

**Many Objective Analysis to Optimize Pumping and Releases in a
Multi-Reservoir Water Supply Network**

by
Rebecca Smith

B.S., University of Alabama in Huntsville, 2011

A thesis submitted to the
faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Master of Science
Department of Civil, Environmental, and Architectural Engineering
2014

This thesis entitled:

Many Objective Analysis to Optimize Pumping and Releases in a Multi-Reservoir
Water Supply Network

written by Rebecca Smith

has been approved for the Department of Civil, Environmental and
Architectural Engineering

Joseph Kasprzyk

Edith Zagona

Balaji Rajagopalan

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Rebecca Smith (M.S., Civil, Environmental, and Architectural Engineering)
Many Objective Analysis to Optimize Pumping and Releases in a Multi-Reservoir
Water Supply Network

Thesis directed by Professor Joseph Kasprzyk

Past research has proven the utility of using multiobjective evolutionary algorithms (MOEAs) to optimize complex water management problems with many conflicting performance objectives. This study expands on the multiobjective optimization methodology by embedding a sophisticated RiverWare model in the algorithm search loop. The main challenges beyond linking the algorithm and model were due to the model's long simulation time and the complexity of the model. Addressing the simulation time necessitated several creative approaches to ensure efficient but thorough algorithm search and intelligent representation of hydrologic variability. Successfully addressing these issues confirms that advanced detailed operations models for civil infrastructure and water management can be used in MOEA-based multiobjective optimization.

The complexity of the Tarrant Regional Water District (TRWD) model offered additional challenges through which this study was able to gain further insight into the MOEA-assisted multiobjective optimization methodology. Initial objectives focused on system-wide reduction in pumping. Through the failure of the initial objectives to account for the performance of individual reservoirs, this study recognized that conflicts exist between objectives not only at a sub-system scale but also between system components and the broader system-wide objectives. Additionally, the incorporation of this information into a second problem formulation, which provided further system insights, confirmed that iterative problem definition is crucial to the decision making process.

The results obtained with this complex model suggest the need for further refinement of problem formulation, but also provided valuable information to TRWD. The implications of climate

forecasting and initial conditions within their model have a significant impact on the performance of suggested management alternatives, and may contribute to ambiguity in the relationships between the decisions made to balance and supplement reservoirs and the performance outcomes. This knowledge may inform TRWD's approach to optimization and decision making in the future and proves the value of the intermediate outcomes in multiobjective optimization.

Acknowledgements

I must first thank my committee members Dr. Joseph Kasprzyk, Dr. Edith Zagona, and Dr. Balaji Rajagopalan- a confluence of expertise that made this research possible. Beyond his professional expertise, Dr. Kasprzyk's support, patience, and good humor were invaluable to this research and my learning experience.

Many thanks are due to Laura Blaylock, hydrologist at Tarrant Regional Water District, for her enthusiastic participation in this research through many conference calls and emails. John Carron and Nick Mander of Hydros Engineering provided extensive model support, without which I would never have been able to properly understand and use the TRWD model.

Finally, thank you to the wonderful staff of CADSWES, who all had some impact on this research. I'm especially appreciative of Bill Oakley and Neil Wilson who were crucial in linking RiverWare and Borg, Jim Pasquotto who handled innumerable technical difficulties with patience and jokes, and Gwen Miller who was always available for support and sanity checks.

Table of Contents

Chapter 1	Introduction	1
Chapter 2	Case Study	5
2.1	Tarrant Regional Water District	5
2.2	TRWD Infrastructure	6
2.3	Operations Overview	7
2.4	TRWD Model	8
Chapter 3	Methods	12
3.1	Policy Variables	12
3.1.1	Bridgeport to Eagle Mountain Balancing (emzones & bpzones)	13
3.1.2	East Texas to Eagle Mountain Supplementation (emtrigs & emcrates)	14
3.1.3	Eagle Mountain to Worth Balancing (worthlevs).....	14
3.2	Model Variables vs. Algorithm Variables	15
3.3	Objectives, Problem Formulation 1	16
3.4	Objectives, Problem Formulation 2	19
3.5	Multiobjective Evolutionary Algorithms	20
3.6	Computational Experiment	22
3.6.1	Simulation Time	22
3.6.2	Hydrology	23
3.6.3	Embedding RiverWare	27
3.6.4	Borg Parameterization	28

