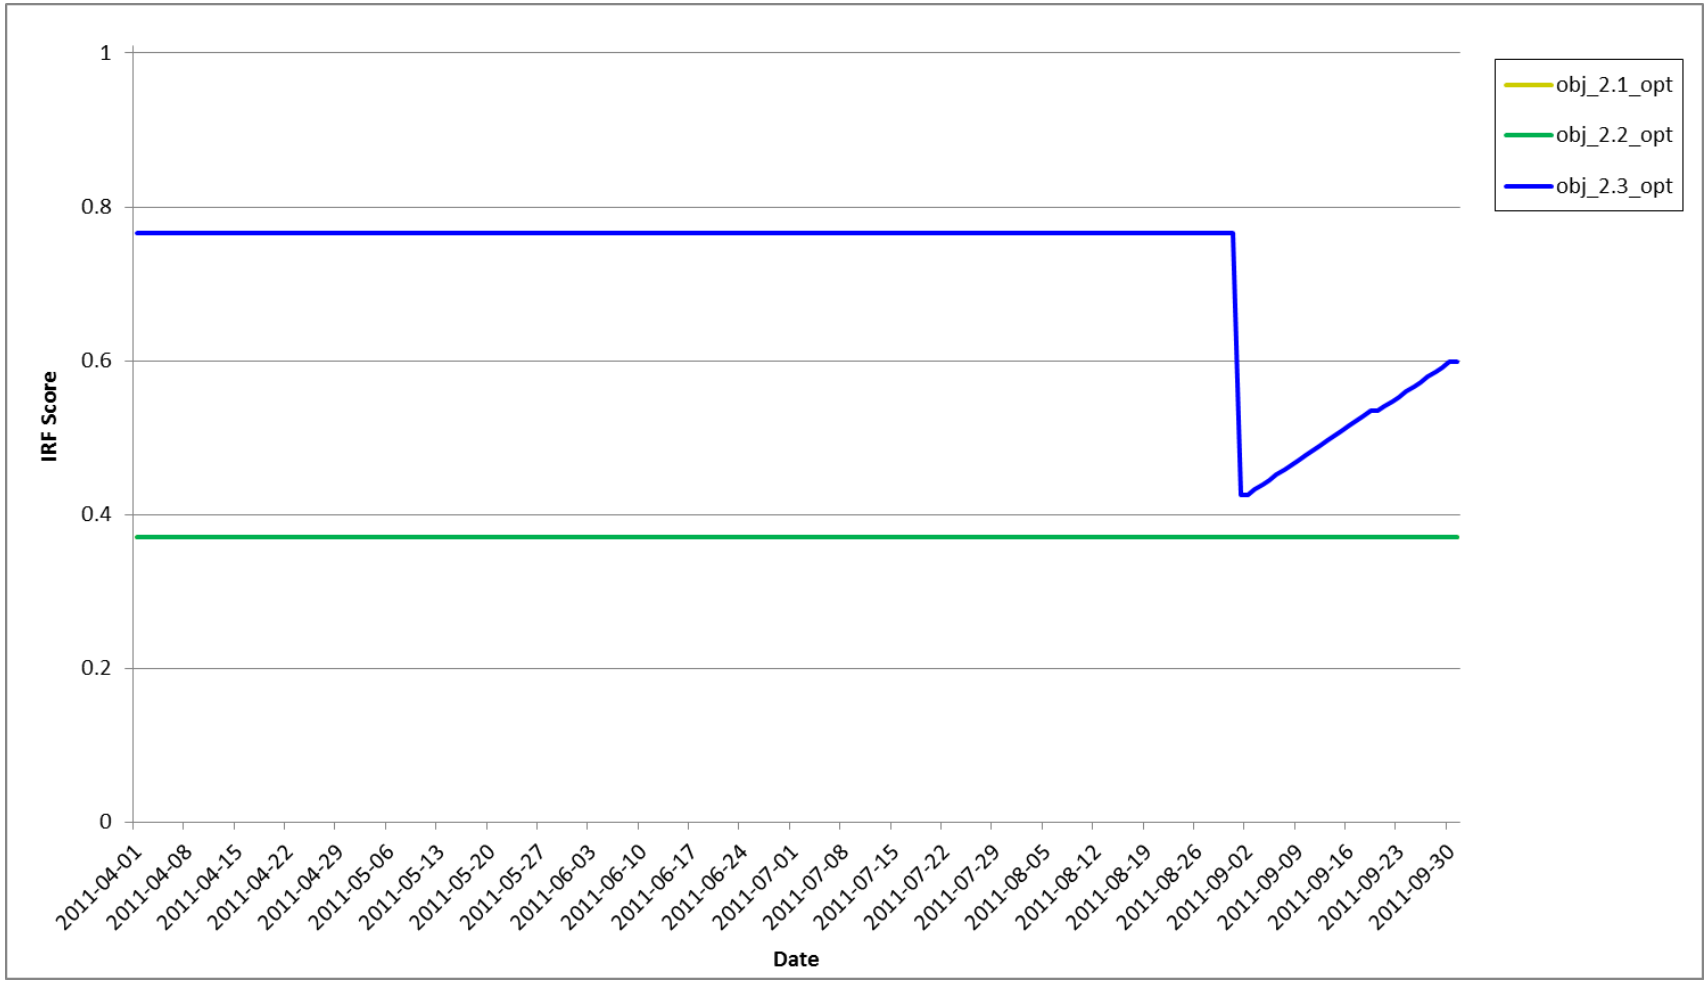

FIGURE 4-120 Daily Objective Performance Score for Objectives 2.1–2.3 based on Optimized WUOT Simulation for the Feather River Using the IRF Planning Tool for April 1, 2011, to September 30, 2011, at the LFC Site (Objectives 2.1–2.3 are described in detail earlier in the result section. Scores for Objectives 2.1 and 2.2 were identical during the modeling period.)

Further reshaping of the hydrograph (beyond what was provided by the Seasonal Hydrosytems Analysis Tool) by the Day-ahead and Real-time Scheduling Tool and its effect on the hydrograph was not considered because IRF was used only by HydroSCOPE and not by the Day-ahead and Real-time Scheduling Tool. Because the Day-ahead and Real-time Scheduling Tool has to use the amount of water allocated by the Seasonal Hydrosytems Analysis Tool with all the operational (including environmental) constraints in place, the IRF-HydroSCOPE optimized hydrograph is close enough to satisfy the environmental objectives so that the Day-ahead and Real-time Scheduling Tool can also meet them, as long as there is an actual scheduler running the model every day to provide that extra bit of guidance when needed. However, because in this modeling exercise the model ran in an automated batch mode (in absence of any course corrections between time steps) for which it was not designed, that extra guidance was not present in the retrospective analysis. As a result, real-time modeling result was not used, because it would impart little meaning to the alterations in the IRF score. The use of the toolset in a batch mode—where course corrections at intermediate timesteps based on guidance from the day-ahead and real-time scheduling tool were lacking—caused the HydroSCOPE tool to rely on actual operation as a guide for improving IRF scores, which sometimes limited the toolset from making substantial gains. If used as intended in a real-life scenario, WUOT's application is anticipated to generate additional gains, especially for environmental performance.

Interpretations of the functional relationships between objective metrics with associated parameters might not be unequivocal. Since the facility staff did not develop these environmental objectives, they should not be held accountable for any unintended errors. At the same time, we would also like to reiterate that we used the best scientific information available to develop those relationships, which were also validated using the IRF routine; however, these relationships have not been explicitly reviewed (CDWR staff reviewed only the environmental constraints applicable at the facility) and have not been tested by the WUOT study team or the facility staff using computational models or field verification for their efficacy and sensitivity. We acknowledge that validated population models are needed to predict the population-level impacts of alternative flow regimes on environmental resources. One reason for controversy related to definition of environmental objectives is their simplification. For example, it can be argued that some of the objectives, like fish spawning and fry development, should not be treated as isolated objectives and need to be integrated into a more comprehensive objective. While the importance of those ecological nuances are fully acknowledged, the IRF component of the WUOT is still in its developmental stages and for the next version additional complexities like linking and sequencing of objectives would be tackled in the working framework. Obviously, as better relationships are identified in the future than those used in this study, the entire optimization routine can be re-run to identify additional improvements in both the environmental and the economic features. The most important aspect of the demonstration studies is that it underscores that the environmental objective definitions are just examples to demonstrate the IRF tool's utility and value, rather than justifying the exact nature of the functional relationships or the resulting outputs.

The intent of integrating the IRF component into the WUOT is to demonstrate that the IRF template can provide a mechanism to implement environmental models vetted by the hydropower facilities, which allows them to go beyond minimum constraints and identify

additional possibilities of environmental and economic enhancement while gaining a better understanding of the involved tradeoffs. Even if mathematically optimal decisions are ignored, the IRF-integrated WUOT's knowledge about the site-specific and pertinent tradeoffs allows the facilities to make informed decisions related to running the facilities from a spatial and temporal context. Moreover, explicit environmental improvements beyond minimum mandated requirements may also be demonstrated as mitigation measures by the facilities during the renewal of licensing agreements with the regulatory agencies.

4.4.2.4 Application of CHEERS for Oroville-Thermalito Complex Day-Ahead and Real-Time Scheduling

The usefulness of the CHEERS model for CDWR applications was demonstrated by a retrospective analysis over a 6-month time span from April 1, 2011, through September 30, 2011, a period selected by CDWR staff. The Oroville-Thermalito Complex model optimization period was selected because the Market Redesign and Technology Upgrade (MRTU) had been implemented by the California Independent System Operator (CAISO) and the CWDR had gained sufficient experience operating under MRTU by April 1, 2011.

The CHEERS modeling approach developed for the retrospective analysis of the Oroville-Thermalito Complex is in many ways similar to the one used for the Colorado River Storage Project (CRSP). Both analyze model water and electricity flows throughout the system; reservoirs are cascaded and operations are constrained; and water delivery obligations are of paramount importance to system operations. In addition, both use functions that relate power production levels to turbine water flow rates.

However, there are a few key differences that make the operation of Oroville-Thermalito Complex different from the CRSP system:

- The Oroville-Thermalito Complex has inflows from a relatively large drainage area with significantly more storage compared to the Aspinall Cascade. The Oroville-Thermalito Complex also has a distinctly different hydrological profile.

- The CDWR participates in and sells power production and ancillary services from the complex to the CAISO wholesale market, which operates as a commodity exchange. Financial product streams modeled for the Oroville-Thermalito Complex include energy, regulation service up, regulation service down, spinning reserves, and non-spinning reserves. In contrast, CRSP transactions are primarily limited to the bilateral energy markets.

- The Oroville-Thermalito Complex model objective is to maximize revenues from the sale of the five product streams listed above. CRSP model runs primarily optimize energy within a set of dispatch guidelines.

- CDWR sells all production to the market since there are no firm loads to serve. CRSP objectives focus on fulfilling energy delivery obligations to its customers.

CHEERS performed a day-ahead optimization of the Oroville-Thermalito Complex for each day in the retrospective analysis period. Each run consisted of 24 hourly time steps. The results were compared to the actual historic operation of the complex. Each model run has the same basic objective, namely to maximize revenue by selling energy and ancillary services to the market, subject to a set of power and reservoir operating criteria and management rules. CHEERS model runs were not performed for real-time dispatch, because the primary use of CHEERS by CDWR staff will be for day-ahead scheduling purposes only. It is important to note that due to the business-sensitive nature of CDWR operating information in the CAISO market, the model results and historical operation information presented in this section are not nearly as detailed as those given for the CRSP analysis in Section 4.4.1.4.

Data Sources and CHEERS Model Inputs. Supporting data describing the characteristics of the complex and detailed actual historic operations were provided by CDWR staff. Hourly data provided by CDWR included unit-level water releases for power production, along with the corresponding power output, the elevation at each reservoir, the flow at Robinson Riffle (low flow channel) and Gridley (high flow channel), fixed reservoir inflows and outflows (such as Palermo Canal, Kelly Ridge, and irrigation withdrawals), station service loads, and market clearing prices for energy and ancillary services. Inflows to Lake Oroville were calculated externally by performing a water balance using the known surface elevation along with the other known system inflows, outflows, and releases. Functions that relate reservoir storage volume to elevation were also provided by CDWR staff. However, 2011 data were not available for either the power flow or generation at the Diversion Dam plant. Non-power reservoir water releases were also not provided for any of the dams.

Water-to-power conversion functions were developed by Argonne staff using 2011 historical data. The historic source of this information was found to have problematic observations and to some degree had to be manually cleaned. Specifically, numerous hours were identified that had records with one of the following issues:

1. Significant water release with zero associated generation;

2. Significant generation with zero or negligible water release;

3. Extremely high generation output for the associated flow rate; or

4. Hours in which a unit was known to be out for maintenance yet still showed generation and/or turbine flows.

The data were cleaned by zeroing out both the unit's water flow and the generation output in most hours that exhibited the issues described above (even after cleaning there remain some hours with values that are questionable, though not definitively incorrect). This approach

for cleaning the data facilitates the fairest comparison between the model results and the historic record, as opposed to, for instance, an approach in which the modelers would fill in missing or invalid gaps with estimated values they deemed suitable. The historic data were observed to have an excessive number of problems during April 2011, and even after being cleaned the historic record would have been at an unfair disadvantage when compared to the model results; therefore, the comparison period was shortened to cover May through September 2011.

The cleaned data were used to construct water-to-power conversion functions. Consistent with fundamental differences in the operating characteristics of generating resources at the Hyatt Power Plant, one conversion function was developed for units 1, 3, and 5, and a second function was developed for units 2, 4, and 6. At Thermalito one function was applied to unit 1, and a second one was developed and applied to each of the three other units. Because the power output of Hyatt units is significantly affected by the wide-ranging head, the Hyatt unit conversion functions were adjusted each week to more accurately represent the conversion efficiency associated with the Lake Oroville surface elevation. Such adjustments were not made for the Thermalito plant because the historical data did not show a clear relationship between head and the water-to-power conversion efficiency. Conversion functions applied to the Hyatt and Thermalito units were restricted by both minimum and maximum operating limits. Consistent with historical 2009 observations, a simple linear conversion function was developed for the unit at the Diversion Dam Plant.

Although the hourly elevation measurements at the Diversion Pool and Thermalito Forebay can differ by up to 2 feet in a given hour (likely due to instrumentation measurement error or high wind and wave action), these two pools are essentially at the same elevation when the power canal gate that connects them is open. Therefore, the two pools are modeled in CHEERS as two storage nodes connected via a bi-directional water link. Modeled water flows in both directions in the canal to equalize the hourly elevation in the two pools. This approach was suggested by CDWR staff and reflects the same assumption implemented in their current models.

Since the CHEERS model is designed for short-term daily operations, each retrospective CHEERS model run was configured such that:

1. The day's modeled initial elevation condition matched the historic record at each reservoir;

2. The day's modeled ending elevation condition matched the historic record at the Diversion Pool and Thermalito Afterbay;

3. The model's hourly flow in the low flow channel must remain at least as high as the historic record;

4. The hourly modeled and historic flow match in the high flow channel; and

5. For each hour, all unit and reservoir operations remained within institutional, environmental, and physical operating limits.

The role of the HydroSCOPE model is to alter these daily values over a longer time horizon. HydroScope outputs were utilized by CHEERS in the integrated model run.

Provided it remained in sync with those historic values, CHEERS was allowed to release water hourly from the reservoirs to maximize the financial value of the complex's energy products, subject to operational and regulatory rules. However, the model was allowed to violate these rules during specific hours when actual operations violated a rule (as measured and recoded), and only to the extent that the actual violation occurred. No daily ending elevation target was explicitly enforced at the Thermalito Forebay because it has the same elevation as the Diversion Pool, which does have such a target.

Initial model runs did not enforce a target elevation at Lake Oroville because all other inflows and outflows to and from Lake Oroville were fixed to the historic record. CHEERS satisfied the other target constraints throughout the complex and high flow channel by releasing water from Lake Oroville, an amount one would expect to match the actual release as closely as possible. Model results using the configuration described above often showed an unrealistically large amount of water released from Lake Oroville. Allowing the large releases was the only feasible way to meet the other targets and constraints throughout the complex. As a result, the end of the day modeled elevation at Lake Oroville was well below the level shown in the historic record.

An additional constraint that requires the daily ending modeled elevation at Lake Oroville to match the historic elevation would not allow the model to satisfy water balance equations and result in infeasible solutions. That is, CHEERS cannot maintain the historical elevation and simultaneously release enough water to both meet the daily ending elevation targets at the lower reservoirs and meet historical water delivery obligations in the high flow channel. Since measurements of turbine water releases and power generation are relatively accurate, the cause of this water balance dilemma is most likely due to measurement error in the reservoirs and high flow channels. Although Lake Oroville elevation values are reported to one one-hundredth of an inch, the actual measurements do not lend themselves to this high level of accuracy, especially during periods of high wave action. The volume of water covered by only a 2-inch margin of error at a full reservoir is equivalent to a sustained release of 1,350 cfs over a day, more than enough to cause the water imbalance problem.

To address this issue, the model was reconfigured with the following changes:

1. On a daily basis, the modeled amount of turbine water released at each plant is not allowed to exceed the amount that was indicated in the historic record;

2. Hourly flows in the high flow channel must meet a minimum level rather than match a fixed level; and

3. The modeled Lake Oroville elevation must match the historic record at the end of each day.

To facilitate the change at Lake Oroville, it was necessary to allow the model to solve for inflows to Lake Oroville rather than specifying these flows exogenously. The inflow amount to Lake Oroville and the amount of water spilled from Lake Oroville were both assigned an extremely severe penalty in the model's objective function, such that the model would minimize water inflows and non-power releases while maintaining daily plant release targets and reservoir elevation ending targets.

It is important to recognize that the historic data are not perfect and therefore cannot present a perfect representation of what occurred in reality. Nonetheless, we have configured the model runs to coordinate as closely as possible with the record presented by the available data after subjecting it to the manual cleanup. As a result of the cleanup, the adjusted historic record likely indicates in aggregate less water released for power production and less generation produced than what actually occurred. However, as configured, the CHEERS model does not have the ability to release any more water for purposes of power production than is indicated in the cleaned historic record. These difficulties encountered in configuring the retrospective analysis demonstrate that the ability of a model to accurately represent a historic condition is directly related to the quality and accuracy of the historic information available. It should be noted that although CHEERS is being used in this retrospective analysis to simulate and analyze historic conditions, the model is instead designed to be used by actual hydro-schedulers as a planning tool for guiding future operations, and when used as a planning tool there is no need to conform to a historic record.

Executing the model using the revised configuration described above, it was then possible to meet all the constraints in place and to coordinate with the historic record. Although large spills from Lake Oroville were often observed in the model results, the daily ending elevation at all reservoirs now matched the historic record; there would have been no way to achieve this balance without using the spill amount to settle the imbalance in the historic record. It is evident that the variable Lake Oroville inflow and spill amounts observed in such situations are not being used to unfairly satisfy downstream flow requirements, as no water is spilled from the Diversion Pool into the low flow channel, nor is any water spilled from the Thermalito Forebay into the Afterbay.

Market Price Forecast. Under actual day-ahead scheduling conditions, future market prices are unknown. Therefore, schedulers use price forecasts when making day-ahead scheduling decisions. It is typical for a scheduler in the WECC footprint to base price forecasts on historical observations and expected weather conditions. The CHEERS retrospective analysis methodology mimics this process by using price forecasts instead of actual observed values. However, historical price forecasts that were made by schedulers during the retrospective analysis period are not available. Therefore, a price forecasting routine was created and implemented specifically for this CHEERS demonstration. This forecasting routine uses an econometric approach to project day-ahead prices for energy, regulation up, regulation down, spinning reserves, and non-spinning reserves.

The forecasting process computes the average price that occurred during the past week, with separate calculations made for weekdays and weekends. For example, the energy forecast

price for Wednesday at 1 AM is initially set equal to the average energy price that occurred at 1 AM during the past five weekdays. Identical computations are made for all other hours of the day. For projections made for a weekend day, an identical process is implemented, except hourly averages are based on past weekend price observations.

Energy Price Forecast. The energy price forecast routine built for the CDWR retrospective analysis employs the approach described above. First, average hourly prices are computed from historical price data from the past week based on the weekday/weekend status of the projection day. The forecast routine then adjusts prices based on next day atmospheric temperatures for three large cities across the state of California: Los Angeles, San Francisco, and Sacramento. Next, the algorithm generates a least-squares linear regression equation that correlates peak daily energy prices over the last 2 weeks to a temperature index. This index is a weighted average of both maximum and minimum temperatures for the three aforementioned cities, whereby the weights vary by city and temperature type. It should be noted that the minimum temperature weights are significantly smaller than the maximum temperature weights. The linear regression equation is then used to project the maximum energy price for the next day based on maximum and minimum temperatures for the forecast day. A similar approach is used to project the minimum energy price for the next day. However, minimum daily prices are linked to a minimum temperate index.

Once next-day maximum and minimum energy prices are projected, the daily energy price range is computed. Next, hourly energy prices are projected. The projection method employed shapes the daily price profile using an index that ranges from zero to one based on the average hourly price profile. The index indicates the relative difference between the daily minimum and maximum price such that the index for the hour with the lowest price is set equal to zero, the hour with the highest hourly price has an index value of one, and a price that equals the average of the minimum and maximum has an index value of 0.5. To obtain a projected energy price for any hour, the daily minimum projected price is added to the product of the daily energy range (i.e., maximum minus minimum projected price) and the hourly price shaping index.

Ancillary Services Price Forecast. Regulation up price projections are simply based on hourly average daily profiles for weekday and weekend days. No adjustments are made for atmospheric temperature because the statistical relationship between atmospheric temperature and prices appears to be weak. On the other hand, there is a very strong historical relationship between regulation up and spinning reserve prices; regulation prices are consistently higher than spinning reserve prices. Therefore, the methodology employed by the price forecast routine generates a least-squares linear regression equation that correlates hourly spinning reserve to regulation up prices for the past week; that is, 168 observations. Spinning reserve price projections for the next day are then estimated via the regression equation.

Both regulation down and non-spinning reserve price projections are simply based on hourly average daily profiles for weekday and weekend days. Non-spinning reserves are typically less than $1/MW. However, during some days, typically when energy prices spike

during super load hours, there is also a spike in non-spinning reserve prices. At some point, the price forecasting routine may be updated to better capture these price spikes.

Price Forecast Error. Monthly absolute price forecast error is presented in Table 4-44 for each product sold to the market. Although the locational marginal price (LMP) forecast has the highest prices and error level in absolute terms, its relative hourly price pattern follows the actual price pattern better than any of the ancillary service price forecasts.

Although price forecasts are imperfect in absolute terms, the daily relative price patterns ultimately drive model results. Therefore, as described in greater detail below, the potential financial loss attributable to imperfect price forecasting is relatively small when considering only energy sales versus sales from the slate of all products.

CHEERS Results—Baseline Scenario. The CHEERS model was run in standalone mode for several retrospective alternatives, and for one WUOT integrated model simulation. In the standalone baseline scenario, CHEERS allowed the generating units to only sell energy because ancillary services were not actually provided during the analysis period. Forecasted market energy prices developed using the aforementioned methodology were used to shape hourly operations. CHEERS was permitted to start each unit at most twice during the day, with a "start" being defined as the transition of the unit from generating zero energy in one hour to nonzero energy in the next. Every startup event was assumed to cost an amount that was specified by CDWR staff. Each unit was also permitted to ramp its output up and down over the course of the day by an amount equal to twice the unit's generating capacity. This guideline was implemented at the suggestion of CDWR in order to avoid potentially erratic price-following generation patterns in favor of something resembling more of a "two-hump" pattern over the course of the day. Individual generating units were available for dispatch in the model only during the same days which they were actually available according to the historic record.

Using identical daily plant-level turbine water release volumes as historically recoded and *projected* market prices, CHEERS solved for an economically optimal dispatch that realized

TABLE 4-44 Absolute Average Error between Forecast and Actual Hourly Price for Energy and Ancillary Services ($/MW)

Month	LMP	SR	NR	RU	RD
April	6.52	5.99	0.03	7.68	5.04
May	5.34	4.96	0.28	5.14	5.15
June	5.35	3.59	0.48	3.59	2.36
July	4.47	3.40	1.18	3.31	2.98
August	3.76	2.30	0.99	2.09	1.92
September	4.96	1.87	0.94	1.88	2.07
Period	5.00	3.59	0.68	3.79	3.19

significant gains in net revenue. In this comparison, revenues are calculated as the hourly modeled or observed generation times the hour's *actual*, not forecasted, market clearing price, minus any applicable startup costs. Relative to the observed revenue, CHEERS achieved a 13.0% increase. CHEERS started and stopped units more frequently than actual operations; however, costs associated with startups are included in the net revenue calculation.

As stated earlier, due to the business-sensitive nature of CDWR operating information, a detailed comparison of model results cannot be presented in this report. However, a general observation about the CHEERS hourly dispatch pattern is that it typically shows power output ramping over a larger daily range than the historic record. From week to week, the revenue improvement ranges from approximately 3% to 35%. The weeks with higher potential improvements coincide with times when the daily turbine water release volumes are relatively low, because the model has greater flexibility to operate the units in a different way than they were operated historically. These weeks contrast with those in which the daily turbine water release volumes are relatively high, in which case the model has less flexibility to find a more optimal operational pattern because the large volume of water necessitates that units produce energy at their maximum output during almost all hours. Throughout the course of the year, such weeks with high turbine water releases might only be expected to occur when the reservoir is near capacity following the spring runoff. Since the retrospective analysis period is only 5 months long and does include this temporary period of high turbine water releases, it is reasonable to expect that over a 12-month period the potential revenue improvement would be even higher than the 13% achieved in this analysis.

Over the entire analysis period, the baseline results show approximately the same water-to-power conversion efficiency (and thus overall generation) as indicated in the historic record. CHEERS had a slightly higher efficiency at the Hyatt plant and a slightly lower efficiency at the Thermalito plant, but the overall difference is so small as to be within the margin of error of the conversion functions and historical measurement error.

CHEERS Results—Alternative Scenarios. During this historic period of the retrospective analysis, some generating units were unavailable for varying lengths of time. It was of particular interest to CDWR to learn what the potential revenue could have potentially been had those units been in service. Therefore, an alternative *availability* scenario was executed in which some units had a higher modeled availability. The results of this scenario also showed overall generation levels that were virtually unchanged from the historic record, but due to the greater unit availability CHEERS was able to improve revenue an additional 10.3% over the baseline scenario, for an overall improvement of 23.3% over historic revenue. Days exhibiting the greatest revenue improvement had both a high turbine water level and a greater number of available units. In such days the model was able to shift more water and generation from low price hours to high price hours, compared to the baseline scenario in which the fewer available units had to almost constantly release the maximum amount of water possible.

Because the availability scenario was of the greatest interest to CDWR, subsequent alternative scenarios were modifications of the availability scenario rather than of the baseline scenario. Another scenario (energy foresight) was executed that was identical to the availability

scenario, except that the model had perfect foresight of the market clearing prices. Again, the generation efficiency remained virtually unchanged, and the revenue improved by an additional 1.6% relative to the availability scenario. This result indicates the soundness of the price forecasting methodology.

Two other scenarios were executed in which the unit dispatch guidelines were either more or less restrictive than those of the availability scenario. In the restricted scenario, each unit was only permitted to start once per day instead of twice, and the limit on the aggregate change in output over the course of a day was reduced by 50% compared to the availability scenario. In the open scenario, both ramping and startup limits were completely eliminated. Both scenarios again resulted in an overall generation efficiency virtually unchanged from the availability scenario. The restricted scenario suffered only a slight drop in revenue, while the open scenario's revenue was unaffected, indicating that these dispatch guidelines are fairly nonbinding in the situations modeled.

Scenarios were also executed in which the Hyatt and Thermalito units were able to produce ancillary services in addition to energy. When modeling both spinning and non-spinning reserves, adequate water storage in upper reservoirs was maintained above the minimum reservoir elevation in order to accommodate the additional water release required for the possible deployment of contingency reserves. Likewise, the lower reservoir upper elevation limit was decreased such that the possible deployment would not violate the maximum elevation constraint. When providing regulation up and/or regulation down services, it was assumed that the unit's average generation output over the course of an hour would increase or decrease by 15% of the regulation service provided. For example, if a hypothetical unit's energy setpoint was 50 MW in a given hour, and that unit was also providing 20 MW of regulation up service in that hour (but no regulation down service), then the energy output of that unit over the hour was assumed to be 53 MWh. Pursuant with the CAISO market settlement process for energy imbalances, the unit in this case would receive an energy payment for 53 MWh rather than the setpoint level of 50 MWh, in addition to the payment for the 20 MW of regulation up service. The additional water release associated with the regulation up service is also tracked by the model and incorporated in water balance equations.

When the availability scenario is modified to allow the production of ancillary services, overall revenue improves by 23.4% over the availability scenario, for a total improvement of 46.7% over historic revenue. Just as in the availability scenario, forecasted market prices for energy and ancillary services are used to guide the decisions made by CHEERS, but the actual market prices and startup costs are used to determine the revenue improvements cited above. Unlike the previous scenarios examined so far, the ancillary service scenario has significantly lower energy output relative to the historic record, approximately 5.3%. The explicit objective in all of these CHEERS scenarios is to maximize revenue, not to maximize energy production. To that end, there are often opportunities during which ancillary service prices are high enough that it is more advantageous to operate a unit at an output level that is less efficient, but that provides more capacity for a particular high-value service, than it is to operate the unit at its most efficient water-to-power conversion point. In the ancillary service foresight scenario, CHEERS has perfect knowledge of the actual energy and ancillary service prices as opposed to just the

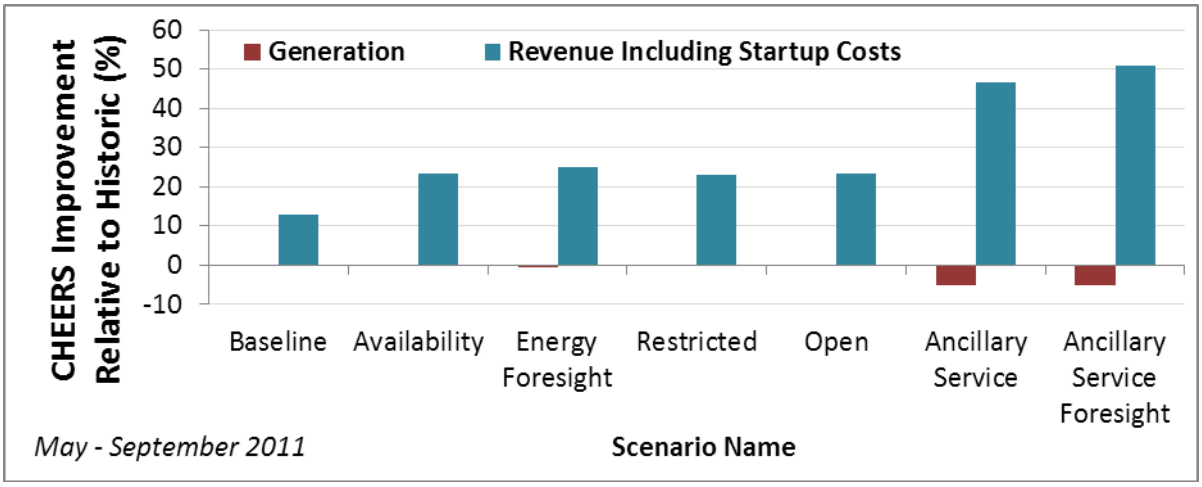

FIGURE 4-121 Relative Improvement in Generation and Revenue of CHEERS Model Scenarios over the Historic Record for CDWR Operations May–September 2011

forecast. Modeled revenue improves to 50.8% over historic amounts. Model results show that having perfect foresight is more valuable when producing ancillary services than when producing only energy, an indication that it is more difficult to predict ancillary service price patterns using the previously described forecasting methodology. Figure 4-121 shows, for each of the scenarios described above, the CHEERS improvement in generation and revenue (including startup costs) relative to the historic record.

CHEERS Results—Integrated Scenarios. Finally, CHEERS was executed using the guidance provided by the other WUOT tools. In the integrated scenario, HydroSCOPE provided model inputs that had been sourced from the problematic historic unit-level data used in the standalone scenarios. Therefore, it was no longer necessary to exclude April from the analysis, and the timeframe was expanded to cover April–September 2011. Inputs provided by HydroSCOPE included the daily initial reservoir surface elevation, evaporation and condensation associated with each of the reservoirs, and the daily water release volume from each of the reservoirs. EHFS provided the system inflow amounts. Due to the long-term nature of environmental objectives relevant to the CDWR system, direct coordination between the IRF and CHEERS tools was not necessary; CHEERS cannot deviate from the daily reservoir release volumes determined by HydroSCOPE, which drive the IRF score, nor can CHEERS violate short-term flow requirements throughout the system.

A critical distinction must be understood between how the daily water release volume was used in the standalone scenarios versus the integrated scenario. In the standalone scenarios, this amount was derived from an examination of the historic unit-level turbine flows. As previously explained, this data was found to be problematic and many records had to be zeroed out. As a result, the amount of water CHEERS was permitted to release through the turbines in the standalone runs was lower than the true amount, but how much lower cannot be determined. Furthermore, since the historic volume of water spilled was not known, the total water release

from the reservoir could not be known. The total release could presumably be back-calculated by examining the reservoir elevations, inflows, and withdrawals over time and calculating a mass water balance, but as previously explained the measurement errors present in those data series make the margin of error in such a calculation so large as to render the result essentially useless. In reality, some unknown amount of water was spilled from the reservoirs during the analysis period, and in the standalone scenarios CHEERS could have increased overall generation had the spilled water amount been known and made available to CHEERS for turbine releases. However, in order to facilitate the fairest comparison possible between the historic record and the standalone scenarios, CHEERS was limited to releasing only the same amount of water through the modeled turbines indicated in the cleaned historic record, and could not use any of the additional turbine water that was missing or eliminated from the historic record; nor could it use any of the water that actually had been spilled.

None of these less-than-ideal circumstances apply in the integrated scenario. Instead, HydroSCOPE determines for CHEERS the amount of water that should be released from each reservoir over the course of the day, and CHEERS is free to release all of that water through either the turbines or bypass routes as it sees fit. The amount of water the historic record indicated was released through the turbines becomes irrelevant, because HydroSCOPE is now creating a “new history” that does not necessarily track what actually happened over time. As a result, CHEERS is able to use the “extra” water to increase generation by 36% and revenue by 48% over the historic record, in a scenario that is related most closely to the availability standalone scenario previously described. On the surface, these numbers seem unrealistically high, but they are due to the caveats explained above, as well as the fact that the historic record against which these results are being compared has not only less turbine water than what was actually used, but less generation as well due to the incomplete and adjusted historic records. Most importantly, the modeled water-to-power efficiency in the standalone case is still essentially the same as that seen in the standalone scenarios. Unfortunately, this overall situation makes it difficult to perform an objective comparison between the results of the CHEERS integrated scenario and the available historic record. However, the described difficulties, which arise from using inconsistent historic data to run the toolset models in a retrospective analysis mode for which they were not designed, should not reflect on the ability of the tools to fulfill their designed purpose. That is, to guide future operations in a computational mode that does not depend on heavy coordination with a historic record.

4.5 REFERENCES

Auble, G.T., M. Wondzell, and C. Talbert, 2009, *Decision Support System for Evaluation of Gunnison River Flow Regimes with Respect to Resources of the Black Canyon of the Gunnison National Park*, U.S. Geological Survey Open-File Report 2009-1126.

Birkhoff, G., M.H. Schultz, and R.S. Varga, 1968, “Piecewise Hermite Interpolation in One and Two Variables with Applications to Partial Differential Equations,” *Numerische Mathematik* 11(3), DOI: 10.1007/BF02161845.

BOR (U.S. Bureau of Reclamation), 1965, *Hydraulic Model Studies of the Oroville Dam Powerplant Intake Structures – California Department of Water Resources – State of California*, Report No. Hyd-509, Division of Research, U.S. Bureau of Reclamation Hydraulics Branch, June 15.

BOR (Bureau of Reclamation), 2003, "Power," Chapter 16 in *Environmental Water Account, Draft Environmental Impact Statement/Environmental Impact Report (EIS/EIR)*, U.S. Department of the Interior, July. Available at http://www.usbr.gov/mp/EWA/docs/DraftEIS-Vol2/Ch16.pdf. Accessed April 23, 2012.

Bui, T., 2013, personal communication from Bui (California Department of Water Resources) to T.S. Lowry (Sandia National Laboratories), Sept. 11.

Busch, D.E., N.L. Ingraham, and S.D. Smith, 1992, "Water Uptake in Woody Riparian Phreatophytes of the Southwestern United States: a Stable Isotope Study," *Ecological Applications* 2:450–459.

CDWR (California Department of Water Resources), undated a, *California State Water Project*. Available at http://www.water.ca.gov/swp/docs/SWPmap.pdf. Accessed April 23, 2012.

CDWR (California Department of Fish and Game), undated b, *Feather River Fish Hatchery Facility*. Available at http://www.dfg.ca.gov/fish/Hatcheries/Feather/Facility.asp. Accessed April 23, 2012.

CDWR (California Department of Water Resources), 1998, *The California Water Plan Update*, Bulletin 160-98, Department of Water Resources.

CDWR (California Department of Water Resources), 2003, *River Flow Effects On Emigrating Juvenile Salmonids In The Lower Feather River, Department of Water Resources*, December. Available at http://www.water.ca.gov/orovillerelicensing/docs/wg_study_reports_and_docs/EWG/12-17-03_env_att_10.pdf. Accessed October 30, 2013.

CDWR (California Department of Water Resources), 2004a, *Matrix of Life History And Habitat Requirements for Feather River Fish Species, SP-F15 TASK 1, SP-F21 TASK 1, Chinook Salmon*, Department of Water Resources, April. Available at http://www.water.ca.gov/orovillerelicensing/docs/wg_study_reports_and_docs/EWG/040528a/04-28-04_fish_chinook_salmon.pdf . Accessed October 30, 2013.

CDWR (California Department of Water Resources), 2004b, *Matrix of Life History And Habitat Requirements for Feather River Fish Species, SP-F15 TASK 1, SP-F21 TASK 1, Steelhead*, Department of Water Resources, April. Available at http://www.water.ca.gov/orovillerelicensing/docs/wg_study_reports_and_docs/EWG/040528a/04-28-04_fish_steelhead.pdf. Accessed October 30, 2013.

CDWR (California Department of Water Resources), 2010a, *Oroville Dam (ORO)*, California Data Exchange Center. Available at http://cdec.water.ca.gov/cgi-progs/profile?s =ORO&type=res. Accessed April 23, 2012.

CDWR (California Department of Water Resources), 2010b, *Thermalito Afterbay (TAB)*, California Data Exchange Center. Available at http://cdec.water.ca.gov/cgi-progs/profile?s =TAB&type=res. Accessed April 23, 2012.

CDWR (California Department of Water Resources), 2012a, *California State Water Project Today*. Available at http://www.water.ca.gov/swp/swptoday.cfm. Accessed April 23, 2012.

CDWR (California Department of Water Resources), 2012b, *Oroville Facilities Complex*. Available at http://www.water.ca.gov/swp/facilities/Oroville/OroFacilitiesComplex.cfm. Accessed April 23, 2012.

CDWR (California Department of Water Resources), 2012c, *California State Water Project: Information about Oroville Facilities*. Available at http://www.water.ca.gov/swp/ facilities/Oroville/LakeDam.cfm. Accessed April 23, 2012.

CDWR (California Department of Water Resources), 2012d, *Edward Hyatt Powerplant*. Available at http://www.water.ca.gov/swp/facilities/Oroville/hyatt.cfm. Accessed April 23, 2012.

CDWR (California Department of Water Resources), 2012e, *Thermalito Facilities*. Available at http://www.water.ca.gov/swp/facilities/Oroville/thermalito.cfm. Accessed April 23, 2012.

CDWR and BOR (California Department of Water Resources and U.S. Bureau of Reclamation), 2012, *Feasibility Investigation of Re-Operation of Shasta and Oroville Reservoirs in Conjunction with Sacramento Valley Groundwater Systems to Augment Water Supply and Environmental Flows in the Sacramento and Feather Rivers*, December. Available at http://www.n-h-i.org/uploads/tx_rtgfiles/NSVCWMP_Report_Final.pdf. Accessed October 30, 2013.

Falaki, F., and S. Loghmanpour, 2009, "Gaining Value by Increasing Efficiency," *HydroWorld.com*, March 1. Available at http://www.hydroworld.com/index/display/article-display/351857/articles/hydro-review/volume-27/issue-1/feature-articles/rehabilitation/gaining-value-by-increasing-efficiency.html. Accessed April 23, 2012.

Fenner, P.R., W.W. Brady, and D.R. Patton, 1984, "Observations on Seeds and Seedlings of Fremont Cottonwood," *Journal of Desert Plants* 6:55–58.

Koczot, K.M., A.E. Jeton, B. McGurk, and M.D. Dettinger, 2005, Precipitation-Runoff Processes in the Feather River Basin, Northeastern California, with Prospects for Streamflow Predictability, Scientific Investigations Report 2004-5202, U.S. Geological Survey.

Mahoney, J.M., and S.B. Rood, 1998, "Streamflow Requirements for Cottonwood Seedling Recruitment—an Integrative Model," *Wetlands* 18(4): 634–645.

McAda, C.W., 2003, *Flow Recommendations to Benefit Endangered Fishes in the Colorado and Gunnison Rivers*, U.S. Fish and Wildlife Service, Grand Junction, CO.

McBain & Trush, Inc. (eds.), 2002, *San Joaquin River Restoration Study Background Report*, prepared for Friant Water Users Authority, Lindsay, CA, and Natural Resources Defense Council, San Francisco, CA.

McCabe, G.J., and L.E. Hay, 1995, "Hydrological Effects of Hypothetical Climate Change in the East River Basin, Colorado, USA," *Hydrological Sciences Journal* 40(3).

Nash, J.E. and J.V. Sutcliffe, 1970, "River Flow Forecasting through Conceptual Models. Part I—A Discussion of Principles," *Journal of Hydrology* 10:282–290.

Nehring, R.B., and D.D. Miller, 1987, "The Influence of Spring Discharge Levels on Rainbow Trout and Brown Trout Recruitment and Survival, Black Canyon of the Gunnison River, Colorado, as Determined by IFIM/PHABSIM Models." *Proceedings of the Western association of Fish and Wildlife Agencies and the Western Division of the American Fisheries Society*.

Poff, N.L., B.D. Richter, A.H. Arthington, S.E. Bunn, R.J. Naiman, E. Kendy, M. Acreman, C. Apse, B.P. Bledsoe, M.C. Freeman, J. Henriksen, R.B. Jacobson, J.G. Kennen, D.M. Merritt, J.H. O'Keeffe, J.D. Olden, K. Rogers, R.E. Tharme, and A. Warner, 2010, "The Ecological Limits of Hydrologic Alteration (ELOHA): a New Framework for Developing Regional Environmental Flow Standards," *Freshwater Biology* 55: 147–170.

Reclamation (U.S. Bureau of Reclamation), 1996, *Record of Decision: Operation of Glen Canyon Dam Final Environmental Statement*, Appendix G, U.S. Department of the Interior, October. Available at http://www.usbr.gov/uc/rm/amp/pdfs/sp_appndxG_ROD.pdf. Accessed April 1, 2010.

Reclamation, 2012a, Final Environmental Impact Statement - Aspinall Unit Operations – Aspinall Unit – Colorado River Storage Project – Gunnison River, Colorado – January 2012, Vols. 1 and 2. Available at http://www.usbr.gov/uc/envdocs/eis/AspinallEIS/index.html. Accessed Dec. 17, 2012.

Reclamation, 2012b, *Record of Decision for the Aspinall Unit Operations Final Environmental Impact Statement – April 2012*. Available at http://www.usbr.gov/uc/envdocs/eis/AspinallEIS/ROD.pdf. Accessed Dec. 17, 2012.

Regonda, S., B. Rajagopalan, M. Clark, and E. Zagona , 2006a, "A Multimodel Ensemble Forecast Framework: Application to Spring Seasonal Flows in the Gunnison River Basin," *Water Resources Research*, 42.

Regonda, S., B. Rajagopalan, and M. Clark, 2006b, "A New Method to Produce Categorical Streamflow Forecast," *Water Resources Research*, 42.

Richter, B.D., J.V. Baumgartner, J. Powell, and D.P. Braun, 1996, "A Method for Assessing Hydrologic Alteration within Ecosystems," *Conservation Biology* 10: 1163–1174.

Sommer, T., D. McEwan, and R. Brown, 2001, *Factors Affecting Chinook Salmon Spawning in the Lower Feather River in Contributions to the Biology of Central Valley Salmonids*. R.L. Brown (ed.), California Department of Fish and Game, Sacramento, California, pp. 269–297.

StateMod, 2012, *State of Colorado's Water Supply Model (StateMod) Version 13.00.00 Users Manual*. Available at http://cdss.state.co.us/software/Pages/StateMod.aspx. Accessed Feb. 27, 2013.

Tao, T., 1999, "Local Inflow Calculator for Reservoirs," *Canadian Water Resources Journal*, 24(1):53–59, DOI: 10.4296/cwrj2401053.

Topozone, 2012, *Topographic Map Dam Features in Butte County, California*. Available at http://www.topozone.com/states/California.asp?county=Butte&feature=Dam. Accessed April 23, 2012.

Trush, W.J., S.M. McBain and L.B. Leopold, 2000, "Attributes of an Alluvial River and their Relation to Water Policy and Management," *Proceedings of the National Academy of Sciences*, 97(22):1158-1163.

Voisin, N., J.C. Schaake, and D.P. Lettenmaier, 2010, "Calibration and Downscaling Methods for Quantitative Ensemble Precipitation Forecasts," *Weather and Forecasting* 25(6):1603–1627.

Voisin, N., F. Pappenberger, D.P. Lettenmaier, R. Buizza, and J.C. Schaake, 2011, "Application of a Medium Range Global Hydrologic Probabilistic Forecast System to the Ohio River Basin," *Weather and Forecasting* 26:425–446.

Yapo, P.O., H.V. Gupta, and S. Sorooshian, 1998, "Multi-objective Global Optimization for Hydrologic Models," *Journal of Hydrology* 204: 83–97.

APPENDIX A:

WUOT NATIONAL LABORATORY TEAM STAFF

Argonne National Laboratory

- John Gasper
- Tom Veselka
- Matthew Mahalik
- Mark Jusko
- John Hayse
- Samrat Saha
- Ronald Black
- Adam Szymanski
- Zhi Zhou
- Canan Uckun
- Sven Leyffer
- Ashutosh Mahajan
- Noam Goldberg

Oak Ridge National Laboratory

- Brennan Smith
- Yetta Jeager

Pacific Northwest National Laboratory

- Mark Wigmosta
- Nathalie Voisin
- Andre Coleman
- Richard Skaggs
- Cindy Rakowski
- Nino Zuljevic

Sandia National Laboratories

- Tom Lowry
- Dirk Vanwestrienen
- Asmeret Bier
- Amy Sun
- Janet Barco
- Daniel Villa
- Will Peplinski
- Scott Griffith
- Marissa Reno-Trujillo

APPENDIX B:

TECHNICAL REVIEW TEAM MEMBERS

California Department of Water Resources

- Tuan Bui
- Molly White
- Glen Solberg
- Aaron Miller
- Maury Miller
- Michael Anderson
- Ryan Wylber
- Jim Wilde
- Twanly Pranger

Western Area Power Administration

- Sam Loftin
- Jennifer Henn
- Jeff Ackerman
- Nancy Scheid
- Chrystal Dean
- Terry Rust
- Chris Brown

U.S. Bureau of Reclamation

- Kerry McCallum
- Toby Steves
- Dave Harpman
- Heather Hermansen
- Mike Pulskampf
- Jane Blair

Southern Company

- Andy Shepard

Brookfield Power

- William Fru
- Tome Mapletoft

Exelon Hydropower

- Ken Poletti
- Paul Adair
- Jay Campbell

PJM Interconnection

- Ed Hoey

APPENDIX C:

ACRONYMS AND ABBREVIATIONS

AF	acre-feet
AMPL	a mathematical programming language
AGC	Automatic Generation Control
AHP	available hydropower
AOP	Annual Operating Plan
Argonne	Argonne National Laboratory
BA	Balancing Authority
C	Celsius
CDEC	California Data Exchange Center
CDWR	California Department of Water Resources
cfs	cubic feet per second
CHEERS	Conventional Hydropower Energy and Environmental Systems
CRSP	Colorado River Storage Project
CRSP MC	CRSP Management Center
DOE	U.S. Department of Energy
DSW	Desert Southwest
EERE	Office of Energy Efficiency and Renewable Energy
EIA	U.S. Energy Information Administration
EMMO	Energy Management and Marketing Office (Western)
ETS	Energy Trading System
F	Fahrenheit
ft	feet
GTMax	Generation and Transmission Maximization (software)
GUI	graphical user interface
GWh	gigawatt-hour(s)
HFC	high-flow channel
hp	horsepower
hr	hour(s)
Hyatt	Edward Hyatt Powerplant
J	Joules
K	Kelvin

L&R	Load & Resource
LFC	low-flow channel
LP	linear programming
LTF	long-term firm
MAF	million acre-feet
MILP	mixed-integer linear program
MINLP	mixed-integer nonlinear program
MSR	Minimum Schedule Requirement
MW	megawatt(s)
MWh	megawatt hour(s)
NERC	North American Electric Reliability Corporation
NP-hard	non-deterministic polynomial-time hard
PGP	pumping-generating powerplant
Reclamation	Bureau of Reclamation
ROD	Record of Decision
rpm	revolutions per minute
SHP	sustainable hydropower
SLCA/IP	Salt Lake City Area Integrated Projects
SOS2	specially ordered set of type 2
SWP	State Water Project (California)
TAB	Thermalito Afterbay
TAF	thousand acre-feet
TDD	Thermalito Diversion Dam
TFB	Thermalito Forebay
THA	Feather River Gauge Station below Thermalito
WALC	Western Area Lower Colorado
WECC	Western Electricity Coordinating Council
Western	Western Area Power Administration
WRP	Western Replacement Power
WSPP	Western Systems Power Pool

APPENDIX D:

DATA REQUIREMENTS

Hydroforecasting
Seasonal Hydrosystems Analysis
Day-Ahead Scheduling and Real-Time Operations

Institutional Arrangements & Functions

Organizations involve with power marketing, ownership, and operations

Long-term planning
Day-ahead scheduling
Real-time dispatch
Plant operation

Other interested parties

Firm customers/ customer organizations
Environmental organizations
Recreation (e.g., river guides, fishing guides)
Public utility commissions
Farming/rancher organizations
Native American tribes
Anthropologists
etc.