3.6.5	Choosing an Appropriate Search Duration	30
Chapter 4	Results	32
4.1	Results: Problem Formulation 1	32
4.2	Results: Problem Formulation 2, stressed & surplus	35
4.3	Results: Problem Formulation 2, surplus	38
4.4	Results: Problem Formulation 2, stressed	39
4.4.1	Objective Performance	39
4.4.2	Connecting Policy Variables and Objective Performance	44
4.5	Impact of Climate State Forecast on Policy Variables	47
4.6	Hybrid Solutions	50
4.6.1	Creation of Hybrids	50
4.6.2	Hybrid Results in Varied Hydrology	53
Chapter 5	Concluding Remarks	56
5.1	Discussion and Future Work	56
5.2	Conclusions	61
References	63

Chapter 1 Introduction

Human works such as reservoirs, diversion channels, and water transfer schemes have facilitated the ability to move water to people. Many cities in the Western United States have taken advantage of this and created attractive destinations for an urbanizing population (U.S. Census Bureau 2012), but it is becoming increasingly difficult to meet growing demands in dry climates (Jacobs 2011; World Water Assessment Programme 2009). Traditionally, new supply sources were seen as the primary means to meet increasing water demands. However, to incorporate additional supply requires energy (for pumping or desalination) and/or infrastructure (e.g. new reservoirs or pipelines). Infrastructure is expensive and can take years to build, not to mention potentially be hydrologically impractical (Gleick 2003; Lund 2012), and the cost of energy has been increasing since the 1970s (World Water Assessment Programme 2009), contributing to the ever growing need for creative management solutions in the region and beyond (Gober 2013). Adding to the challenges are the as yet unclear implications of climate change; in the new reality of nonstationarity, wherein the past is not a reliable predictor of future hydrologic conditions (Milly et al. 2008), utilities are pressed to characterize the limitations of their current operations and find innovative ways to address deficits.

The circumstances facing the Tarrant Regional Water District (TRWD) are a perfect example of the confluence of these major challenges. TRWD supplies raw water to water treatment plants in the North Central region of Texas. The region has faced dry conditions in seven of the last 10 years (Blaylock 2014a), indicating the possibility that “dry” is the new “average”. Nonstationarity notwithstanding, the entire state of Texas has been in a drought of record since 2011 (Texas Water Development Board 2014a). TRWD owns several reservoirs in the East Texas climate zone, which exhibits relatively stable hydrologic inflows, but the large population centers served by the utility are located 50 miles or more west of the reliable water sources and up 400 feet of elevation. The energy required to move water to

these demands currently constitutes over 99% of TRWD's annual expenses, and the combination of drought and population growth in these areas has motivated the pursuit of additional supply and improved operations. TRWD has begun construction of a new pipeline to shore up supply, but substantial completion of the infrastructure is not expected before 2020. In the meantime, TRWD faces the challenge common to many utilities in the western United States (Western States Water Council 2012)- they must do what they can to improve the efficiency of their existing storage and delivery system under the combined stresses of increasing demands and less reliable supply.

As is the case for an increasing number of water utilities, modeling plays a major role in helping TRWD plan and make decisions. With the advancements in river and watershed management models in the last 15 years, the models are now able to provide managers a systematic way to explore "what if" hydrologic and demand scenarios and thoroughly vet management alternatives (J. W. Labadie and Baldo 2000; Matrosov, Harou, and Loucks 2011; Yates et al. 2005; Zagona et al. 2001). TRWD models its system in RiverWare, a generalized, detailed operational modeling platform (Zagona et al. 2001). RiverWare is often used with a graphical user interface (GUI), in which users can point and click to modify their system and visualize system outputs. The model also operates in so-called batch mode, where thousands of simulations can occur without invoking the GUI. Two examples of how TRWD uses its model are short- and long-term pumping and storage planning using stochastic hydrology and demands and development of climate forecasting functionality within the model.

The level of customization now available in models lends itself to very accurate representations of infrastructure dynamics and operational policies, creating the opportunity to use these models to speak for the actual physical system in suggesting new management alternatives. The historic approach to making water management decisions is for a limited number of options to be developed, tested, and evaluated based on an aggregated cost function (Harou et al. 2009), creating a situation where a "least bad" option is adopted. If instead, managers were able to enumerate their various performance goals

and “ask” the system how it could best meet each individual goal while simultaneously considering all objectives, the potential of the system (and the shortcomings) could be much better understood, and a more informed management decision could be made.