System Topology

Demand, supply, market, interchange agreements, etc. (i.e., nodes/objects)
Water channels and energy links (transmission) among demand and supply resources (i.e., links/arches)

Economic/Financial

Capacity market ($/MW-month) and terms/qualification
Spot market prices ($/MWh)
Transaction costs ($/MWh)
Spinning reserve price ($/MW)
Non-spinning reserve price ($/MW)
Replacement reserve price ($/MW)
Regulation services ($/MW)

Contractual Obligations

Long-term firm contract loads (MWh)
Short-term firm contract loads (MWh)
Bilateral purchase contracts (MWh)
Bilateral sales contracts (MWh)
Full-loads service loads (MWh)
Variable resource loads (MWh)

Project use loads (MWh)

<u>Reservoirs (including upper and lower reservoirs for pumped storage)</u>
Location: River basin, state, latitude & longitude, etc.
Initial reservoir state (ft)
Where applicable, historical (e.g., week) hourly reservoir elevations (ft)
Top of inactive pool elevation (ft above sea level)
Top of dead storage elevation (ft above sea level)
Storage (TAF)/elevation (ft) /area (acres) curves
Reservoir rule curves (maximum and minimum reservoir forebay elevation [ft] limits)
Maximum water surface elevation (ft)
Minimum water surface elevation (ft)
Future value of water – i.e., opportunity costs ($/AF)
Maximum weekly release (AF)
Minimum weekly release (AF)
Maximum daily release (AF)
Minimum daily release (AF)
Maximum elevation change over time (e.g., ft/day)
Water temperature (Fahrenheit) as a function of depth (ft)
Specific weight of water (lbs/ft^3)
Hourly evaporation (cfs)
Hourly ± bank storage (cfs)

<u>Dam</u>
Structure height (ft)
Crest elevation (ft above sea level)
Spillway crest (ft)
Spillway capacity (cfs) versus elevation (ft) table
Minimum intake structure elevation(s) (ft)
Non-power outlet works capacity (cfs) versus elevation (ft) table
Streambed at dam axis (ft above sea level)
Where applicable, minimum intake structure elevation for municipal/irrigation (ft)
Multiple input structure for temperature control

<u>Power plant</u>
Topology (turbine, switch, transformer connections)
Maximum transformer capability (MW)
Switch limit (MW)
Ancillary service requirements
- Spinning (MW)
- Non-spinning (MW)
- Replacement reserves (MW)
- Regulation (MW)

Ancillary service arrangements (pooling, formal markets, bilateral agreements, etc.)
Tail water equation – water release rate (cfs) versus tailwater elevation (ft) versus

Special considerations (e.g., must-run, energize transmission line requirement, etc.)

Turbines

Number of turbines
Outage
- Equivalent forced outage rate (%)
- Scheduled maintenance (hours) and timing in/out of service schedule

Initial outage status (in or out of service)
Initial output level (MW)
Where applicable, historical (1-day) outage status (on/off)
Cumulative generation level (MWh) during a designated time period (i.e., past year)
Rated head (ft)
Turbine efficiency curves (% efficiency versus production level [MW])
Turbine name-plant capacity (MW)
Turbine rough zone lower and upper limits (MW)
Minimum and maximum power output levels (MW)
Start-up costs ($)
Shut-down costs ($)
Minimum up time (minutes)
Minimum downtime (minutes)
Wicket gate leakage when closed (cfs)
Maximum turbine release (cfs) as a function of hydraulic head (ft)
Water release requirements for a zero-load spinning state (cfs)
Implication of turbine state transitions over time (e.g., efficiency loss)
Where applicable, unit/water flow interactions

Pumps

Pump/turbine configuration (e.g., reversible turbines or separate pumps)
Outage
- Equivalent forced outage rate (%)
- Scheduled maintenance (hours) and timing in/out of service schedule

Initial outage status (in or out of service)
Initial output level (MW)
Maximum pumping rate (cfs)
Minimum pumping rate (cfs)
Power consumption versus pumping elevation/pumping rate function (MWh/AF)
Start-up costs ($)
Shut-down costs ($)
Minimum up time (minutes)
Minimum downtime (minutes)

Operational Objectives/Preferences/ Guidelines

Objective options – weighted goals:
- Maximize financial benefits
- Maximize economic value

Maximize environmental benefits
Maximize load following

Guidelines:

On-peak & off-peak block purchase preferences
Variable spot market purchase preferences
Equalize turbine generation over time
Maximum number of turbine start-ups & stops per day
Turbine loading order preferences
Minimum profit level for off-peak to on-peak hydropower shifting
etc.

Environment Constraints

Reservoir releases

Hourly minimum (cfs) & averaging time (e.g., second, minute, hour)
Hourly maximum (cfs) & averaging time (e.g., second, minute, hour)
Hourly maximum up-ramp rate (cfs/hr)
Hourly maximum down-ramp rate (cfs/hr)
Maximum daily release (AF)
Minimum daily release (AF)
Maximum daily release change (i.e., daily maximum – daily minimum) (cfs)
Maximum and minimum water release temperature (Fahrenheit)
Water dissolved oxygen content (mg/l) versus tailwater function
etc.

Downstream reaches as measured at specific points (seasonal specific)

Minimum flow (cfs) and channel stage (ft)
Maximum flow (cfs) and channel stage (ft)
Maximum daily flow change (cfs/day) and daily stage change (ft/day)
Maximum and minimum water temperature (Fahrenheit)
Maximum and minimum dissolved oxygen content (mg/l)
HSI and/or WUA functions/relationships
(a) HSI and/or WUA bounds or (b) environmental cost functions

Topology Links

Water channels

Hourly water sideflows between reservoirs (cfs)
Water time-travel distribution functions (WTTD) in % upstream released per hour

Energy/Transmission lines

Hourly total transfer capability (TTC) (MW)
Transmission line loss (%) as a function of power flow (MW)
Minimum hourly power flow (MW)
Maximum hourly power flow (MW)

Power flow costs ($/MW)
Power flow hourly up-ramp costs ($/MW per hour)
Power flow hourly down-ramp costs ($/MW per hour)

Day-ahead Scheduling and Real-time Model Output

Economic/Financial

Firm capacity sale (MW) and revenue ($)
Hourly spot market purchase/sale (MWh) and expenses/revenues ($)
Energy/transmission line flows (MW) and costs ($)
Transaction costs ($)
Total costs for turbine start-ups ($)
Total costs for turbine shut-downs ($)
Total costs for pump start-ups ($)
Total costs for pump shut-downs ($)
Hourly energy consumption costs for pumping ($)
Hourly Spinning reserve purchase/sale (MW) and expenses/revenues ($)
Non-spinning reserve purchase/sale (MW) and expenses/revenues ($)
Replacement reserve purchase/sale (MW) and expenses/revenues ($)
Regulation service purchase/sale (MW) and expenses/revenues ($)
Hourly long-term firm contract sales (MW) & revenue ($), curtailment amount (MW) & penalties ($)
Hourly short-term firm contract sales (MW) & revenue ($), curtailment amount (MW) & penalties ($)
Hourly bilateral purchase amounts (MW) and expenses ($)
Hourly bilateral contract sales (MW) and revenue ($)
Full-loads service load sales (MW) & revenue ($), curtailment amount (MW) & penalties ($)
Variable resource loads sales (MW) & revenue ($), curtailment amount (MW) & penalties ($)
Project use sales (MW) and revenue ($)

Reservoirs /Dams

Hourly reservoir elevation (ft)
Hourly and daily change in reservoir elevation (ft)
Hourly tailwater elevation (ft)
Hourly effective head (ft)
Hourly storage (TAF)
Hourly reservoir area (acres)
Hourly water releases (cfs) - turbine specific, river outlet, spillway, municipal/irrigation outlet, etc.
Weekly marginal value of water ($/AF)
As needed, marginal cost of hourly and daily constraints (e.g., min & max releases)

Turbines/Powerplant
Hourly transformer specific load (MW)
Hourly production level (MW)
Hourly turbine efficiency (MW/cfs)
Number of turbine startups and shutdowns during the simulation period
Hourly wicket gate leakage (cfs)
Hourly operating state (out-of-service, in production, spinning, non-spinning)
Hourly ancillary services
As needed, marginal cost of powerplant constraints

Pumps
Pumping load (MW)
Pumping rate (cfs)
Pump efficiency (MWh/AF)
Number of pump (or pump mode) startups and shutdowns during the simulation period
As needed, marginal cost of pumping constraints

Environment
Reservoir releases
- Hourly ramping (cfs/hr)
- Daily release change (cfs/day)
- Hourly water release temperature (Fahrenheit)
- As needed, marginal cost of reservoir environmental constraints

Downstream reaches as measured at specific points (seasonal specific)
- Hourly dissolved oxygen content (mg/l) of water exiting the tailwater
- Hourly water channels flow rates (cfs) and stage (ft) at designated downstream locations
- Daily flow (cfs/day) and stage change (ft/day) at designated downstream locations
- Water temperature (Fahrenheit) at designated downstream locations
- Dissolved oxygen content (mg/l) at designated downstream locations
- Hourly HSI and/or WUA values at designated downstream locations
- As needed, marginal cost of downstream environmental constraints

Topology Links
As needed, marginal cost of link constraints
Environmental Performance

APPENDIX E:

EQUATION NOMENCLATURE

β	0.47, Bowen ratio for conductive heat exchange (unitless);
γ	62.4, the specific weight of water at 50 degrees Fahrenheit (lbs/cubic foot);
σ	5.6704×10^{-8}, Stefan-Boltzmann constant for back radiation ($W/(m^2K^4)$);
ε	0.97, emissivity of water (unitless);
λ	Light penetration extinction coefficient (1/m);
λ_s	Light penetration extinction coefficient for upper layer (1/m);
π_T^t	Cost of starting up turbine T;
ρ	Water density as a function of temperature (kg/m^3);
$a, b, and\ c$	Calibration parameters for wind function calculation (unitless);
$a_d^{k(E)}$	k^{th}-order coefficient (k = 1,n) for describing the reservoir elevation $\boldsymbol{E}$, as a function of total reservoir volume $\boldsymbol{V}$ at dam $\boldsymbol{d}$;
$a_d^{k(TW)}$	$\boldsymbol{k}^{th}$-order coefficient (k = 1,n) describing the rate of tail water change $\boldsymbol{T}$ as a function of total reservoir release $\boldsymbol{RD}$, at dam $\boldsymbol{d}$;
$a_d^{k(TE)T,H_d^m}$	k^{th} order coefficient describing the turbine efficiency as a function of release water rate at turbine $\boldsymbol{T}$, of dam $\boldsymbol{d}$, for reservoir head H_d^m;
A_{1-6}	Diffusion mixing coefficients;
A_{sT}^t and A_{nT}^t	Market value of spinning and non-spinning reserves;
$b_d^{(E)}$	Intercept of the polynomial function for describing the reservoir elevation $\boldsymbol{E}$.
b_d^{TW}	Intercept of the polynomial function for describing the rate of tail water change.
$b_d^{(TE)T,H_d^m}$	Intercept of the polynomial function for the turbine efficiency.
B_n^t	Quantity of resource entering a network from the boundary node n at time t;

C_{bi}	Withdrawal structure boundary interference coefficient (unitless);
C_c	Cloud cover scale (0-10, 0 = clear, 10 = totally cloudy);
C_{lT}^t and C_{uT}^t	Whether a unit is turned on or off;
C_p	4186.8, specific heat of water (J/(kg*C))
C_{pi}	2093.4, specific heat of water ice (J/(kg*C))
CF^{kW}	737.5, horsepower to kilowatt conversion factor (kW/ft-lbs/sec);
d	Dam index;
D_c	Effective diffusion rate coefficient (m^2/s);
D_{min}	Minimum diffusion rate coefficient (m^2/s);
D_{sT}^t and D_{nT}^t	Level of reserves provided;
e_a	Atmospheric vapor pressure (mmHg);
e_s	Saturation vapor pressure (mmHg);
E_{Vh}	The elevation of the water surface in the reservoir;
E_{dh}^{Res}	Elevation (ft) of the surface of the reservoir at dam ***d***, during hour ***h***;
EF_{dh}^T	Efficiency factor of an individual turbine;
E_{dh}^{TW}	Tail water elevation (ft) below the reservoir at dam ***d***, during hour ***h***;
g	9.81, acceleration of gravity (m/s^2);
G_T^t	Power generated at turbine T at time t,
h	Hour index;
H_{BR}	Back radiation heat exchange rate across water surface (W/m^2);
H_C	Conductive heat exchange rate across water surface (W/m^2);

H_{dh}	Head (ft) or forebay minus tailwater elevation at dam ***d***, during hour ***h***;
H_E	Evaporative heat exchange rate across water surface (W/m^2);
H_{LW}	Long wave heat exchange rate across water surface (W/m^2);
H_{RF}	Heat exchange rate from rainfall across water surface (W/m^2);
H_S	Total heat exchange rate across water surface (W/m^2);
H_{SW}	Short wave heat exchange rate across water surface (W/m^2);
i, j	Node index;
N	Internal buoyancy frequency for inflow and withdrawal dynamics (Hz);
N_T^t	1 if turbine T's status switched to on at time t;
O_T^t	if the turbine T is on at time t;
p	Plant index;
P	Market prices vector;
P^t	Market price per unit of energy at time t,
PF_{dh}	Power factor for the power plant at dam ***d***, during hour ***h*** (fraction);
PQ_{dh}^T	Real power production from turbine *T* of dam ***d***, during hour ***h*** (MW);
Q_{ij}^t	Amount of water released from node i to node j at time t;
Q_{out}	Reservoir outflow rate used in heat balance equation (m^3/s);
Q_z	Vertical flow rate of water for reservoir heat balance (cfs);
R_{dh}^T	Water release rate through turbine ***T*** of dam ***d*** during hour ***h*** (cfs);
R_{SL}	Maximum ramp-down limit;
R_{SU}	Maximum ramp-up limit;
RD	Amount of water released from upstream reservoirs or uncontrolled river inflow;
RS_u^t and RS_{lT}^t	Level of regulation services;

RT	Amount of water released from a reservoir and past turbine blades;
RT_{dh}^{T}	Total release rate through turbine T, reservoir ***d***, during hour ***h***;
RW	Amount of water released from a reservoir and through bypass tubes and spillways;
$(RD_{dh})^{k}$	Total release (AF) from reservoir ***d***, during hour ***h***, raised to the ***k***th power;
S	Storage index;
S_a	Percentage of total incoming solar radiation absorbed (%);
S_r	Total incoming solar radiation (W/m^2);
S_T^t	Quantity of resource stored in a storage node n at the end of time t;
S_{uT}^t and S_{lT}^t	Market value of regulation up and down services, respectively,
SF	Amount of water from side flow;
SOS-2	Specially-Ordered-Set of type 2;
t	General time index;
t_{lag}	River reach time travel lag (s);
T	Turbine index; (CHEERS model only)
T	Water temperature (C); (HydroSCOPE model only)
T_a	Air temperature (C);
T_{aK}	Air temperature in Kelvin (K);
T_d	Air dewpoint temperature (C);
T_{in}	Temperature of boundary water inflow (C);
T_s	Water surface temperature (C);
TE_{dh}^{T}	Efficiency for turbine ***T*** of dam ***d***, during hour ***h*** (fraction);
V	Reservoir index;

$\cancel{V}$	Reservoir volume in heat balance equation (ft^3);
$(V_{dh})^k$	Volume of water (AF) at reservoir ***d***, during hour ***h***, raised to the ***k***th power;
W_f	Wind function value used in conductive and evaporative heat exchange calculations ($J/(C*m^2)$);
W_{uT}^t and W_{lT}^t	Turbine startup and shutdown cost;
Z	Net economic value of system operations over the optimization period.

APPENDIX F:

ADAPTIVELY REFINED DYNAMIC PROGRAM FOR LINEAR SPLINE

ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, IL 60439

Adaptively Refined Dynamic Program for Linear Spline Regression

Noam Goldberg, Youngdae Kim, Sven Leyffer, and Thomas Veselka

Mathematics and Computer Science Division

Preprint ANL/MCS-P3040-0912

September 25, 2012

This work was supported by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357.

Adaptively Refined Dynamic Program for Linear Spline Regression

Noam Goldberg* Youngdae Kim† Sven Leyffer‡ Thomas D. Veselka§

September 25, 2012

Abstract

The linear spline regression problem is to determine a piecewise linear function for estimating a set of given points while minimizing a given measure of misfit or error. This is a classical problem in computational statistics and operations research; dynamic programming was proposed as a solution technique more than forty years ago by Bellman and Roth (1969). The algorithm requires a discretization of the solution space to define a grid of candidate breakpoints. This paper proposes an adaptive refinement scheme for the grid of candidate breakpoints in order to allow the dynamic programming method to scale for larger instances of the problem. We evaluate the quality of solutions found on small instances compared with optimal solutions determined by a novel integer programming formulation of the problem. We also consider a generalization of the linear spline regression problem to fit multiple curves that share breakpoint horizontal coordinates, and we extend our method to solve the generalized problem. Computational experiments verify that our nonuniform grid construction schemes are useful for computing high-quality solutions for both the single-curve and two-curve linear spline regression problem.

Keywords: Piecewise regression, least squares, change point detection, dynamic programming, mixed integer programming.

1 Problem Description and Survey of Literature

The linear spline regression problem, given a dataset $D = (x\ y) \in \mathbb{R}^{n\times 2}$, is to determine a continuous piecewise linear function with m line segments that best fits the data as measured by a given error function $e : D \to \mathbb{R}$. The data consists of an *independent variable* vector $x \in \mathbb{R}^n$ and a *dependent variable* vector $y \in \mathbb{R}^n$; without losing generality it is assumed that $x_1 \le \cdots \le x_n$. Continuity implies that every pair of incident line segments must intersect at a breakpoint. Also, note that the methods herein generalize for the case that x is a matrix but there is a single ordered dimension (e.g., time) in which the function is piecewise linear with multiple pieces and otherwise the function is linear.

The estimated function has a set of $m+1$ breakpoints denoted by $b = [b_{ij}] \in \mathbb{R}^{(m+1)\times 2}$. The breakpoints must be determined within a domain of interest $I \subset \mathbb{R}$; typically $I = [x_1, x_n]$ or $I = [0, x_n]$. The error over the line segments for given data D is measured by one of the following:

Absolute error:
$$e_1(D,b) = \sum_{j=1}^{m} \sum_{\substack{i=1,\ldots,n:\\ b_{j1}\le x_i\le b_{j+1,1}}} \left| y_i - \left(\frac{b_{j+1,2}-b_{j,2}}{b_{j+1,1}-b_{j1}} (x_i - b_{j1}) + b_{j1} \right) \right|. \tag{1}$$

Sum of squared error:
$$e_2(D,b) = \sum_{j=1}^{m} \sum_{\substack{i=1,\ldots,n:\\ b_{j1}\le x_i\le b_{j+1,1}}} \left[y_i - \left(\frac{b_{j+1,2}-b_{j2}}{b_{j+1,1}-b_{j1}} (x_i - b_{j1}) + b_{j1} \right) \right]^2. \tag{2}$$

Maximum of absolute errors:
$$e_\infty(D,b) = \sum_{j=1}^{m} \max_{\substack{i=1,\ldots,n:\\ b_{j1}\le x_i\le b_{j+1,1}}} \left| y_i - \left(\frac{b_{j+1,2}-b_{j2}}{b_{j+1,1}-b_{j1}} (x_i - b_{j1}) + b_{j1} \right) \right|. \tag{3}$$

*Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, noamgold@mcs.anl.gov.
†Department of Computer Sciences, University of Wisconsin-Madison, youngdae@cs.wisc.edu.
‡Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, leyffer@mcs.anl.gov.
§Decision and Information Sciences, Argonne National Laboratory, Argonne, IL 60439, tdveselka@anl.gov.

Bellman and Roth [3] proposed a dynamic programming algorithm for minimizing (3) over a given (uniform) grid of candidate breakpoints. Guthery [6] proposed an efficient dynamic program to fit a segmented curve that is not required to be a continuous function in order to minimize (2). Ertel and Flowlkes [4] developed heuristics to solve both problems. More recently, in the econometrics community, Bai and Perron [2] revisited the problem considered by Guthery. Also, Aronov et al. [1] suggest an approximation algorithm for the problem considered herein and similar to that of [3] with a runtime complexity that is a high degree polynomial in the number of data points; hence it may not be efficient in practice. The problem's computational complexity remains an open problem according to [1]. In recent work, Toridello and Viemla [8] proposed integer programming formulations for a related curve-fitting problem.

In the following section we propose an integer programming formulation for the linear spline regression problem. In Section 3 we give an overview of the dynamic programming method. In Section 4 we consider two approaches to constructing tractable grids of candidate breakpoints for applying the dynamic program: the first is a static scheme, and the second is a dynamic scheme that adaptively refines the grid based on a current solution. We conclude with numerical experiments in Section 6.

2 A Mixed Integer Programming Formulation

In this section, we consider a mixed-integer programming (MIP) formulation for the problem. The following tables define the variables and constants of the formulation.

Variable	Meaning
β_j	The slope of the j-th line segment, i.e., $\beta_j = \frac{b_{j2}-b_{j-1,2}}{b_{j1}-b_{j-1,1}}$.
α_j	The intercept of the j-th line segment, i.e., $\alpha_j = b_{j-1,2}+\beta_j(b_{j1}-b_{j-1,1})$.
ϕ_{ij}	Indicating whether the error of the i-th data point should should contribute to the error of the jth line segment, i.e., $\phi_{ij} = \begin{cases} 0 & \text{if } b_{j-1,1} \le x_i \le b_{j1} \\ 1 & \text{otherwise} \end{cases}$
ξ_i	$\xi_i = \lvert y_i - (\beta_j x_i + \alpha_j)\rvert$, the difference between y value of the ith data point and the estimated y value of the line segment.

Constant	Meaning
M_1, M_2, M_3	Large ("big-M") constants – to satisfy the associated inequalities for $i = 1,\dots,n$. For example, it suffices to set $M_1 \ge \max_{\substack{i,j,k=1,\dots,n \\ i\ne j}} \left\lvert y_k - y_j - \frac{y_i-y_j}{x_i-x_j}(x_k-x_j)\right\rvert$.

We first consider a nonconvex nonlinear formulation for the general problem:

$$\underset{\beta,\phi,\xi,b}{\text{minimize}} \quad \sum_{i=1}^{n} \xi_i^2 \tag{4a}$$

$$\text{subject to} \quad |y_i - b_{j2} - \beta_j x_i| \le \xi_i + M_1\phi_{ij} \qquad i=1,\dots,n, \quad j=1,\dots,m \tag{4b}$$

$$\sum_{j=1}^{m} \phi_{ij} = m-1 \qquad i=1,\dots,n \tag{4c}$$

$$-M_2\phi_{ij} \le b_{j1} - x_i \qquad i=1,\dots,n, \quad j=1,\dots,m \tag{4d}$$

$$-M_2\phi_{ij} \le x_i - b_{j-1,1} \qquad i=1,\dots,n, \quad j=2,\dots,m \tag{4e}$$

$$b_{j1} = \frac{b_{j+1,2}-b_{j2}}{\beta_j - \beta_{j+1}} \qquad j=1,\dots,m-1 \tag{4f}$$

$$\phi_{ij} \in \{0,1\} \qquad i=1,\dots,n, \quad j=1,\dots,m. \tag{4g}$$

Here we minimize the error loss function $e_2(D,\cdot)$. Note that by replacing the objective function with other loss functions such as (1) or (3), we can instead consider, respectively, the minimization of the sum of absolute errors, or the minimization of the maximum of absolute errors. The constraints (4b) imply that for $i = 1,\dots,n$, ξ_i measures the deviation of y_i and $\alpha_j + \beta_j x_i$ for the segment j to which data point i is assigned; in other words, if $\phi_{ij} = 0$, then the constraints (4b) require that $\xi_i = |y_i - (\beta_j x_i + \alpha_j)|$. Otherwise,

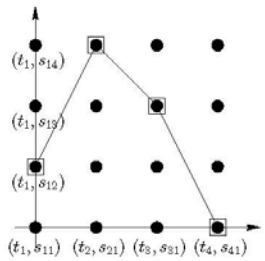

Figure 1: Example of a uniform grid $\Omega = \bigcup_{i=1}^{4} \bigcup_{j=1}^{4} \{(t_i, s_{ij})\}$.

the constraints are automatically satisfied by using the sufficiently large constant M_1. The constraints (4d)-(4e) force each data point to be assigned to exactly one of the m line segments and its error to be respectively evaluated by using that line segment.

Assuming that the fitted curve is either concave or convex, we can rewrite (4) as a convex mixed integer quadratic program (MIQP). The MIQP formulation is obtained from (4) by substituting the b_{j1} variables in (4d) and (4e) and replacing M_2 by a constant M_3 satisfying $M_3 > (\beta_j - \beta_{j+1})M_2$ for $j = 1, \ldots, m$; if the function to be estimated is Lipschitz with constant L, then we may let

$$M_3 = L \cdot M_2 \geq \max_{j=1,\ldots,m} \beta_j M_2 \geq \max_{j=1,\ldots,m} \{\beta_j - \beta_{j+1}\} M_2.$$

The resulting formulation is as follows:

$$\begin{array}{llll}
\underset{\alpha,\beta,\phi,\xi}{\text{minimize}} & \displaystyle\sum_{i=1}^{n} \xi_i^2 & & \text{(5a)} \\
\text{subject to} & |y_i - (\beta_j x_i + \alpha_j)| \leq \xi_i + M_1 \phi_{ij} & \text{for } i = 1, \ldots, n, \;\; j = 1, \ldots, m & \text{(5b)} \\
& \displaystyle\sum_{j=1}^{m} \phi_{ij} = m - 1 & \text{for } i = 1, \ldots, n & \text{(5c)} \\
& -M_3 \phi_{ij} \leq \alpha_j - \alpha_{j+1} - (\beta_{j+1} - \beta_j) x_i & \text{for } i = 1, \ldots, n, \;\; j = 1, \ldots, m-1 & \text{(5d)} \\
& -M_3 \phi_{ij} \leq \alpha_j - \alpha_{j-1} + (\beta_j - \beta_{j-1}) x_i & \text{for } i = 1, \ldots, n, \;\; j = 2, \ldots, m & \text{(5e)} \\
& \beta_{j+1} \leq \beta_j & \text{for } j = 1, \ldots, m-1 & \text{(5f)} \\
& \phi_{ij} \in \{0, 1\} & \text{for } i = 1, \ldots, n, \;\; j = 1, \ldots, m. & \text{(5g)}
\end{array}$$

The constraints (5d)-(5e) follow from (4d)-(4e) given that $\beta_j \geq \beta_{j+1}$ for $j = 1, \ldots, m$ (when the function to be estimated is a concave piecewise linear function). Note that if instead of (5f) we required that $\beta_j \leq \beta_{j+1}$ for all j and replaced the inequalities in (5d)-(5e) with the reverse inequalities, then (5) could be used to estimate a convex piecewise linear function. The constraint (5f) ensures that the resulting continuous piecewise linear function is concave.

Note that an MIQP formulation for a related problem has been proposed by Toriello and Vielma [8]. A key difference between their formulation and (5) is that it does not require the estimated curve to correspond to a function over the domain defined by the data points; in other words, in [8] the error is computed as a function of the vertical distances from the data points to the closest line segments, whereas in our case a set of points between adjacent breakpoints must always be assigned to the same line segment. We use formulation (5) as a benchmark for numerical results in Section 6.

Next we review the dynamic programming method for solving this problem.

3 Dynamic Programming for Linear Spline Regression

Bellman and Roth [3] proposed dynamic programming (DP) for determining a continuous piecewise linear function that minimizes the error function $e_\infty(\cdot)$. The Bellman optimality equations of the DP are expressed for a given grid system consisting of a finite number of candidate breakpoints over which the error is minimized [3]. Figure 1 shows an example of the uniform grid system considered in [3] along with a piecewise linear solution curve with three line segments.

We now write the Bellman equations of the linear spline regression problem DP [3]. First, let $\Omega \in \mathbb{R}^2$ denote a finite set of candidate breakpoints. For a positive integer p let $t \in \mathbb{R}^p$ denote the vector of breakpoint x-coordinates of a grid Ω, with $t_1 < \cdots < t_p$. Define $S : \mathbb{R} \to 2^n$, where $S(\bar{x})$ is the set of indexes that share the same x-coordinate value $\bar{x}$. For example, in Figure 1, $S(t_1) = \{1, 2, 3, 4\}$. For convenience we also write the breakpoints as $\Omega = \bigcup_{i=1}^{p} \{(t_i, s_{ij}) \mid j \in S(t_i)\} \in \mathbb{R}^2$; we assume that $s_{i1} < \cdots < s_{i|S(t_i)|}$ for $i = 1, \ldots, p$.

For $k, l \in \{1, \ldots, p\}$ with $t_k < t_i$, $l \in S(t_k)$ and $j \in S(t_i)$, we denote the error of segment $(t_k, s_{kl}), (t_i, s_{ij})$ by $E(t_k, s_{kl}, t_i, s_{ij})$; for the squared error function $e_2(\cdot)$ it is given by

$$E(t_k, s_{kl}, t_i, s_{ij}) = \sum_{\substack{r=1,\ldots,n: \\ t_k \le x_r \le t_i}} \left\{ y_r - \left(\frac{s_{ij} - s_{kl}}{t_i - t_k} (x_r - t_k) + s_{kl} \right) \right\}^2 .$$

Let $F_{\widehat{m}}(i, j)$ denote the optimal error of $\widehat{m}$ line segments when the right breakpoint of the last segment is (t_i, s_{ij}). Then, letting $\widehat{m} \in \{2, \ldots, m\}$, the optimality equations of the DP are given by:

$$F_{\widehat{m}}(i, j) = \min_{k=\widehat{m},\ldots,i-1, l \in S(t_k)} \{ E(t_k, s_{kl}, t_i, s_{ij}) + F_{\widehat{m}-1}(k, l) \} \text{ for } i = \widehat{m}+1, \ldots, p-m+\widehat{m}, j \in S(t_i). \tag{6a}$$

$$F_1(i, j) = \min_{l \in S(t_1)} E(t_1, s_{1l}, t_i, s_{ij}) \text{ for } i = 2, \ldots, p-m+1, j \in S(t_i). \tag{6b}$$

The method requires the evaluation of $F_1(\cdot), \ldots, F_m(\cdot)$ to determine an optimal continuous piecewise linear function with m line segments given Ω. To obtain the optimal solution value using (6), we solve for

$$\min_{j \in S_{t_p}} F_m(p, j). \tag{7}$$

To determine the optimal solution, we must also store the breakpoints $(t_k, s_{k\ell})$ at which the minimum of $F_{\widehat{m}}(t_i, s_{ij})$ is attained for $\widehat{m} = 1, \ldots, m$, for each $i \in \{1, \ldots, p\}$ and each $j \in S(t_i)$.

4 Constructing a Tractable Grid for Dynamic Programming

The density of the grid directly affects the performance of the DP method. A denser grid with many candidate breakpoints yields lower error values, but this result is at the expense of significant degradation of runtime performance. Table 1 shows that lower error values could be achieved by increasing the number of candidate breakpoints, although, at the expense of significant increases in the runtime. For example, with a total of 10,000 candidate breakpoints, we could lower the error value by 1.9% compared with 500 candidate breakpoints. However, the runtime increased nearly 400-fold.

Table 1: Running time vs. lower error values for the UCI MPG dataset [5] with $n = 398$ and $m = 5$.

$\lvert\Omega\rvert$	Time (sec)	Objective Value
500	0.056	6,830.611
1,000	0.201	6,812.490
5,000	4.981	6,718.609
10,000	19.917	6,697.479

If the grid has $|S(t_i)| = q$ for $i = 1, \ldots, p$, then letting $N = |\Omega| = pq$, the number of candidate breakpoints is $O(N)$. If the segment errors are precomputed, then the DP running time is bounded by $O(N^2 m)$; for a detailed analysis see Appendix A. Hence, it is essential to construct a grid that will be tractable in terms of running time satisfactory in terms of the quality of the solutions.

Bellman and Roth [3] do not consider a particular grid system. In the example introduced in [3], a uniform grid is assumed, as depicted in Figure 1. For low-accuracy solutions a static uniform grid may be adequate. In order to improve on the accuracy, however, a finer grid may be required, which may not be tractable using the uniform grid approach.

Figure 2 shows that there are many cases in which the range of useful y-coordinate values of candidate breakpoints can be significantly narrowed down to lie within close proximity to the given data. Based on this observation we suggest two types of techniques for constructing (nonuniform) grid systems. In Section 4.1, we describe a grid system which generates a fixed number of candidate breakpoints near data points. In Section 4.2 we consider an adaptive construction of a grid. First we solve using a coarse grid. Then, we generate additional candidate breakpoints that lie within a neighborhood of the current solution, and we iterate until a termination condition is satisfied.

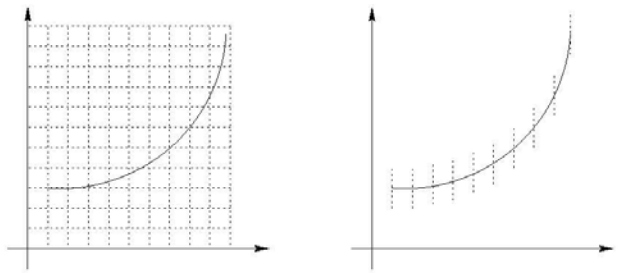

Figure 2: Uniform grid vs. nonuniform grid. In the uniform grid, many of the candidate breakpoints do not help much in approximating the given curve, depicted as a solid line. In the nonuniform grid, however, we generate candidate breakpoints that are expected to improve the approximation of this curve.

4.1 DP-static: Placing a Fixed Number of Candidate Breakpoints near the Data Points

Here we propose a scheme for constructing a nonuniform grid of candidate breakpoints; we call this scheme DP-static. The input of the DP-static scheme consists of the dataset D, an integer p, and an odd integer q.

DP-static generates p equally spaced x-coordinate values $t_1 < \cdots < t_p$ that subdivide the interval $[\min_i x_i, \max_i x_i]$ into $(p-1)$ sub-intervals. The values $t_1, \ldots, t_p$ constitute the x-coordinate values of the candidate breakpoints.

For each t_i we determine

$$k \in \operatorname*{argmin}_{r=1,\ldots,n} \{|x_r - t_i| \mid x_r \le t_i\} \qquad \text{and} \qquad \ell \in \operatorname*{argmin}_{r=1,\ldots,n} \{|x_r - t_i| \mid t_i \le x_r\}.$$

Then, for $i = 1, \ldots, n$ we compute $s_{i1}, \ldots, s_{iq}$. Since q is odd, it follows that $\lceil \frac{q}{2} \rceil$ is the median index of $s_{i1}, \ldots, s_{iq}$. For $c = \frac{\max_i y_i - \min_i y_i}{2n}$ and $j = 1, \ldots, q$ we let

$$s_{ij} = s_{i\lceil \frac{q}{2} \rceil} + \left(j - \left\lceil \frac{q}{2} \right\rceil\right) c.$$

Note that $s_{i\lceil \frac{q}{2} \rceil} = \frac{y_k + y_l}{2}$. The resulting breakpoints are the $\bigcup_{i=1,\ldots,p} \{(t_i, s_{i1}), \ldots, (t_i, s_{iq})\}$. An example of the resulting grid is depicted on the right of Figure 2.

4.2 DP-adaptive: Adaptive Refinement

The statically determined grid may not be adequate when the data points do not lie close enough to the optimal breakpoints. When a region of data space requires a larger number of candidate breakpoints in order to obtain adequate solutions, it may be useful to employ an iterative procedure based on the current solution. Also, through an adaptive refinement scheme it may be possible to attain competitive error values by using a sparser grid and smaller number of candidate breakpoints.

Based on this intuition, we propose an adaptive refinement scheme. The algorithm is initialized with a coarse grid. At each iteration, we apply the DP algorithm to the current grid and adaptively refine only those rectangles that intersect the DP solution curve. Algorithm 1 describes our simple procedure for refining a subset of the rectangles, subdividing each into four subrectangles by halving each one of its sides. Figure 3 shows an example that applies a few steps of the adaptive refinement algorithm. At the first iteration we start with a large rectangle with four candidate breakpoints (denoted by circles) corresponding to the rectangle's vertices. After applying the DP, we refine the rectangles that intersect the line segments, by creating more candidate breakpoints at the center of the rectangles as well as midpoints of each rectangle's sides. The algorithm repeats the refinements step until a termination criterion is met. Algorithm 2 implements the DP-adaptive scheme.

We use $\{[(t_k, s_{kl}), (t_i, s_{ij})]\}$ with $t_k < t_i$ and $s_{kl} < s_{ij}$ to denote a rectangle with bottom left point (t_k, s_{kl}) and top right point (t_i, s_{ij}). Algorithm 2 is typically initialized with a rectangle set consisting of a single rectangle:

$$R \leftarrow \left\{ \left[(\min_{i=1,\ldots,n} x_i, \min_{i=1,\ldots,n} y_i), (\max_{i=1,\ldots,n} x_i, \max_{i=1,\ldots,n} y_i) \right] \right\}. \tag{8}$$

In the following for a finite $C = \{p_1, p_2, \ldots\} \subseteq \mathbb{R}^2$ assume $p_{11} \le p_{21} \le \cdots \le p_{|C|1}$, and define

$$\mathcal{L}(C) = \{\lambda p_i + (1-\lambda) p_{i+1} \mid i = 1, \ldots, |C|, 0 \le \lambda \le 1\};$$

$\mathcal{L}(C)$ corresponds to the line segments that connect the points in C that are pairwise adjacent along the x-axis.

Algorithm 1 refineGrid($R, \hat{R}, \Omega$), refine a subset of rectangles of a given grid

Input: A set of rectangles R, current grid Ω, subset of rectangles $\widehat{R}$ to be refined

for each $[(t_k, s_{kl}), (t_i, s_{ij})] \in \widehat{R}$ **do**

  $r_1 \leftarrow [(t_k, s_{kl}), \left(\frac{t_k+t_i}{2}, \frac{s_{kl}+s_{ij}}{2}\right)]$, $r_2 \leftarrow [\left(t_k, \frac{s_{kl}+s_{ij}}{2}\right), \left(\frac{t_k+t_i}{2}, s_{ij}\right)]$

  $r_3 \leftarrow [\left(\frac{t_k+t_i}{2}, \frac{s_{kl}+s_{ij}}{2}\right), (t_i, s_{ij})]$, $r_4 \leftarrow [\left(\frac{t_k+t_i}{2}, s_{kl}\right), \left(t_i, \frac{s_{kl}+s_{ij}}{2}\right)]$.

  $R \leftarrow R \cup \{r_1, r_2, r_3, r_4\} \setminus \{[(t_k, s_{kl}), (t_i, s_{ij})]\}$

  $b_1 \leftarrow \left(t_k, \frac{s_{kl}+s_{ij}}{2}\right)$, $b_2 \leftarrow \left(\frac{t_k+t_i}{2}, s_{kl}\right)$, $b_3 \leftarrow \left(\frac{t_k+t_i}{2}, \frac{s_{kl}+s_{ij}}{2}\right)$, $b_4 \leftarrow \left(\frac{t_k+t_i}{2}, s_{ij}\right)$

  $b_5 \leftarrow \left(t_i, \frac{s_{kl}+s_{ij}}{2}\right)$.

  $\Omega \leftarrow \Omega \cup \{b_1, b_2, b_3, b_4, b_5\}$.

end for

Output: R, Ω.

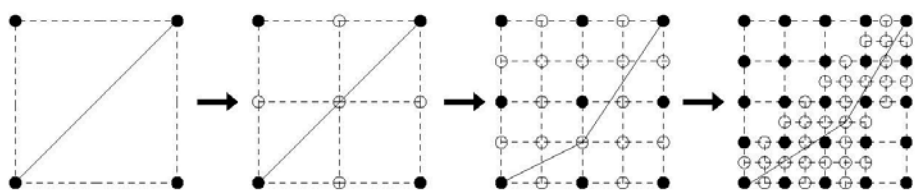

Figure 3: Adaptive refinement steps to find a continuous piecewise linear function with two line segments. Unfilled circles denote newly added candidate breakpoints at each iteration. The graph of the piecewise linear function in each step consists of the line segments connecting the circles. The curve minimizes the error with respect to the current grid consisting of both filled and unfilled circles. For the sake of clarity, the data points are not displayed in this diagram.

Algorithm 2 Adaptive refinement for a single curve (DP-adaptive)

Input: Dataset $D \subset \mathbb{R}^2$, the number of line segments m, iteration limit T,
  initial grid $\Omega \subset \mathbb{R}^2$, and rectangle set R

for $k = 1, \ldots, T$ **do**

  $m' \leftarrow \min\{m, k\}$

  Evaluate (7) on Ω and let $C \subseteq \Omega$ denote the optimal solution.

  $\widehat{R} \leftarrow \{[(t_i, s_{ij}), (t_k, s_{kj})] \in R \mid \operatorname{conv}\{(t_i, s_{ij}), (t_i, s_{kj}), (t_k, s_{ij}), (t_k, s_{kj})\} \cap \mathcal{L}(C) \neq \emptyset\}$

  $(R, \Omega) \leftarrow$ refineGrid$(R, \hat{R}, \Omega)$

end for

Output: C

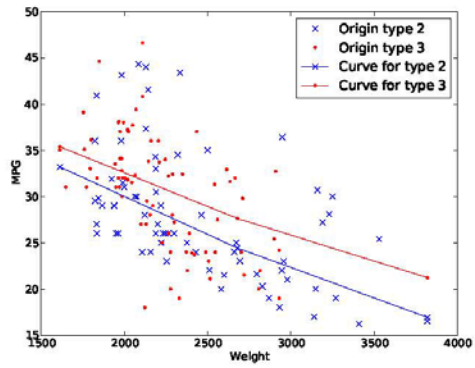

Figure 4: Least squares fit of two piecewise linear functions to a partition of the UCI MPG data [5] into domestic and foreign made cars. Note that each breakpoint of the two piecewise linear functions share the same x-coordinate value.

5 Extension of the Problem: Fitting Multiple Curves that Share a Breakpoint Coordinate

The data points may be associated with one or more categories. For example, as in Figure 4 one may consider domestic vs. foreign vehicles for estimating gas consumption. To each category we may choose to fit a distinct curve. However, the breakpoints' horizontal coordinates, also known as change-points or timestamps for time series data [7], may be required to be equal over all categories. For example, with financial data the change-point may correspond to an event such as a market crash, and each category may correspond to a different financial index.

In the following, we consider an extension of the dyanamic program and grid construction methods for multiple curves. To simplify the presentation, we limit the current discussion to two curves, although Algorithm 3 describes the general method for $v \geq 2$ curves. Let n^w denote the number of data points in part w of the data, for $w \in W \stackrel{\text{Def}}{=} \{1, \ldots, v\}$, and so that $\sum_{w \in W} n^w = n$. Denote the corresponding data $D^w = (x^w \;\; y^w) \in \mathbb{R}^{n^w \times 2}$ for $w \in W$. As for the single curve case, a grid of candidate breakpoints is constructed in order to apply the DP method. We generate distinct sets of candidate breakpoints, each of which is used to approximate one of the curves. We follow the same procedure to generate the candidate breakpoints for each curve as described in Sections 4.1 and 4.2 except that the candidate breakpoints of each curve must share the same x-coordinate values.

5.1 Dynamic Programming for Multiple Curves

Now we write the Bellman optimality equations for the multiple-curve problem. For a curve $w \in W$ and each pair of x-coordinate values (t_k, t_i) with $t_k < t_i$, we have a v-tuples of line segments

$$\left\{ \times_{w \in W} (t_k, s^w_{k\ell}, t_i, s^w_{ij}) \;\mid\; \ell \in S^w(t_k), j \in S^w(t_i) \right\};$$

for each index pair of x-coordinates (k, i) there are $\prod_{w \in W} |S^w(t_k)| \times |S^w(t_i)|$ line segments. For example, with $v = 2$ there are $\left(|S^1(t_k)| \times |S^1(t_i)|\right)\left(|S^2(t_k)| \times |S^2(t_i)|\right)$ possible combinations of line segments. For convenience, for a pair of indices $(k, i) \in \{1, \ldots, p\} \times \{1, \ldots, p\}$ with $k \neq i$, $l_w \in S^w(t_k)$ and $j_w \in S^w(t_i)$ for each $w \in W$, we write a v-tuple of segments as $(t_k, \times_{w \in W} s^w_{kl_w}, t_i, \times_{w \in W} s^w_{ij_w})$. The segment v-tuple error is then defined as

$$E(t_k, \times_{w \in W} s^w_{kl_w} t_i, \times_{w \in W} s^w_{ij_w}) = \sum_{w \in W} E(t_k, s^w_{kl_w}, t_i, s^w_{ij_w}).$$

For $\widehat{m} = 2, \ldots, m$ the optimality equations of the generalized multiple curve problem are

$$F_{\widehat{m}}(i,j) = \min_{\substack{k=\widehat{m},\ldots,i-1,\\ l_w \in S^w(t_k):\\ w\in W}} \left\{ E(t_k, \times_{w\in W} s^w_{l l_w}, t_i, \times_{w\in W} s^w_{i j_w}) + F_{\widehat{m}-1}(k, \times_{w\in W} l_w) \right\} \quad \text{for } i = \widehat{m}+1, \ldots, p-m+1, \quad j_w \in S^w(t_i) : w \in W. \tag{9a}$$

$$F_1(i,j) = \min_{l_w \in S^w(t_1): w\in W} E(t_1, \times_{w\in W} s^w_{1 l_w}, t_i, \times_{w\in W} s^w_{i j_w}) \quad \text{for } i = 2, \ldots, p-m+1, \quad j_w \in S^w(t_i) : w \in W. \tag{9b}$$

If $S^w(t_i) \le q$ for $w \in W$ and $i = 1, \ldots, p$, then the running time of evaluating (9) is $O(p^2 q^{2v} m)$. The running time analysis and proof are given in Appendix A.

5.2 Grid Construction for Multiple Curves

For DP-static we determine a single set of t_i's by subdividing the entire range of data (including both sets) into $(q-1)$ subintervals of equal lengths. To generate a distinct set $S^w(t_i)$ for each i and each curve w, we apply a procedure similar to the one described in Section 4.1 and using the same rectangle refinement subroutine given by Algorithm 1.

Algorithm 3 describes the extension of Algorithm 2 to the case of multiple curves. For the initial grid we typically use four

Algorithm 3 Adaptive refinement for multiple curves

Input: Datasets $D_w \subset \mathbb{R}^2$, for $w \in W$, the number of line segments m, iteration limit T, initial grid Ω_w, and rectangles R_w for $w \in W$.

1: **for** $k = 1, \ldots, T$ **do**
2: $m' \leftarrow \min\{m, k\}$
3: Evaluate $\min_{j \in \times_{w\in W} S^w(t_{2^{k-1}})} F_{m'}(2^{k-1}, j)$ on $\Omega_1, \ldots, \Omega_v$ and let $C_1 \subseteq \Omega_1, \ldots, C_v \subseteq \Omega_v$ denote the optimal solutions.
4: $\widehat{R}_w \leftarrow \{[(t_i, s_{ij}), (t_k, s_{kj})] \in R_w \mid \operatorname{conv}\{(t_i, s_{ij}), (t_i, s_{kj}), (t_k, s_{ij}), (t_k, s_{kj})\} \cap \mathcal{L}(C_w) \neq \emptyset\}$ for $w \in W$
5: refineGrid$(R_w, \hat{R}_w, \Omega_w)$ for $w = 1, \ldots, v$
6: **end for**

Output: $C_w \in \Omega_w$ for $w \in W$.

candidate breakpoints for each Ω_w for $w = 1, \ldots, v$. Accordingly, for $w = 1, \ldots, v$, we initialize

$$R_w \leftarrow \left\{ \left[(\min_{w'=1,\ldots,v} \min_{i=1,\ldots,n^w} x_i^w, \min_{i=1,\ldots,n^w} y_i^w), (\max_{w'=1,\ldots,v} \max_{i=1,\ldots,n^w} x_i^w, \max_{i=1,\ldots,n^w} y_i^w) \right] \right\}.$$

After the initialization, the algorithm follows the same procedure described in Algorithm 2 for refinement while maintaining the distinct sets of candidate breakpoints for each curve. As the two curves are initialized with a single set of t_i's and the refinement maintains a single set of t_i's for both curves thereafter, it follows that the breakpoints of the two solution curves will share the same set of x-coordinate values.

Algorithm 3 implements the adaptive refinement procedure for multiple curves. In particular, with $v = 2$, we have an adaptive refinement algorithm for two curves. The computational experiments of the current paper, discussed in Section 6, focus on the cases of $v = 1, 2$.

6 Numerical Results

In this section we evaluate the performance of our proposed techniques. We experimented with six datasets including standard benchmark UCI data [5] and hydrological data of water reservoirs in Colorado. Table 2 lists the datasets which we used: The first three datasets are derived from the Auto MPG Dataset [5], with different columns as the independent variable. The fourth and fifth datasets include hydrological data for estimating a reservoir's volume based on water elevation measurements. The sixth dataset is a hydrological time series.

Table 2: Datasets used in the computational experiments.

Dataset	n	Description
(weight,MPG)	398	Vehicle weight as the independent variable used to explain MPG [5].
(horsepower, MPG)	392	Vehicle horsepower as the independent variable used to explain MPG [5].
(displacement, MPG)	398	Vehicle engine displacement as the independent variable used to explain MPG [5].
(elevation, volume)	50	A subset of a dataset used to estimate the volume of the Morrow Point reservoir as a function of the water elevation.
(elevation, volume)	834	A dataset used to estimate the volume of the Morrow Point reservoir as a function of the water elevation.
(time, elevation)	1216	A dataset of water elevation measurements in the Blue Mesa reservoir over 1216 time periods during the years 2008-2011.

The DP method and Algorithm 2 were implemented in C++. All experiments were run in serial using Intel Q8400 2.66 GHz CPUs with a 2 MB cache size.

6.1 Comparison of the DP Solution Quality vs. MIP

We now compare the solution quality of the proposed DP variants with the solution of the MIP. In this section we limit our consideration to a fixed setting of the parameters of DP-static and DP-adaptive and focus on comparing the solution quality with respect to the optimal solution. The MIP formulations were solved using CPLEX 12.04. A time limit of one hour was applied to all runs. Table 3 compares the running times and solution values of DP with error loss function $e_2(\cdot)$ with optimal solutions computed by (5). Here, the initial rectangle of DP-adaptive is set to be three times the height of the default setting of (8) in order to allow for a larger range of y-intercepts that seemed to be required with $m = 1$. The CPLEX relative optimality gap is set to 1%, and we set $M_1 = M_3 = 1e5$. Table 4 compares the running times and solution values of DP with error loss function $e_1(\cdot)$ using optimal solutions that are computed by (5) with (5a) replaced by $\sum_{i=1}^{n} \xi_i$. Both tables indicate that in most cases the solution quality of the DP methods closely approximates the optimal solution. An exception is the DP-static method with $m = 1$. However, overall DP-adaptive shows an advantage, albeit small, over DP-static in terms of the solution quality. Tables 3 and 4 also display the running times of the three methods. The MIQP can only be solved up to $m = 2$. As shown in Table 4, the MILP can be solved with more breakpoints, but the table also shows that the running times rapidly increase to exceed 50 minutes with 4 breakpoints.

Table 3: Computational results for (elevation,volume), $n = 50$, (#xGrid,#yGrid)=(100,5) for DP-static, and $T = 9$ for DP-adaptive.

m	MIQP		DP-Static			DP-Adaptive		
	Time (sec)	Objective	Time (sec)	Objective	% Error	Time (sec)	Objective	% Error
1	0.01	4,677.89	0.020	18,720.035	300.18	0.744	4,696.153	0.39
2	1.31	289.878	0.022	371.700	28.22	2.053	295.099	1.80
3	Time Limit		0.021	58.096		3.276	60.571	

Table 4: Computational results for (elevation,volume), $n = 50$, (#xGrid,#yGrid)=(100,5) for DP-static, and $T = 9$ for DP-adaptive.

m	MILP		DP-Static			DP-Adaptive		
	Time (sec)	Objective	Time (sec)	Objective	% Error	Time (sec)	Objective	% Error
1	0.46	405.000	0.020	855.400	111.21	0.712	405.527	0.13
2	2.41	100.083	0.021	120.972	20.87	2.125	100.569	0.49
3	193.31	43.179	0.024	44.439	2.92	3.211	45.035	4.30
4	2934.81	18.997	0.025	26.874	41.46	6.263	25.981	36.76

6.2 DP-static vs. DP-adaptive

We now discuss experiments with larger datasets and larger m values to compare our two schemes for constructing tractable grids of candidate breakpoints. In our experiments we first ran DP-static and then DP-adaptive. DP-adaptive was run for sufficiently many iterations so that its objective value (or error) matches that of DP-static. In Tables 5-10, the last column shows the number of iterations of DP-adaptive that were needed to achieve the indicated results.

Tables 6-8 show that typically as m increases, the running time of DP-adaptive may increases in order to match and exceed the quality of the DP-static solution. This is explained by the fact that DP-static allows the breakpoints to be located at or near the y-coordinate of the actual data points. When $m = n$, this heuristic in fact yields the optimal solution with an optimal objective value (i.e., error) of 0. When m is small the quality of this heuristic is poor; but as m is increased while n is held constant, the quality of the heuristic typically improves. Table 6 shows a clear performance advantage of DP-adaptive over DP-static in terms of the running times that are needed to exceed the error performance of the DP-static alternative. Table 7 shows that with a denser grid the quality of DP-static's solutions slightly improves but DP-adaptive easily exceeds the quality of DP-static's solutions within a fraction of the running time.

In Tables 9–10 it becomes apparent that as $m \geq 6$ the running time required of DP-adaptive can significantly increase relative to DP-static; as mentioned this is due to the improved performance of the DP-static heuristic for larger values of m relative to fixed values of n. However, for a wide range of values of m, DP-adaptive remains faster than DP-static while providing a superior solution quality. In Table 9 in cases in which the running time of DP-adaptive exceeds that of DP-static, the improvement in solution quality, on the other hand, also tends to be significant (ranging from 8%-40%).

Figure 5 graphically displays the solutions of the linear spline regression experiments. The solutions of DP-static and DP-adaptive seem to substantially differ for the noisier MPG datasets. For the less noisy water-elevation datasets it appears that the estimated curves are quite close even when the estimated breakpoints do not coincide. Finally, Figure 6 displays an example of the grid that is evaluated by DP-static compared with the grid that is evaluated by DP-adaptive. The figure shows that the DP-Adaptive algorithm evaluates points that seem to be sparsely dispersed around the final solution curve but a little more densely distributed within a close proximity to the curve. The DP-static grid on the other hand is very densely distributed around the datapoints (and final solution curve).

Table 5: (weight,MPG) dataset, $n = 398$, (#xGrid,#yGrid)=(100,5) for DP-static.

m	DP-Static			DP-Adaptive			
	Time (sec)	Objective	$\lvert\Omega\rvert$	Time (sec)	Objective	$\lvert\Omega\rvert$	T
2	0.055	7,031.140	500	0.010	7,019.904	138	5
3	0.055	6,980.827	500	0.012	6,970.296	150	5
4	0.056	6,915.180	500	0.014	6,783.982	165	5
5	0.057	6,830.611	500	0.014	6,783.670	165	5

Table 6: (horsepower,MPG) dataset, $n = 392$, (#xGrid,#yGrid)=(100,5) for DP-static.

m	DP-Static			DP-Adaptive			
	Time (sec)	Objective	$\lvert\Omega\rvert$	Time (sec)	Objective	$\lvert\Omega\rvert$	T
2	0.054	14,129.924	500	0.0002	12,210.006	9	2
3	0.054	8,192.229	500	0.002	7,694.997	60	4
4	0.055	7,467.350	500	0.012	7,176.181	147	5
5	0.055	7,378.477	500	0.012	7,162.479	147	5

Table 7: (horsepower,MPG) dataset, $n = 392$, (#xGrid,#yGrid)=(200,20) for DP-static.

m	DP-Static			DP-Adaptive			
	Time (sec)	Obj Value	$\lvert\Omega\rvert$	Time (sec)	Obj Value	$\lvert\Omega\rvert$	T
2	2.9391	12,250.488	4000	0.0002	12,210.006	9	2
3	2.9780	7,737.867	4000	0.002	7,694.997	60	4
4	3.0164	7,306.523	4000	0.012	7,176.181	147	5
5	3.0555	7,191.656	4000	0.012	7,162.479	147	5

Table 8: (displacement,MPG) dataset, $n = 398$, (#xGrid,#yGrid)=(100,5) for DP-static.

m	DP-Static			DP-Adaptive			
	Time (sec)	Objective	$\|\Omega\|$	Time (sec)	Objective	$\|\Omega\|$	T
2	0.050	8,811.040	500	0.002	7,808.003	54	4
3	0.051	7,328.727	500	0.050	7,092.556	328	6
4	0.055	6,817.053	500	0.073	6,811.856	406	6
5	0.058	6,741.545	500	0.346	6,593.420	855	7

Table 9: (elevation,storage) dataset, $n = 834$, (#xGrid,#yGrid)=(200,5) for DP-static.

m	DP-Static			DP-Adaptive			
	Time (sec)	Objective	$\|\Omega\|$	Time (sec)	Objective	$\|\Omega\|$	T
2	0.352	71,814,118,000	1000	0.002	25,628,928,000	57	4
3	0.356	13,302,963,000	1000	0.016	5,620,973,400	153	5
4	0.358	3,630,796,100	1000	0.016	2,054,457,400	152	5
5	0.362	1,230,120,200	1000	0.082	664,427,240	393	6
6	0.363	546,859,270	1000	0.086	406,544,860	397	6
7	0.365	221,637,540	1000	0.101	221,608,790	431	6
8	0.369	100,619,500	1000	5.259	60,009,459	2960	8
9	0.373	47,851,391	1000	5.569	44,004,352	3046	8
10	0.376	24,115,875	1000	35.878	21,359,717	7370	9

Table 10: (time,elevation) dataset, $n = 1216$, (#xGrid,#yGrid)=(200,5) for DP-static.

m	DP-Static			DP-Adaptive			
	Time (sec)	Objective	$\|\Omega\|$	Time (sec)	Objective	$\|\Omega\|$	T
2	0.473	541.8554	1000	0.003	285.8593	54	4
3	0.477	127.3319	1000	0.015	66.5020	137	5
4	0.485	28.4243	1000	0.073	17.4254	333	6
5	0.483	14.6801	1000	0.083	11.6842	352	6
6	0.489	9.7675	1000	0.083	8.368	352	6
7	0.491	5.0368	1000	3.387	4.587	2276	8
8	0.494	4.2606	1000	3.569	4.034	2312	8
9	0.497	3.7497	1000	4.250	3.551	2504	8
10	0.500	3.0380	1000	4.323	2.998	2515	8

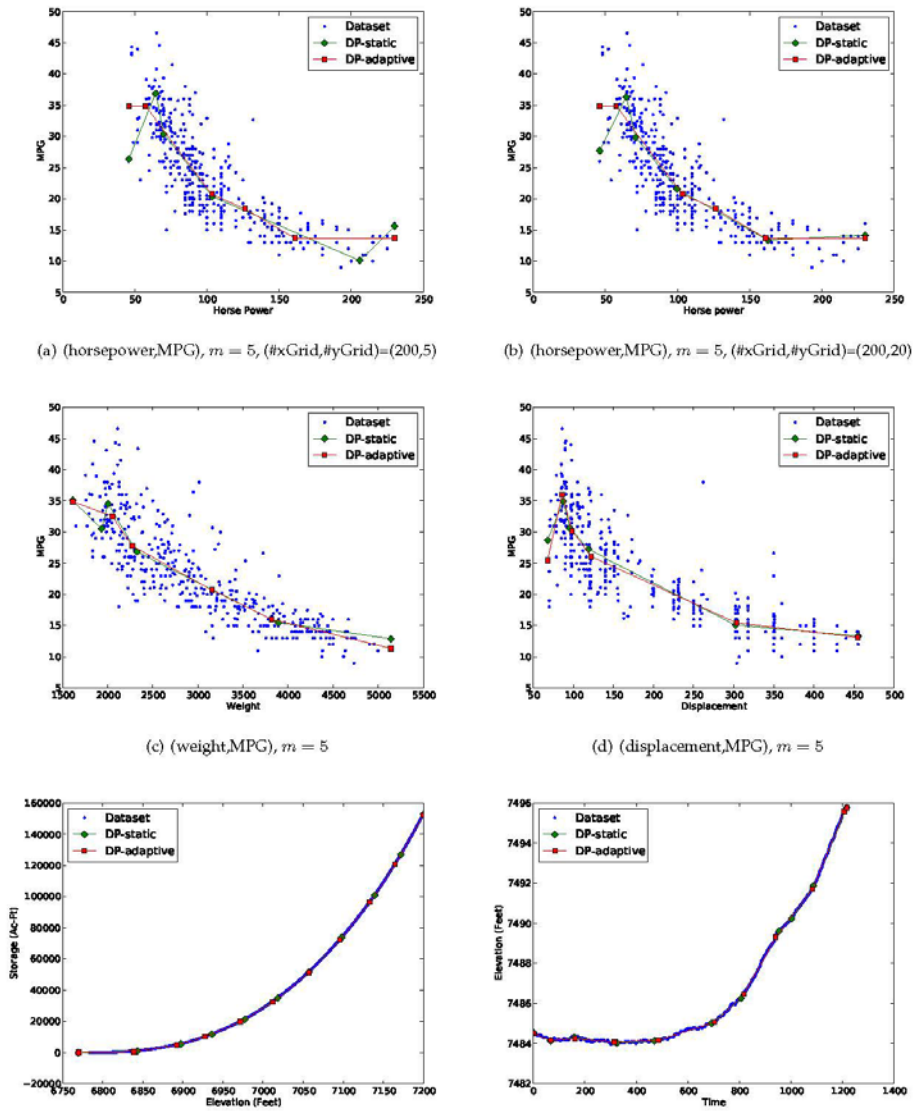

(a) (horsepower,MPG), $m = 5$, (#xGrid,#yGrid)=(200,5)

(b) (horsepower,MPG), $m = 5$, (#xGrid,#yGrid)=(200,20)

(c) (weight,MPG), $m = 5$

(d) (displacement,MPG), $m = 5$

(e) (elevation,storage), $m = 10$

(f) (time,elevation), $m = 10$

Figure 5: Graphical representation of numerical results of DP-static and DP-adaptive on several datasets. Figure 5(b) show the results for the same dataset and setting as Figure 5(a) but with DP-static using a denser static. Also, note that both methods result in similar solutions in Figure 5(e) where the approximating curve is convex.

Finally, Figure 6 shows the static grid next to the grid resulting from adaptive refinement.

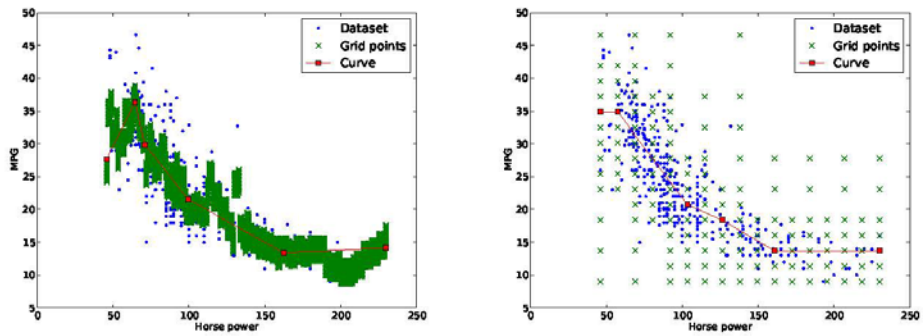

(a) DP-Static (b) DP-Adaptive

Figure 6: Comparison of the dense static grid and the adaptively refined grid for the (horsepower,MPG) dataset.

6.3 DP-Static vs. DP-Adaptive with Two Curves

We now compare the performance of the static grid and adaptive refinement for the extended problem of determining two curves that share breakpoints with identical horizontal coordinates. In Table 11, the column $(|\Omega_1|, |\Omega_2|)$ indicates the total number of candidate breakpoints generated for the first and second curves, respectively. The results of the table indicate that DP-adaptive is competitive in terms of the running time required to obtain a given solution quality. The running time of DP-adaptive significantly increases, however, for $m = 5$ with $T = 6$. The improved solution quality however is also significant in this case and Figure 7 displays the different estimated curves associated with the DP-static and DP-adaptive solutions.

Table 11: Comparison of DP-static and the DP-adaptive Algorithm 3 with $\nu = 2$ (two curves) using the (weight,MPG) dataset. Here D_1 consists of vehicles of origin type 2 ($|D_1| = 70$), and D_2 consists of vehicles of origin type 3 ($|D_2| = 79$). Also, (#xGrid,#yGrid)=(30,5) for DP-static.

m	DP-Static			DP-Adaptive			
	Time (sec)	Objective	$\lvert\Omega_1\rvert, \lvert\Omega_2\rvert$	Time (sec)	Objective	$\lvert\Omega_1\rvert, \lvert\Omega_2\rvert$	T
2	0.039	4431.550	(150,150)	0.003	4363.912	(22,22)	3
3	0.040	4230.532	(150,150)	0.064	4071.355	(66,57)	4
4	0.041	4108.698	(150,150)	0.064	4043.569	(66,57)	4
5	0.042	4033.527	(150,150)	1.168	3864.417	(175,158)	5

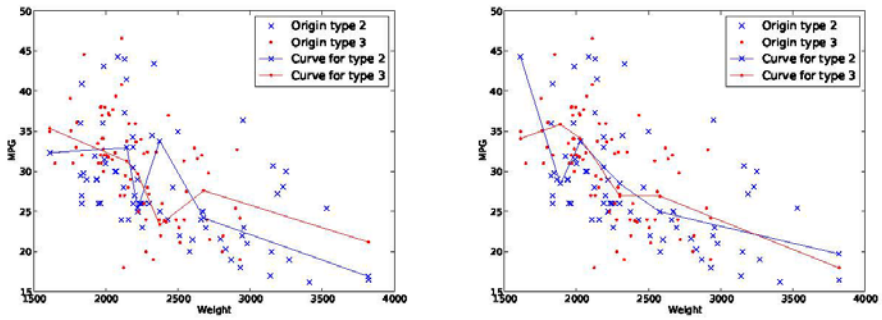

(a) DP-static (b) DP-adaptive

Figure 7: Graphical representation of numerical results of DP-static and DP-adaptive for the UCI MPG dataset [5]. Here D_1 contains the subset of vehicles with origin type 2, D_2 contains the subset of vehicles with origin type 3, and $m = 5$.

7 Conclusion

We have shown the efficacy of using nonuniform grids for dynamic programming in solving linear spline regression. In particular, we devise an adaptive refinement scheme for solving linear spline regression with an arbitrary location of breakpoints. We compared our methods with the optimal solutions computed using a novel MIP formulation of the problem. The advantage of the adaptive refinement scheme is particularly apparent for (difficult) instances whose points are well approximated using $m \geq 2$ line segments where m remains small relative to n. When m is large relative to n, an alternative heuristic that determines a small set of candidate breakpoints near the data points is competitive in terms of solution quality and speed.

Finally, we have extended the DP algorithm and our grid construction approach to fit multiple piecewise linear curves that share the breakpoint horizontal coordinates (or change points). We also find the proposed adaptive refinement scheme useful in this setting.

Acknowledgements

We thank Jed Brown for comments and Ashutosh Mahajan for pointing us to [2].

A Analysis of the Running Time of the Dynamic Program

Proposition 1. *Suppose that for $w \in W = \{1, \ldots, v\}$, and $i = 1, \ldots, p$ we have $S^w(t_i) \leq q$. Then, if we assume $E(\cdot)$ is given, the running time of the DP (9) is $O(p^2 q^{2v} m)$.*

Proof. The DP requires the evaluation of $F_{\widehat{m}}(i,j)$ for $\widehat{m} = 1, \ldots, m$, $i = 1, \ldots, p$, and $j \in S(t_i)$.

First, $F_1(i,j)$ is evaluated for $i = 2, \ldots, p-m+1$ and $j \in \times_{w \in W} S^w(t_i)$ and each such evaluation is $O(1)$ assuming that E is given. For each $i \in \{1, \ldots, p-(m-1)\}$, $|S^w(t_i)| \leq q$ for each $w \in W$, so that q^{2v} operations are used to enumerate elements of $S^w(t_1) \times S^w(t_i)$ in order to determine $F_1(i,j) = \min_{t_w \in S^w(t_1): w \in W} E(t_1, \times_{w \in W} s^w_{1t_w}, t_i, \times_{w \in W} s^w_{ij_w})$, for $j \in \times_{w \in W} S^w(t_i)$ and $w \in W$. Hence, in total at most $(p-m)q^{2v}$ evaluations are required to determine $F_1(i,j)$ for $i = 1, \ldots, p-m+1$ and $j \in \times_{w \in W} S^w(i)$.

Next, for $\widehat{m} \in \{2, \ldots, m\}$, $F_{\widehat{m}}(i,j)$ is evaluated for $i \in I \stackrel{\text{Def}}{=} \{\widehat{m}+1, \ldots, p-m+\widehat{m}\}$ and $j \in \times_{w \in W} S^w(i)$. Note that $|I| = p-m$. For a given $i \in I$ and $j \in S(t_i)$, $F_{\widehat{m}}(i,j)$ evaluates $F_{\widehat{m}-1}(k, \cdot)$ for $k \in \{\widehat{m}+1, \ldots, i-1\}$. For each pair $(k,i) \in I \times I$ with $k \neq i$, and $w \in W$, at most q^2 evaluations are required. Then, the number of segment combinations over $w \in W$ is $q^{2|W|} = q^{2v}$. There are at most $\binom{(p-m)}{2}$ such pairs (where $k \neq i$), so that at most $\frac{(p-m)(p-m-1)}{2} q^{2v}$ evaluations are required to compute $F_{\widehat{m}}(i,j)$ for $i = \widehat{m}, \ldots, p-m+\widehat{m}$ and $j \in \times_{w \in W} S^w(t_i)$.

Hence, the number of operations is bounded by

$$(p-m)q^{2v} + (m-1)\frac{(p-m)(p-m-1)}{2}q^{2v} = (p-m)q^2\left(\frac{(m+1)+(m-1)(p-m)}{2}\right)$$
$$\in O((p-m)^2q^{2v}m) \subseteq O(p^2q^{2v}m). \qquad \square$$

If $v = 1$ then the total running time may be dominated by the computation of the segment errors $E(\cdot)$. If N is the total number of breakpoints then note that the running time of computing the segment errors is bounded by $O(N^2n)$. As (9) with $v = 1$ reduces to (6), the following corollary establishes the running time of evaluating (6) following the precomputation of $E(\cdot)$.

Corollary 1. *Suppose that for $i = 1, \ldots, p$ we have $S^w(t_i) \leq q$ and let $N = pq$. Then, if we assume $E(\cdot)$ is given, the running time of the DP (6) is $O(N^2m)$.*

References

[1] B. Aronov, T. Asanov, N. Katoh, and K. Mehlhorn. Polyline fitting of planar points under min-sum criteria. *International Journal of Computational Geometry and Applications*, 16(2&3):97–116, 2006.

[2] J. Bai and P. Perron. Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 18:1–22, 2003.

[3] R. Bellman and R. Roth. Curve fitting by segmented straight lines. *Journal of the American Statistical Association*, 64:1079–1084, 1969.

[4] J.E. Ertel and E.B. Fowlkes. Some algorithms for linear spline and piecewise multiple linear regression. *Journal of the American Statistical Association*, 71(355):640–648, 1976.

[5] A. Frank and A. Asuncion. UCI machine learning repository, 2010.

[6] S.B. Guthery. Partition regression. *Journal of the American Statistical Association*, 348:945–947, 1974.

[7] A. Kehagias, E. Nidelkou, and V. Petridis. A dynamic programming segmentation procedure for hydrological and environmental time series. *Stochastic Environmental Research and Risk Assessment*, 20(1):77–94, 2005.

[8] A. Toriello and J.P. Vielma. Fitting piecewise linear continuous functions. *European Journal of Operational Research*, 219(1):89–95, 2012.

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.

APPENDIX G:

PREDICTED WATER-TO-POWER CONVERSION EFFICIENCY CURVES FOR HYDROPOWER GENERATING UNITS IN THE ASPINALL CASCADE

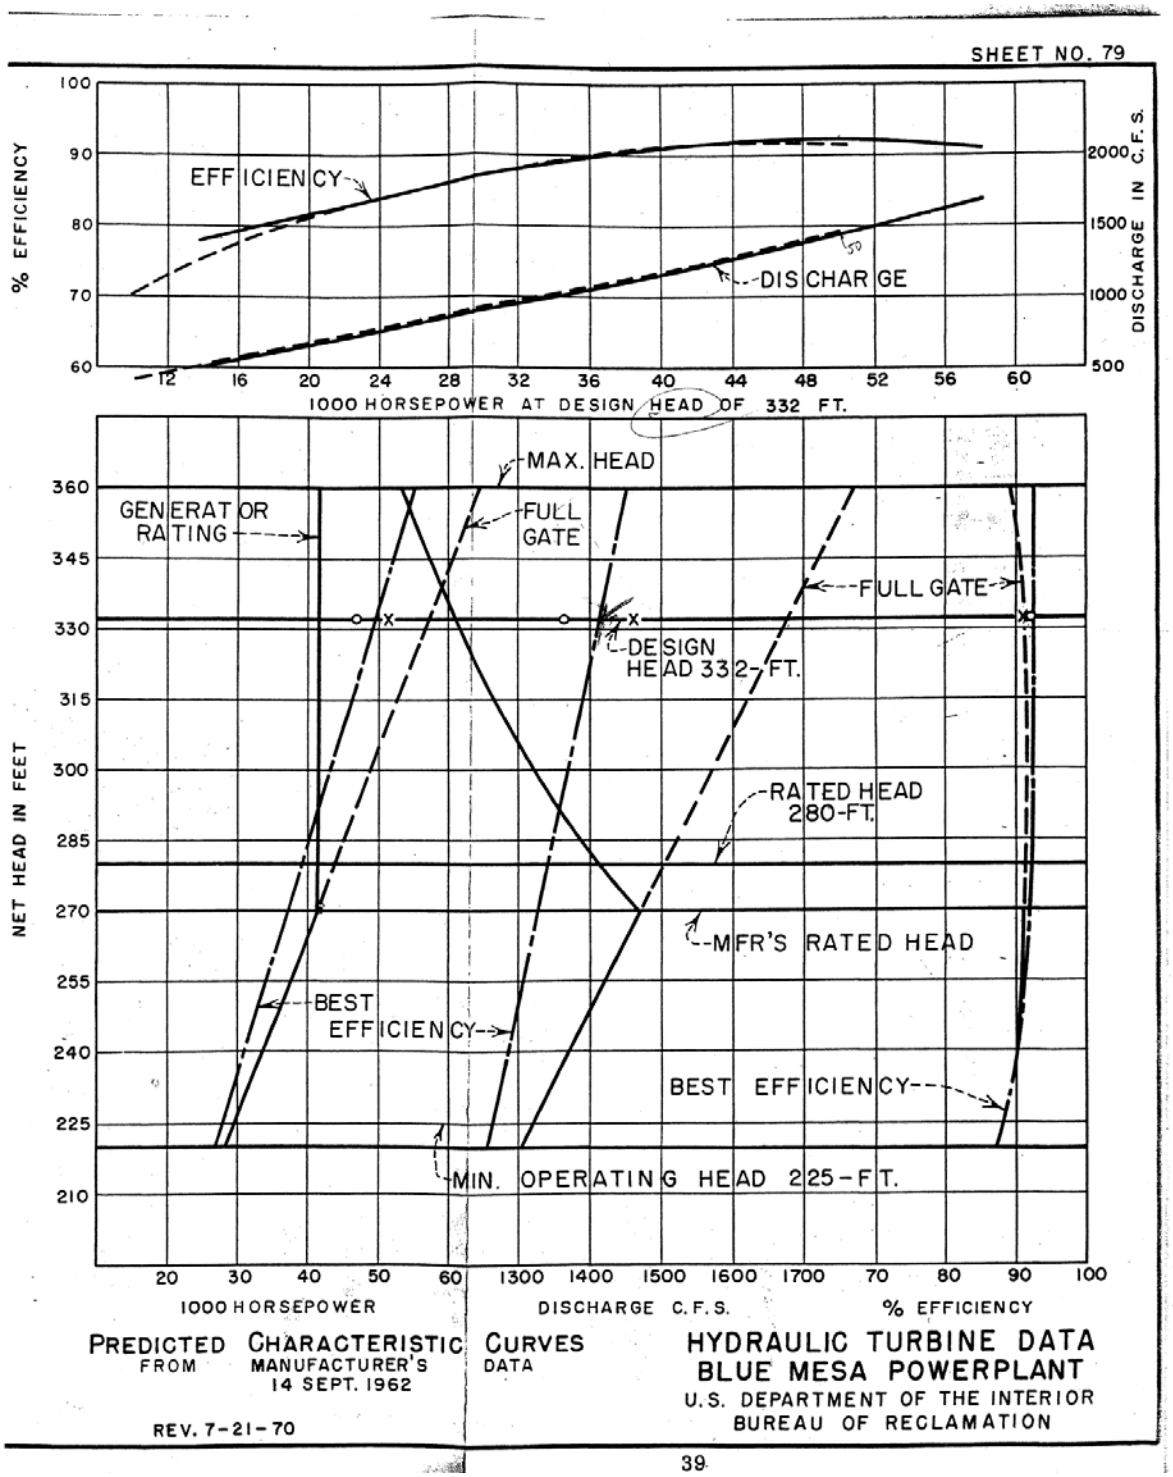

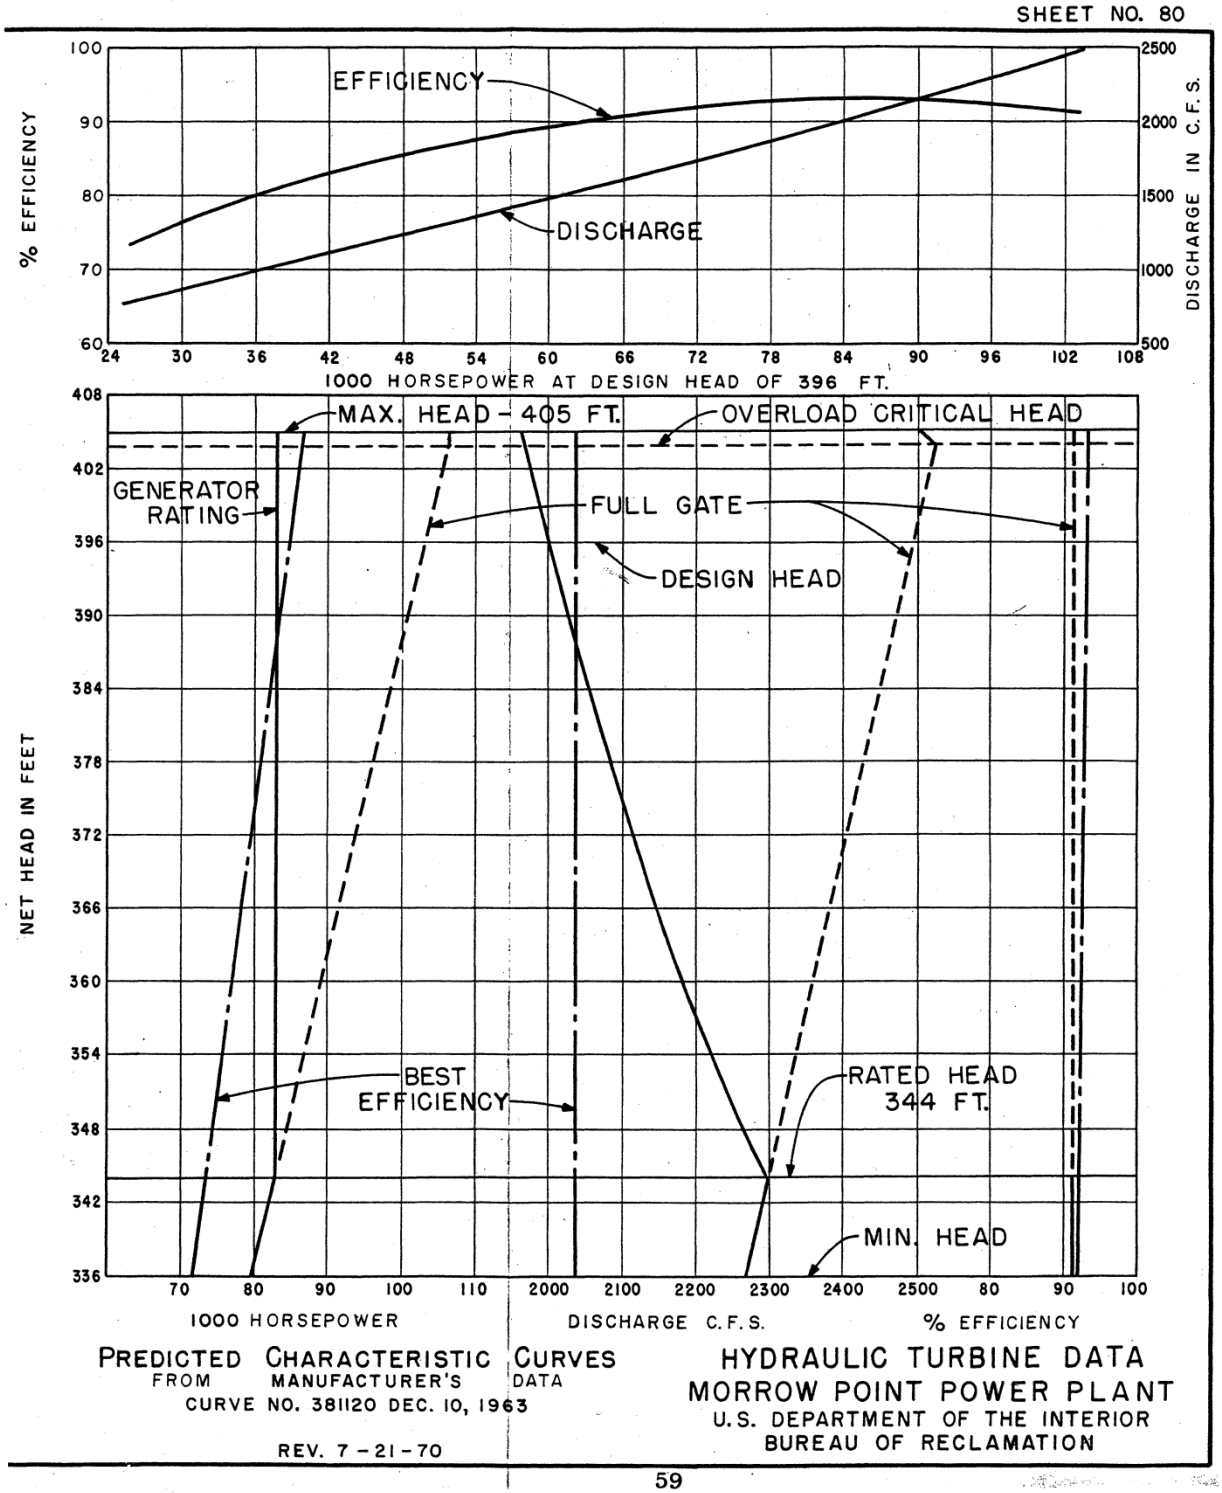
SHEET NO. 80
EFFICIENCY
DISCHARGE
% EFFICIENCY
DISCHARGE IN C.F.S.
1000 HORSEPOWER AT DESIGN HEAD OF 396 FT.
MAX. HEAD - 405 FT.
OVERLOAD CRITICAL HEAD
GENERATOR RATING
FULL GATE
DESIGN HEAD
BEST EFFICIENCY
RATED HEAD 344 FT.
MIN. HEAD
NET HEAD IN FEET
1000 HORSEPOWER
DISCHARGE C.F.S.
% EFFICIENCY
PREDICTED CHARACTERISTIC CURVES
FROM MANUFACTURER'S DATA
CURVE NO. 381120 DEC. 10, 1963
REV. 7-21-70
HYDRAULIC TURBINE DATA
MORROW POINT POWER PLANT
U.S. DEPARTMENT OF THE INTERIOR
BUREAU OF RECLAMATION

59

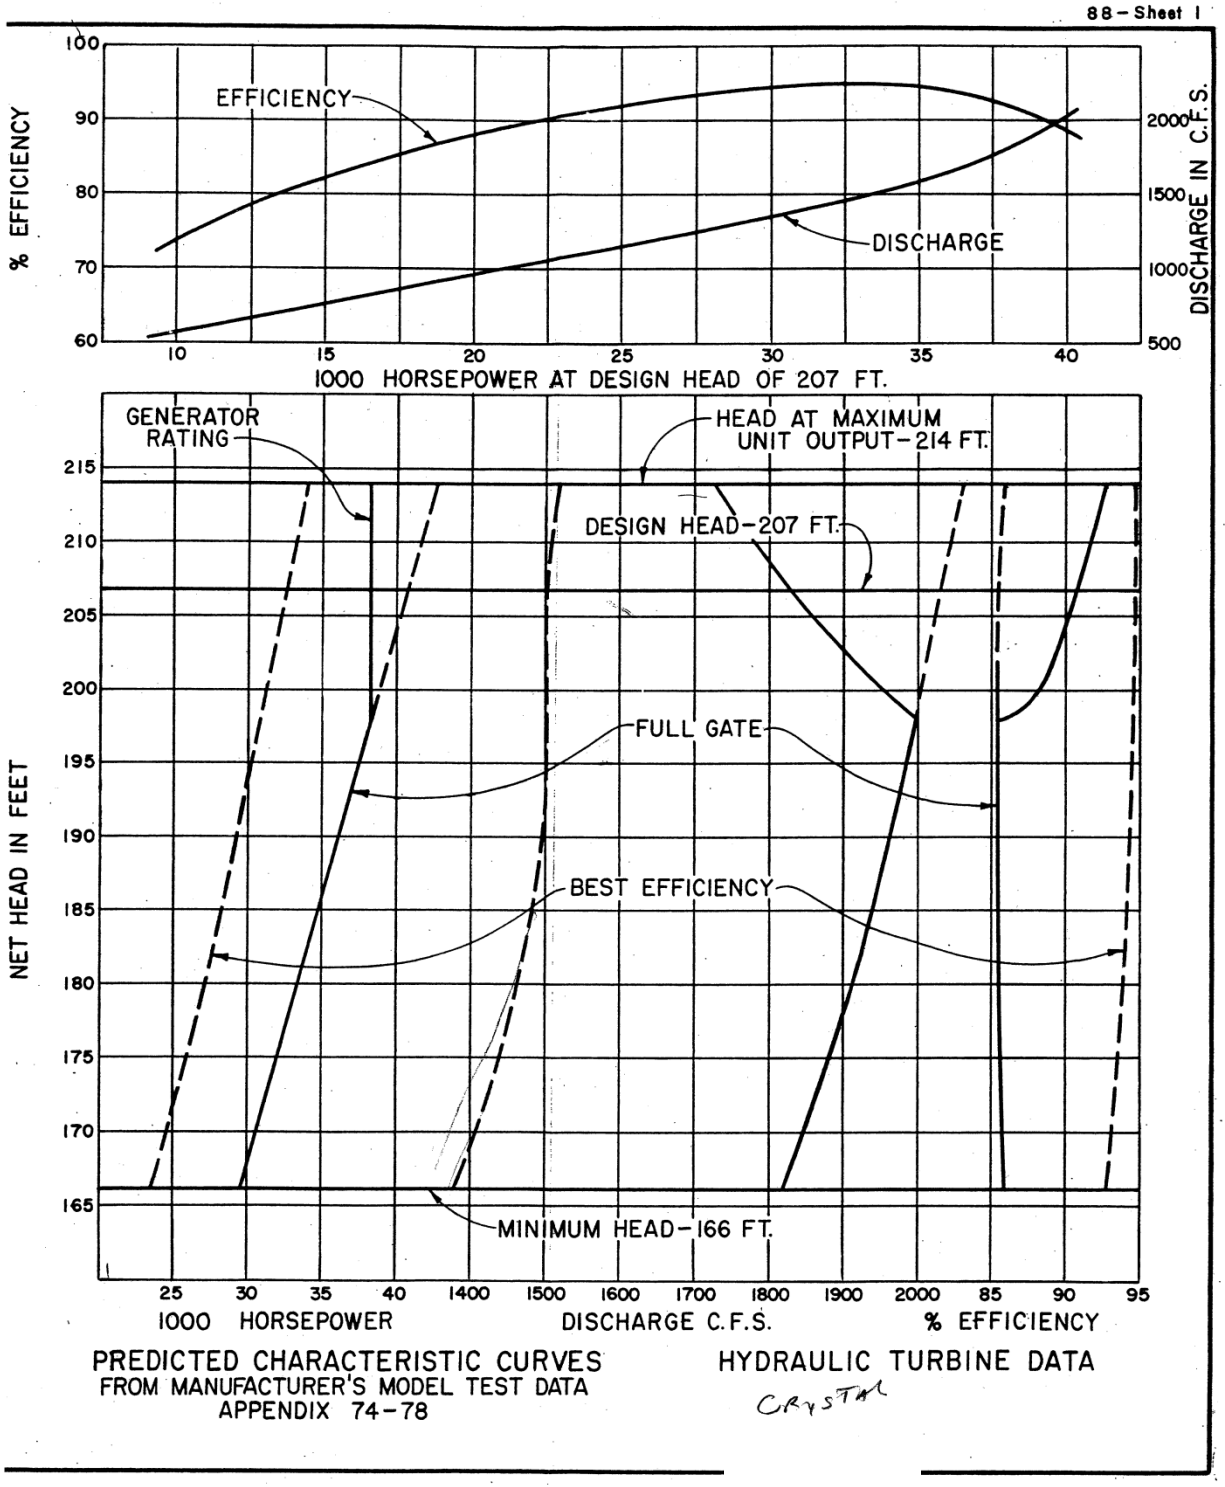
88 – Sheet 1
% EFFICIENCY
EFFICIENCY
DISCHARGE
DISCHARGE IN C.F.S.
1000 HORSEPOWER AT DESIGN HEAD OF 207 FT.
GENERATOR RATING
HEAD AT MAXIMUM UNIT OUTPUT – 214 FT.
DESIGN HEAD – 207 FT.
FULL GATE
BEST EFFICIENCY
MINIMUM HEAD – 166 FT.
NET HEAD IN FEET
1000 HORSEPOWER
DISCHARGE C.F.S.
% EFFICIENCY
PREDICTED CHARACTERISTIC CURVES
FROM MANUFACTURER'S MODEL TEST DATA
APPENDIX 74-78
HYDRAULIC TURBINE DATA
CRYSTAL

41

APPENDIX H:

HEAD LOSS APPROXIMATIONS AT BLUE MESA AND MORROW POINT

As shown above, power output is computed as a piecewise linear function of water release rate for each unit of Blue Mesa and Morrow Point power plants. However, if both units are operating at the same time, using individual piecewise linear curves overestimate the power output. This is due to ignoring the head loss when both units are operating together. To estimate this difference, head-generation relationship is investigated. Figures H-1 and H-3 represent the head-generation curves for different water release rate levels for Blue Mesa Unit 1 and Morrow Point Unit 1, respectively. Slopes and intercepts of the linear power productions functions for Blue Mesa unit 1 and Morrow Point unit 1 are given in Tables H-1 and H-2. By using these relationships, an upper bound and a lower bound on the water release rate-generation curve are created as follows. Find the maximum and minimum head in a month. For each head, find the generation level given the water release rate. If the water release rate is not defined by a line, interpolate between two lines and find the generation level.

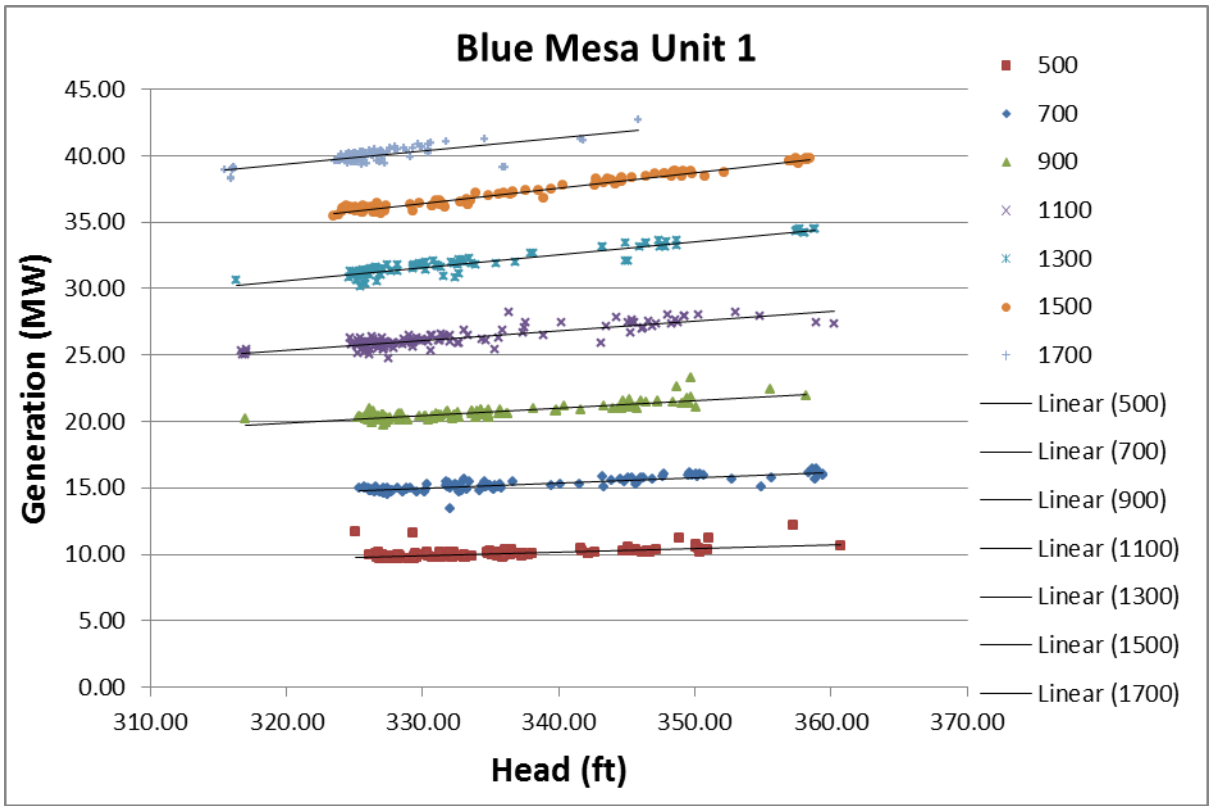

FIGURE H-1 Blue Mesa Unit 1 Head and Generation Relationship for Different Water Release Rates

TABLE H-1 Intercepts and Slopes of Head-Generation Linear Approximations, Blue Mesa Unit 1

Water Release Rate (cfs)	Intercept	1st Order
500	1.23E+00	2.64E-02
700	1.57E+00	4.06E-02
900	1.64E+00	5.70E-02
1100	1.60E+00	7.42E-02
1300	–9.17E-01	9.84E-02
1500	–1.77E+00	1.16E-01
1700	7.87E+00	9.85E-02

Figure H-2 represents the water release rate–generation relationship for maximum and minimum head for Blue Mesa Unit 1. The resulting water release rate–generation curves for maximum and minimum head in April 2009 for Blue Mesa unit 1 and Morrow Point unit 1 are shown in Figures H-3 and H-4.

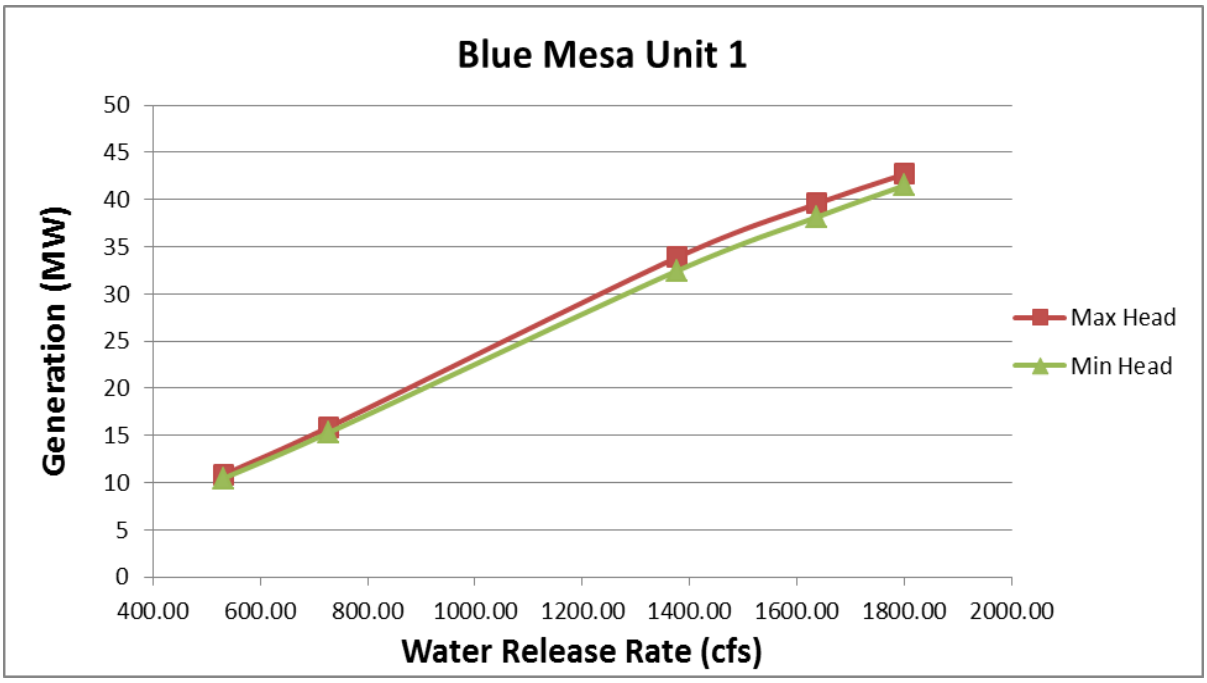

FIGURE H-2 Blue Mesa Unit 1 Head Loss Approximation, April 2009, Min Head = 320.77 ft, Max Head = 334.27 ft

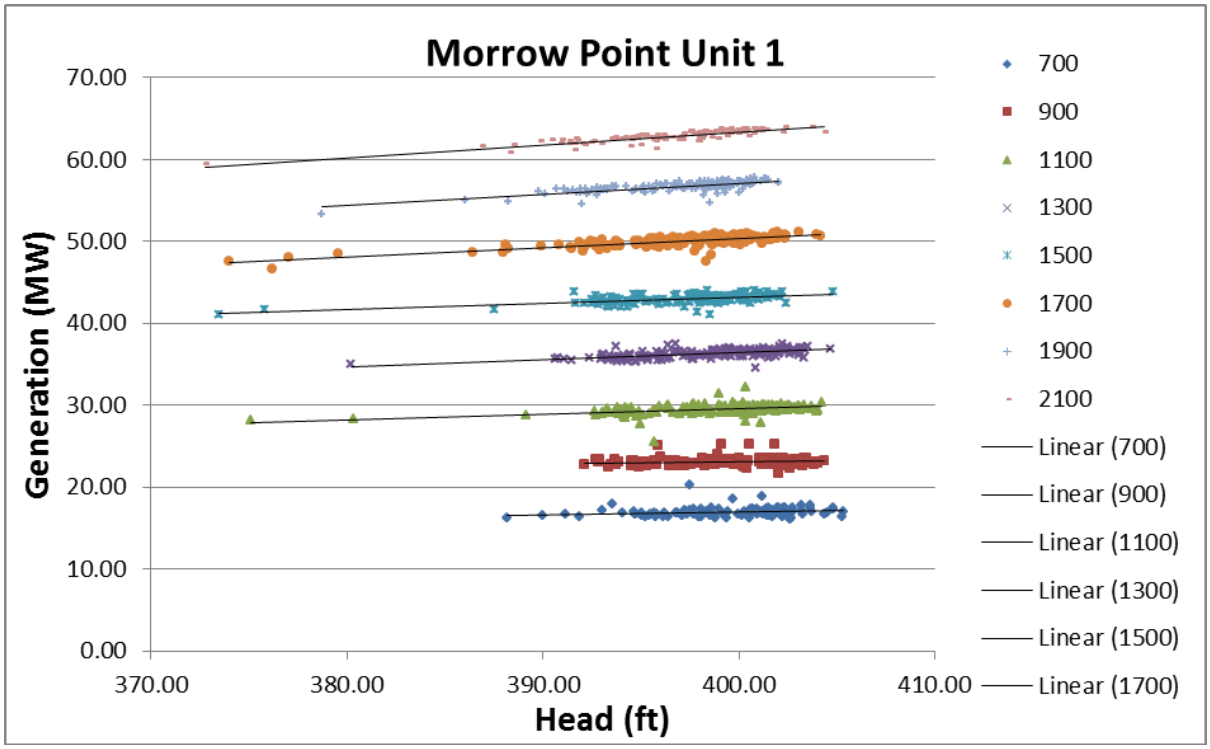

FIGURE H-3 Morrow Point Unit 1 Head and Generation Relationship for Different Water Release Rates

TABLE H-2 Intercepts and Slopes of Head-Generation Linear Approximations, Morrow Point Unit 1

Water Release Rate (cfs)	Intercept	1st Order
700	4.36E+00	3.15E-02
900	1.61E+01	1.75E-02
1100	2.39E+00	6.80E-02
1300	1.49E+00	8.73E-02
1500	1.38E+01	7.34E-02
1700	5.53E+00	1.12E-01
1900	2.33E+00	1.37E-01
2100	1.17E+00	1.55E-01

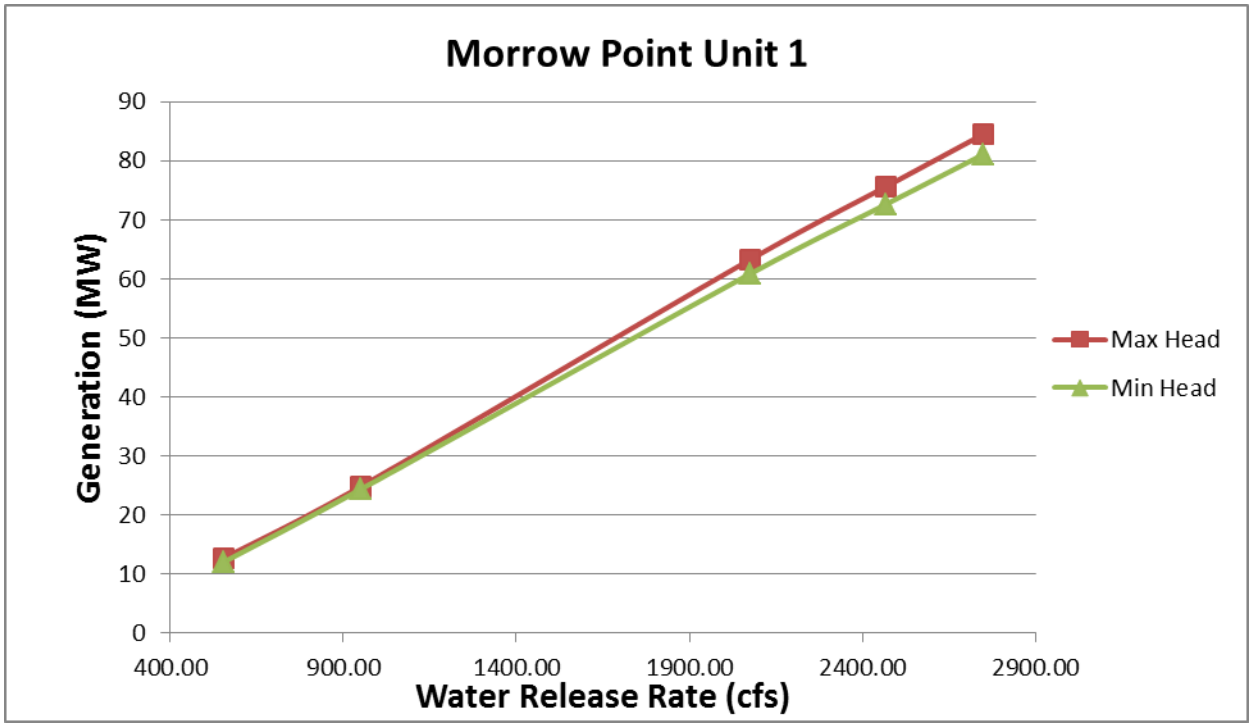

FIGURE H-4 Morrow Point Unit 1 Head Loss Approximation, April 2009, Min Head = 388.99 ft, Max Head = 404.83 ft

APPENDIX I:

HOURLY CHEERS PRICES

Hourly electricity prices were estimated using a multistep process — the first of which is to collect and process historical market prices. As shown in Figure I-1, this process uses market prices posted by two independent entities, namely, the Intercontinental Exchange (ICE) and the CRSP Management Center in Salt Lake City, Utah. Data from ICE contains daily minimum and maximum prices, along with an average daily price for several WECC market hubs. The Western data originates from the Energy Management and Marketing Office (EMMO) and is reported as weekly on-peak and off-peak price ranges for areas; namely Four Corners, Central Rockies, and the Southwest. Monthly on-peak and off-peak price data were also collected for actual purchases made by the EMMO in Montrose, Colorado. These data were only used for comparative purposes and did not factor into the hourly price estimates used for economic evaluations.

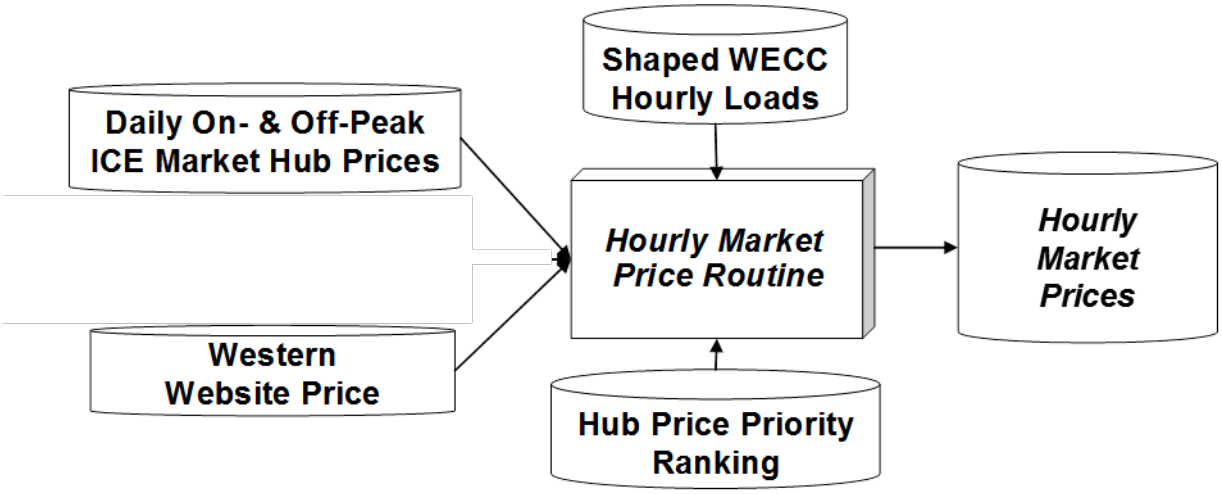

FIGURE I-1 Process for Estimating Hourly WECC Market Hub Prices

ICE serves are the primary price data source, while information found on the Western Web site is used to fill in data gaps. These daily prices are then used as the basis for estimating day-ahead STF market prices. Hourly real-time prices input into CHEERS also rely on this daily information. Hourly prices are estimated by a price-shaping algorithm such that the resultant price vector has the same basic shape as WECC loads.

Hourly loads for WECC subregions to support this hourly price-shaping routine are estimated from historical data collected by FERC (in Form-714) and WECC statistics. As shown in Figure I-2, the process uses hourly loads collected for all control areas in WECC that are located the United States. Consistency checks are performed on the data and adjustments made when errors are found and data are missing. Control area loads are then grouped and aggregated into the following four WECC subregions: (1) Northwest Power Pool (NWPP), (2) Rocky Mountain Power Area (RMPA), (3) Arizona, New Mexico, and Southern Nevada Power Area

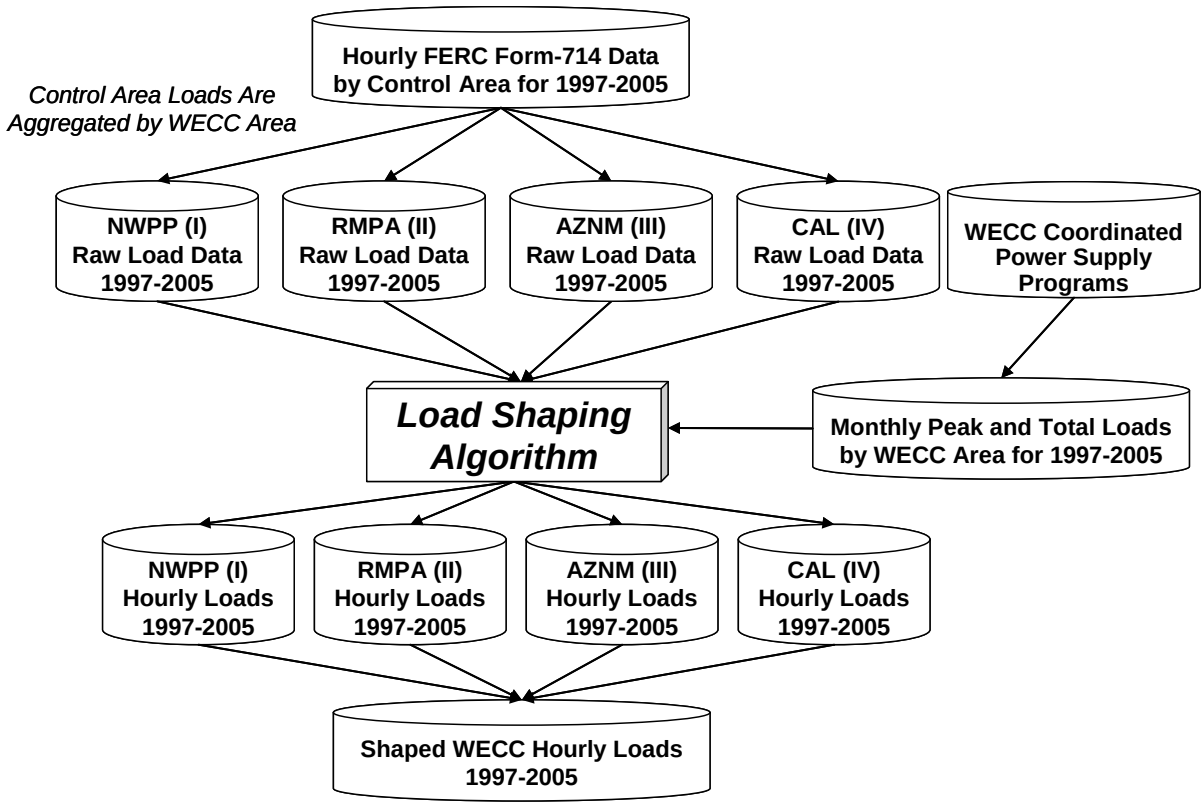

FIGURE I-2 Process for Estimating Total Hourly Loads in the U.S. Portion of WECC

(AZNM), and (4) California and Mexico Power Area (CAMX). Next, a Load Shaping Algorithm adjusts aggregated hourly load profiles to exactly match monthly peak and total load values that are reported for each WECC subregion.

The Loads Shaping Algorithm uses a QP technique that minimizes differences between a normalized Load Duration Curve (LDC) constructed from historical data and a reshaped LDC generated by the model. Figure I-3 shows the original LDC, constructed from control area historical loads for August 2005 in the AZNM sub-region and the reshaped LDC. The reshaped curve is consistent with a monthly load factor computed from the peak and total load values reported by WECC. Upper and lower load constraints are specified by the user to bind the model's solution. For each point in the LDC, a scaling factor is then computed as the ratio of the reshaped load to the original load. Finally, the algorithm constructs a scaled chronological hourly profile based on the load scaling factors and an associated original hourly load. The end product, as shown in Figure I-4, is a chronological load profile that exactly matches WECC monthly statistics.

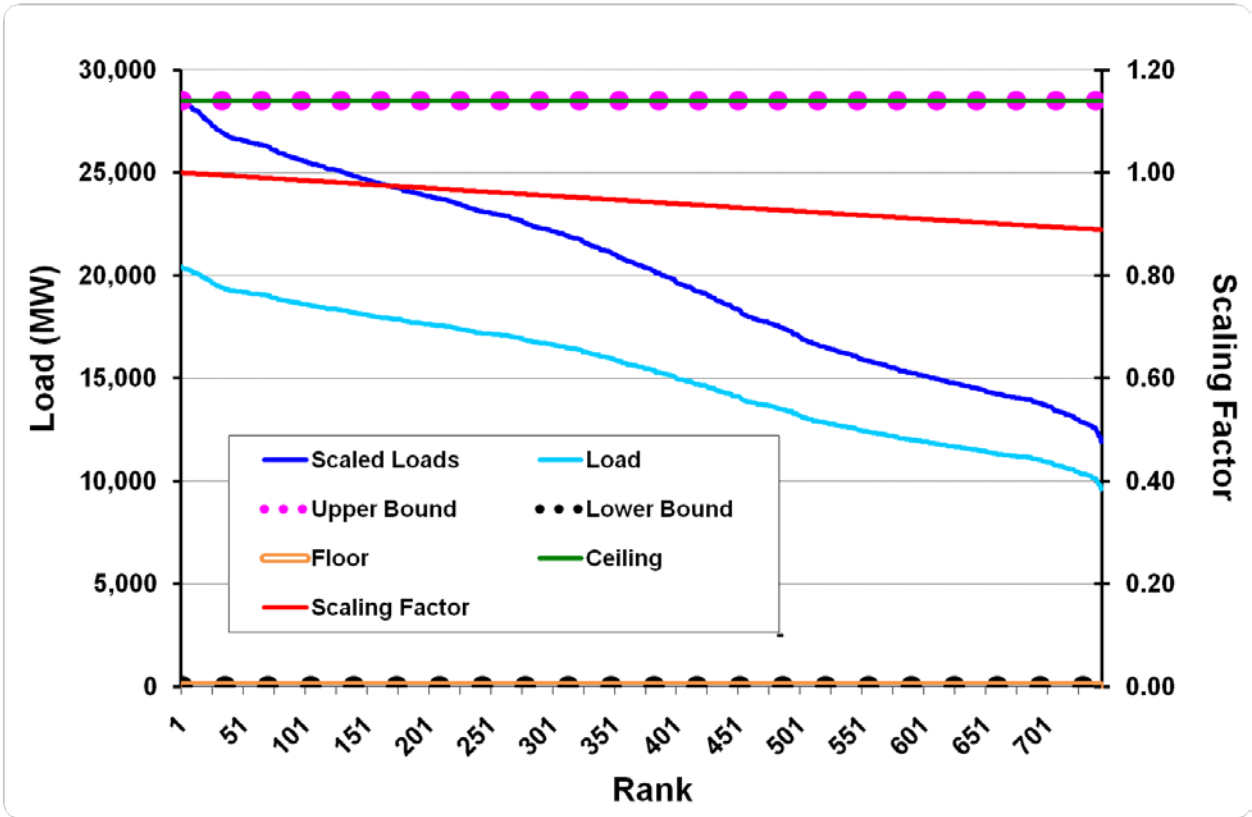

FIGURE I-3 Illustration of Load Duration Curve Shaping to Match a Target Load Factor

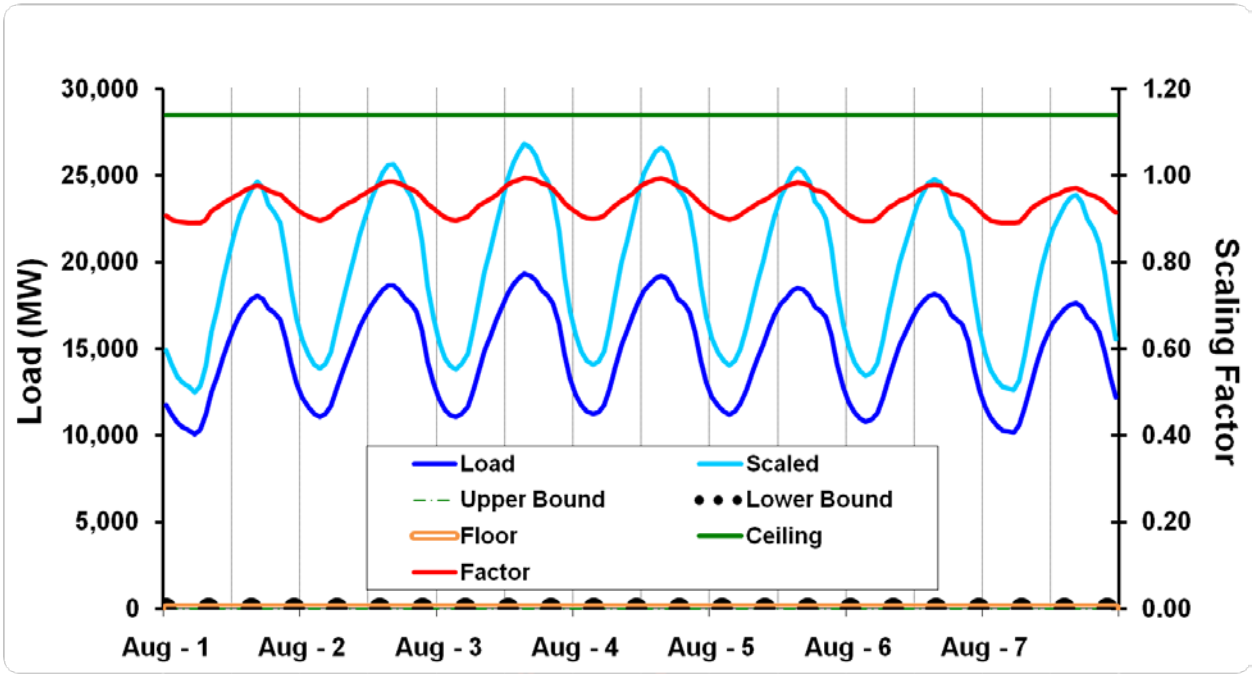

FIGURE I-4 Original and Shaped Chronological Load Curve

APPENDIX J:

CUSTOMIZED INFLOW FORECAST ERROR RESERVOIR BUFFERS

One statistical error distribution does not fit all hydrological conditions. There are seasonal differences. Therefore, to improve upon the construction and use of inflow forecast error probability distributions, data are stratified by month to capture the nuances of each season. For the Aspinall Cascade, forecast errors in Spring tend to be the highest. During snowmelt months data are further stratify into continuous six hour periods. Since errors during Spring are distinctly different over the course or a day. The daytime has significant snowmelt while the night the snowpack melt is dramatically reduced as temperatures drop below the freezing level.

Forecasts in the mid to late afternoon tend to underestimate inflows while nighttime forecasts tend to overestimate inflows.

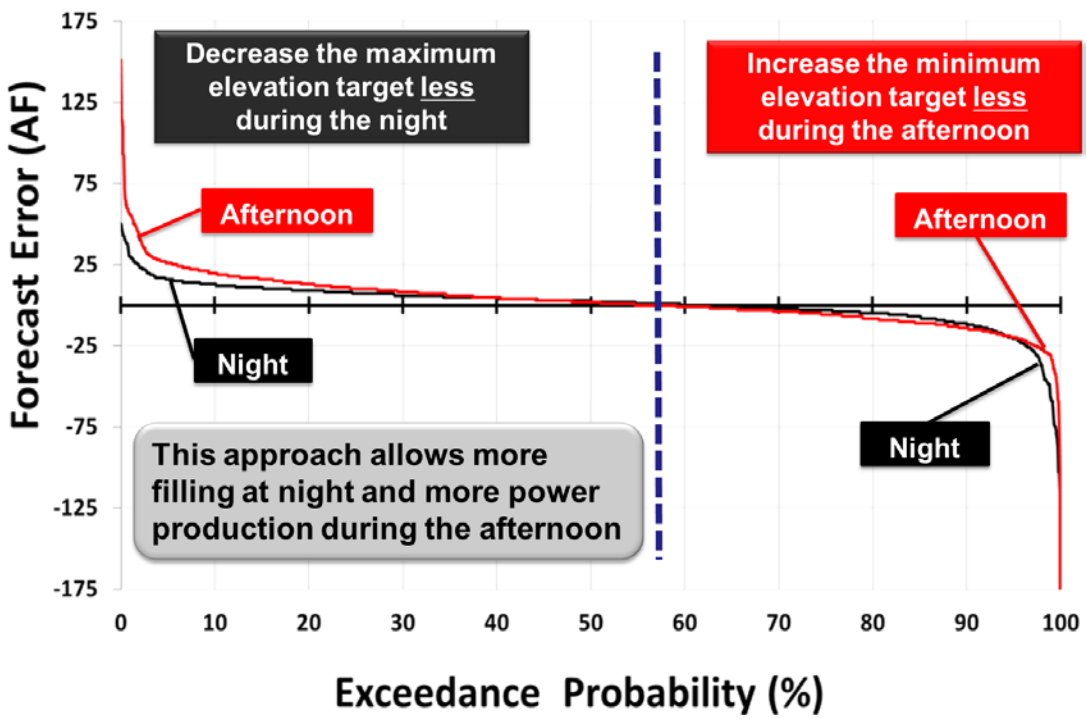

FIGURE J-1 Forecast Error versus Exceedance Probability

APPENDIX K:

HISTORICAL CRYSTAL RESERVOIR OPERATING VIOLATIONS

The Crystal reservoir is subject to a number of "rolling" elevation change constraints. During the first three months of the retrospective analysis, the reservoir level may not change by 4, 5, 6, or 7 feet in any 24, 48, 72, or 96 hour period, respectively. During the last 3 months of the retrospective analysis, the constraints change to 10 feet over 24 hours and 15 feet over 72 hours. A constraint is referred to as "rolling" when the temporal period over which it is enforced can be any *n*-hour period, say 3 PM to 3 PM, rather than just over calendar days (midnight to midnight). Each Conventional Hydropower Energy and Environmental Resource System (CHEERS) run in the retrospective analysis includes 4 full days of unchangeable historic values that precede the first hour to be optimized by the model. This history is input into CHEERS so that optimized values for future reservoir elevations do not violate these rules.

During the actual historic operation of the Crystal Reservoir, constraints were violated several times. These violations occurred due to inflow/side flow forecast errors. As shown in Figure K-1, these errors are typically the largest when inflows and side flows rapidly change over time. Because the stand-alone CHEERS retrospective analysis is following the same historic daily release volumes and uses historical operations as a modeling starting point, experiencing some of the same violations in the optimized runs is unavoidable. However, when the magnitude of the historic violation is small, CHEERS will often find an optimization regime that completely eliminates the violation. When the model cannot find a compliant solution, it minimizes the violation level. Crystal Reservoir elevations also encounter some violations when running the integrated Water Use Optimization Toolset (WUOT). CHEERS runs in spite of the fact that the daily release volumes are not locked in to the historic amounts. In such cases, the violations are due to inflow forecast uncertainty; when a large amount of water enters unexpectedly into the system, the plants might have to be operated in such a way that a violation is unavoidable.

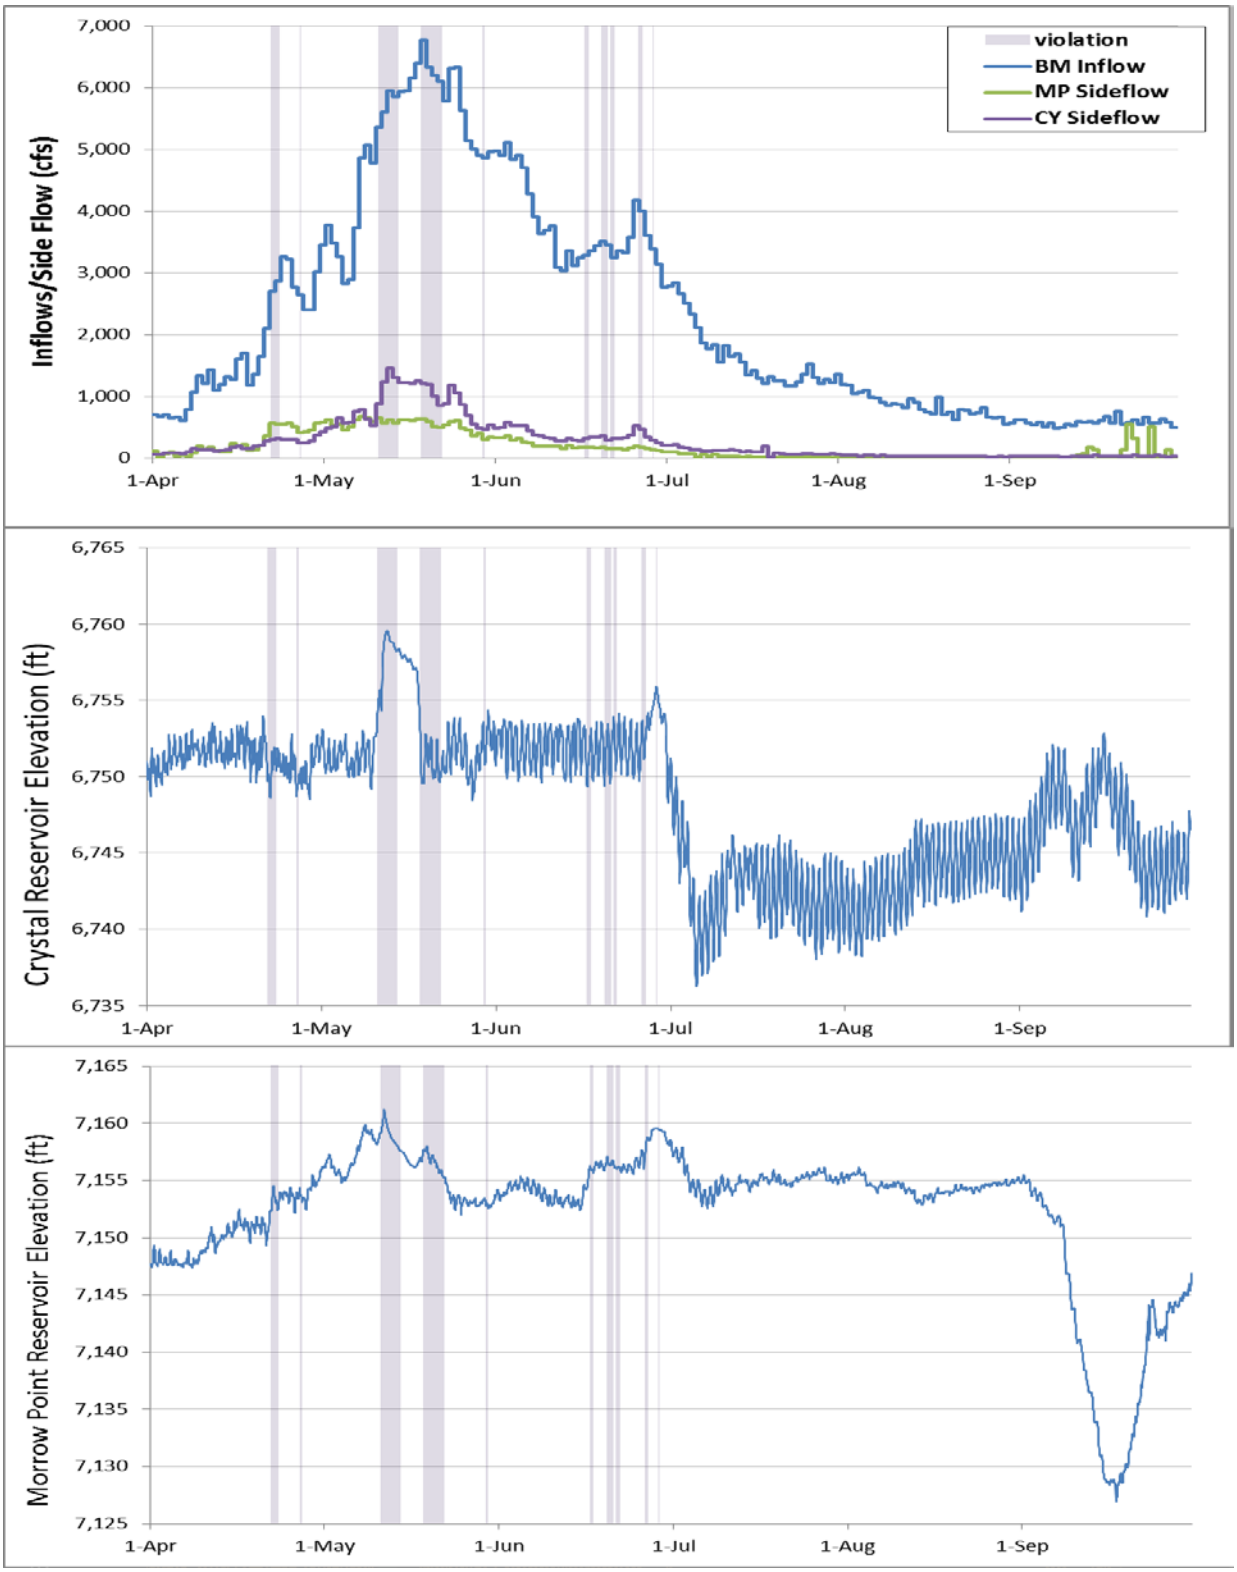

FIGURE K-1 Stretches of Time during the Retrospective Analysis in which Historic Information Indicates Reservoir Constraints Were Violated at Crystal, Together with Inflow/Sideflow Amounts and Reservoir Elevation Levels

APPENDIX L:

COMPARISONS OF CHEERS MILP AND NONLINEAR POST-PROCESSOR RESULTS

As discussed in previous sections, Conventional Hydropower Energy and Environmental Resource System (CHEERS) mixed-integer linear program (MILP) results for water release rates through various channels at a reservoir can optionally be input into the Nonlinear Post-processor Routine. Recall that the Nonlinear Routine is a more accurate representation because MILP requires that all equations be linear. Releases are used in a series of nonlinear equations and a convergence algorithm to compute reservoir water elevations and unit-level generation levels. Figure L-1 compares hourly reservoir elevations for April 20, 2009, computed by the CHEERS MILP and the Nonlinear Routine. The CHEERS run assumes a maximum of one start/stop cycle per day, with start and stop costs of $50 each. For all practical purposes, computations of elevations at all three reservoirs made by both methods are identical.

A comparison of hourly unit-level generation for April 20, 2009, is shown in Figure L-2. Generation levels are nearly identical for most hours. However, there are a handful of hours in which unit-level generation levels differ by 1 MW or less. The largest errors tend to be periods when generation levels are the largest. Table L-1 shows that, over the 24-hour period, total daily generation levels for the Aspinall Cascade are only 3 MW apart, a difference of less than 0.13%.

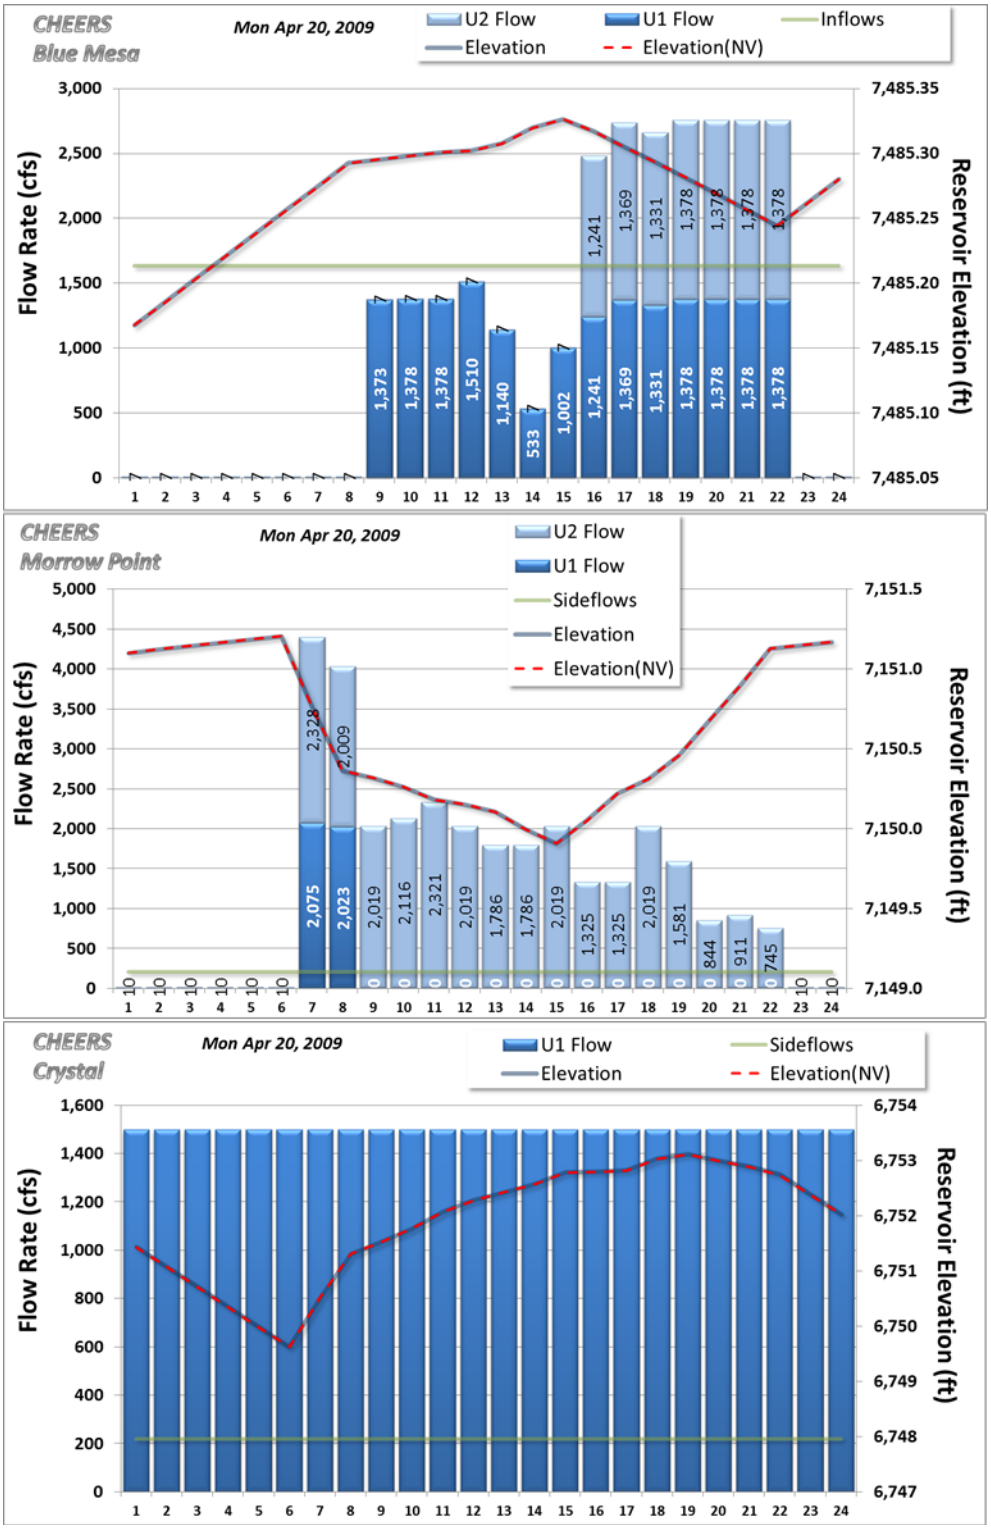

FIGURE L-1 Comparison of Hourly Aspinall Reservoir Elevations Computed by the CHEERS MILP and the Nonlinear Routine (April 20, 2009)

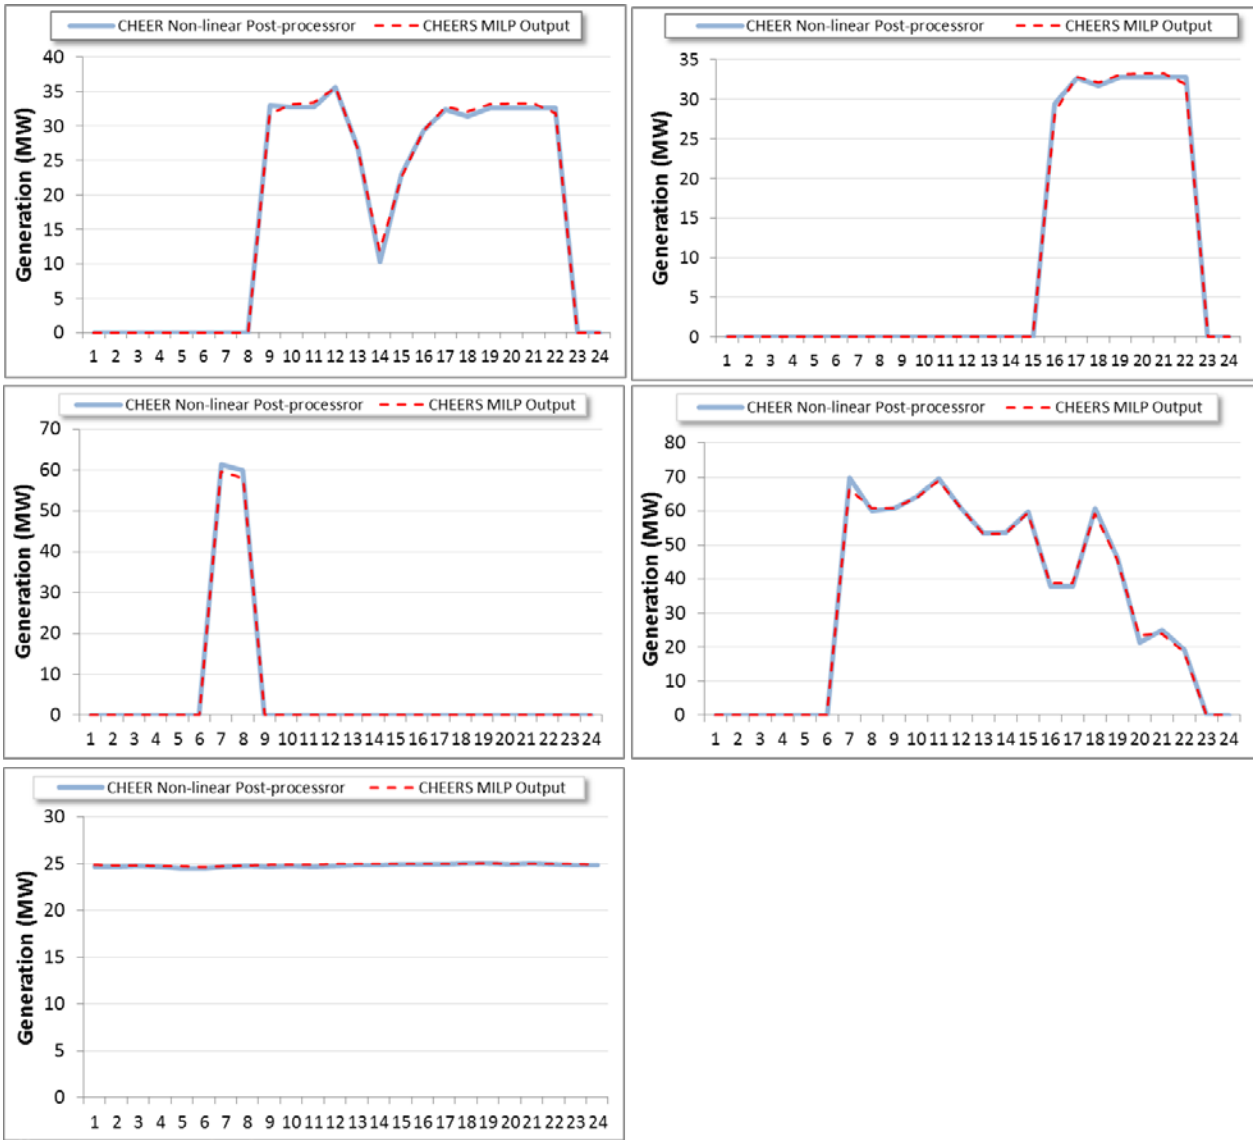

FIGURE L-2 Comparison of Hourly Unit-Level Generation Levels Computed by the CHEERS MILP and the Nonlinear Routine (April 20, 2009)

TABLE L-1 Comparison of Total Daily Unit-Level Generation Levels Computed by the CHEERS MILP and the Nonlinear Routine (April 20, 2009)

Parameter	CHEERS MILP Output (MW)	CHEERS Nonlinear Post-Processor (MW)
BM 1	420.8	417.9
BM 2	224.9	224.8
MP 1	117.6	121.3
MP 2	796.5	800.9
CY 1	596.9	594.8
Total	2,156.7	2,159.7

APPENDIX M:

STREAMFLOW ASSIMILATION IN EHFS

Uncertainty in forecasted flows stems from model errors due to the representation of the physical world in model space, errors in the forecasted meteorology, and the moisture state that is used in the hydrological model as the initial state at the start of the forecast period. Because water stored as soil moisture and snow creates significant memory in hydrological systems, this initial state often provides much of the predictability, especially for short-lead-time hydrological forecasts. This also means that errors in the initial state can have a significant negative impact on hydrological forecast skill.

Errors in the initial state result from model errors and errors in the observed meteorological forcings. The latter can be particularly problematic because a systematic bias in, for example, precipitation can lead to a long-term bias in the initial model state used in the forecast process. Because streamflow represents the integrated hydrologic response of a basin, streamflow observations contain indirect information about the moisture state of the basin. These observations can potentially be used to correct modeled soil moisture storage values and provide better initial states for forecast generation. Thus, assimilation of observed streamflow to update the modeled soil moisture state has the potential to improve forecast skill, especially at short forecast lead times.

We have implemented a streamflow assimilation procedure that uses near-real-time observed streamflow to update the model's initial moisture state. In this first implementation, soil moisture values in the deepest model soil layer, the layer responsible for the generation of baseflow, are replaced by direct substitution with values estimated from the observed streamflow. Soil moisture values are updated for all areas upstream of a given streamflow gauge in such a way that the spatial patterns of wetter and drier cells are preserved.

In EHFS, baseflow is modeled as drainage from a linear reservoir when the moisture storage is less than a threshold amount of storage. Above this threshold, baseflow generation is modeled as a nonlinear reservoir. In addition to this baseflow generation mechanism, the model has a fast response runoff mechanism that generates flow based on a parameterized saturation curve for the top soil layer. It is this latter curve that gives the model its name, but flow is generated through this mechanism mostly during high-flow events.

The assimilation procedure consists of two parts. The first part establishes a relationship between modeled streamflow and the spatial mean moisture storage in the deepest model soil layer (the third layer in our setup for the Gunnison and Feather river basins). This relationship is developed based on a retrospective model simulation. Figure M-1 shows an example of this relationship for the area upstream of the gauge at Blue Mesa in the Gunnison basin. A single third-order polynomial was used to model the relationship, with the same curve used for all

seasons. The values used to establish the curve are smoothed using an *n*-day forward-looking window to diminish the effects of timing errors on the assimilation process. Note that the update

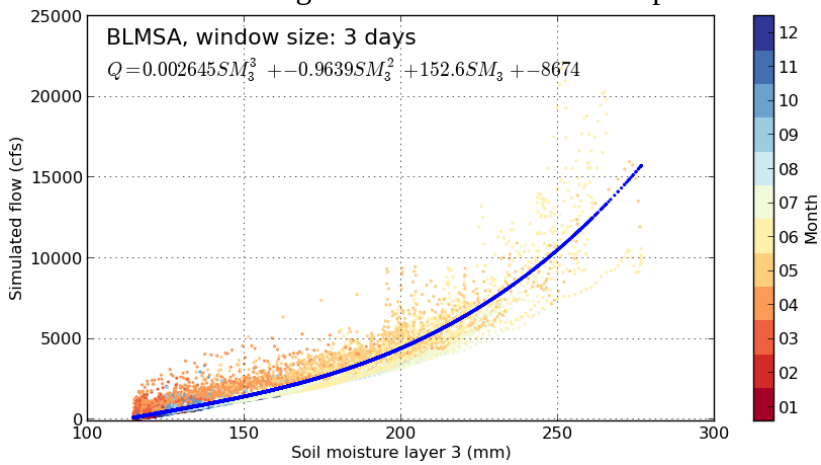

FIGURE M-1 Relationship between Simulated Streamflow and Spatially Averaged Model Soil Moisture in Layer 3 for the Area Upstream of Blue Mesa (BLMSA) in the Gunnison Basin

is performed a few days in advance of the forecast date because it will take some time for the updated soil moisture to impact the flow at the gauge.

This curve is then used to estimate model soil moisture from observed streamflow, so that the model simulated flow will be closer to the observed. During the forecasting process, the soil moisture updates are performed in the following way:

- Let t_f be the date of the forecast.

- The update will be performed on day t_f minus t_c, where t_c is the time of concentration for the subbasin for which the update will be performed. In this initial implementation, it is assumed that the time of concentration is the same for all the subbasins.

- Let *n* be the averaging period for the soil moisture and streamflow. Because we use a forward-looking window for the relationship between soil moisture and flow, we use $n = t_c + 1$ (basically the time period between the update and the forecast).

- For the *n* days preceding the forecast time (t_f), obtain the following:
  - Simulated soil moisture (SM_3) for layer three, averaged over all upstream grid cells; and
  - Observed flows at the gauge of interest gauge (Q_{obs}).

- Note that simulated, routed flows (Q_{sim}) are not strictly needed to perform the update; however, as explained later, they will be used anyway, not to determine the magnitude of the update but to provide a diagnostic on whether to perform an update at all.

- Calculate the *n*-day average for SM_3 and Q_{obs}.

- Solve for the updated value of SM_3 that will produce Q_{obs}. The resulting value is the soil moisture update SM_3^{up}.

- Let SM_3^{max} be the maximum SM_3 in the long-term time series. Then constrain SM_3^{up} so that $SM_3^{up} = \min(SM_3^{up}, SM_3^{max})$.

- Determine whether an update should be performed by comparing SM_3^{up} with SM_3 and Q_{sim} with Q_{obs}. Do not perform an update if this would increase existing bias, that is, if $Q_{sim} > Q_{obs}$ and $SM_3^{up} > SM_3$ or $Q_{sim} < Q_{obs}$ and $SM_3^{up} < SM_3$. If no update is to be performed, do nothing; otherwise, continue with the next step.

- Calculate the ratio (R) between SM_3^{up} and SM_3.

- Apply this ratio to the soil moisture in the third layer for each grid cell, with the constraint that soil moisture cannot exceed the maximum value. This update should be made in the state file.

- Restart the simulation using the updated state file on t_f minus t_c and create a new state file for the date of the forecast (t_f).

- Use the new state file to initialize the forecast.

Figure M-2 shows an example of the soil moisture update for three subbasins in the Gunnison Basin. Early results from the application of the assimilation procedure show that the assimilation tends to improve the correlation between the forecasted and observed flows and reduce the RMSE in the forecasts, but that it somewhat increases the bias. Reduction of the bias at short forecast lead times is the focus of current work to post-process the forecasted streamflow sequences.

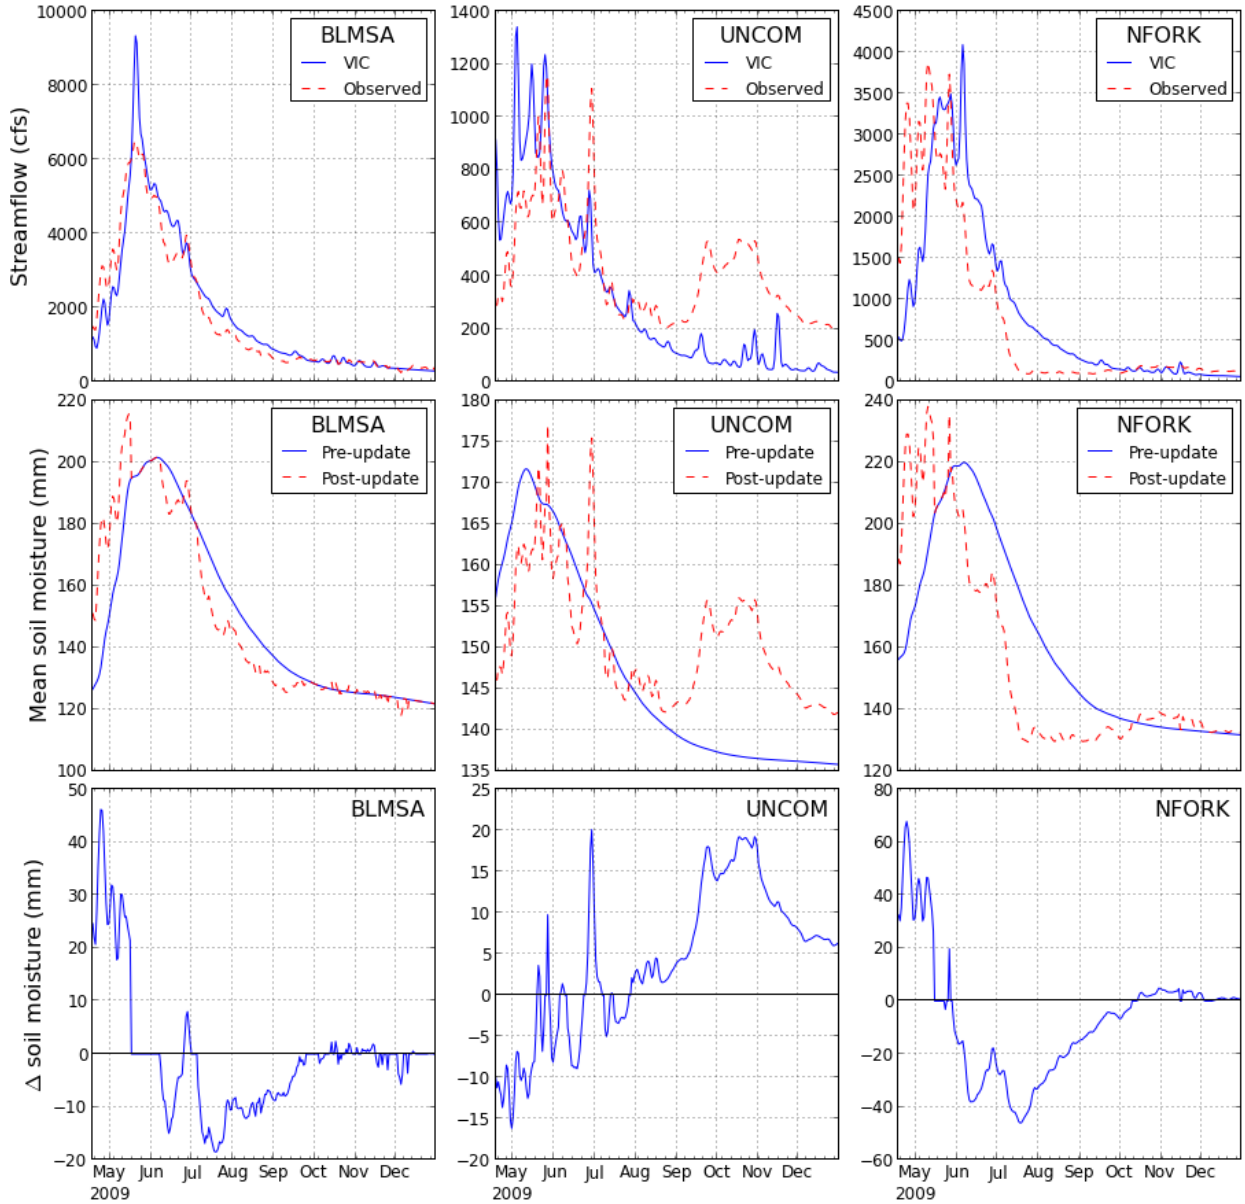

FIGURE M-2 Effect of Streamflow Assimilation on the Model Soil Moisture State for Three Subbasins in the Gunnison River Basin (BLMSA = Blue Mesa; UNCOM = Uncompaghre; NFORK = North Fork of the Gunnison. Top: observed and simulated streamflow for the period from mid-April through the end of 2009. Middle: soil moisture in the third layer pre- and post-update. Bottom: change in soil moisture in that layer as a result of the updates. As explained in the text, no update is performed when the update would increase the deviation of the simulated flow from the observed. Note that these updates were performed in a non-sequential manner to test the performance of the update algorithm.

Environmental Science Division
Argonne National Laboratory
9700 South Cass Avenue, Bldg. 240
Argonne, IL 60439-4847

www.anl.gov

Argonne National Laboratory is a U.S. Department of Energy
laboratory managed by UChicago Argonne, LLC