To be able to take advantage of management models as described above, it is necessary to be able to efficiently suggest and evaluate many alternatives. Multiobjective evolutionary algorithms (MOEAs) are well suited for this task because simulation models can be embedded within their search (Nicklow et al. 2010; Reed et al. 2013) . The MOEA automatically suggests alternatives, feeds them to the simulation model, and evaluates each alternative based on model outputs (system response). The ability of MOEAs to solve a variety of complex water management problems has been proven in past research, of which a few examples follow. In Kasprzyk et al. (2009) an MOEA was used in a multiobjective optimization analysis of tradeoffs associated with incorporating various levels of water market activity to manage urban water supply risks in Texas. In Mortazavi et al. (2012), researchers employed multiobjective optimization to characterize short- and long-term drought security options for Sydney, Australia, in light of the conflicting objectives of water security, economic costs, and environmental factors. Finally, Giuliani et al. (2014) featured MOEA-based many objective optimization in a framework designed to improve individual reservoir operation by replacing traditional rule curves with algorithm-assisted policy recommendations that incorporated many objectives and uncertainty.

Not only is MOEA optimization capable of finding creative solutions to water management problems with many conflicting objectives, it also lends itself to a process called “constructive decision aiding.” In constructive decision aiding, problem formulations and system representations are not static, but instead evolve as information about the problem (system) is gained through results (Woodruff, Reed, and Simpson 2013; Zeleny 1989). The incorporation of iterative problem formulation to refine multiobjective optimization results is an example of the concept that decision making can best be supported through multiple problem reformulations as new information is discovered and

incorporated. The benefits of this strategy were recently shown in the context of MOEA-optimized groundwater remediation when the initial problem formulation (set of goals and limitations) did not fully account for the creative, yet impractical, ways the algorithm would find to remediate the aquifer. By examining the shortcomings of initial results, more information about the remediation process was obtained and a better problem formulation developed that produced high quality, feasible results (Piscopo, Kasprzyk, and Neupauer 2014). The TRWD problem is an appropriate testbed through which to demonstrate constructive decision aiding because of its high number of complex interactions between different system components, and the conflicting management goals identified by the utility. The study presented in this thesis builds on Piscopo et al. by including the preferences and feedback of real decision makers within the TRWD system throughout the research.

The following document presents research which combined a powerful RiverWare model, MOEA-assisted multiobjective optimization, and close collaboration with a water utility to improve the management of an existing complex water storage and distribution system. A preliminary version of the decision making framework was presented at the American Geophysical Union Fall Meeting 2013 (Smith et al. 2013), and some results were presented at the Environmental Water Resources Institute congress in Spring 2014 (Smith et al. 2014). Results in this thesis will be developed into a journal article with a target submission date of Fall 2014. The thesis is organized as follows: a Case Study in chapter 2 will introduce the geographical region, infrastructure, operations, and model for the TRWD utility; chapter 3, “Methods”, will describe the details of the multiobjective problem formulations, provide background on MOEAs, and explain the mechanics of how the algorithm and simulation were linked; the Results and analysis of results are presented in chapter 4; finally, chapter 5 features a discussion of findings, future work, and conclusions.

Chapter 2 Case Study

2.1 Tarrant Regional Water District

This research was performed with the cooperation of the Tarrant Regional Water District (TRWD), a water utility located in the Trinity River Basin in North Central Texas, USA. As of 2011, Texas has been experiencing its most severe drought on record, and the past 24 months have been classified as “severely dry” for much of the Trinity River Basin (Texas Water Development Board 2014b). In addition to dry conditions, the state is expected to increase its population 33% between 2010 and 2030, with TRWD’s regional population increase over 37% in the same time span (Texas Water Development Board 2014a).

TRWD is a raw water supplier to 30 water treatment plants which span 11 counties. It is the second largest water utility in Texas and provides water to over 1.8 million people, including those in the densely populated cities of Fort Worth and Arlington. The region served by TRWD lies mostly in the North Central Climate Division, with the exception of one reservoir which lies in the East Texas Climate Division. The North Central Division is generally more dry and variable while the East Texas Division has more reliable reservoir inflows (Texas Water Development Board 2014b), a dynamic that has major operational implications for the management of TRWD’s system. Due to the relative abundance of water on the eastern end of the system and insufficiency in the west, TRWD is constantly pumping water from east to west, which includes up to 400 feet of elevation gain. This pumping requires a considerable amount of energy: 300 million kWh in 2012 at a cost of \$17.6 million (99.5% of the annual expenditures) (Deloitte & Touche, LLP 2013).

As a large purchaser of power, the pricing scheme for the utility is complex and dictates that they buy blocks of electricity months in advance, use what they buy or try to sell it on a spot market, and minimize any large variation in energy use lest they incur penalties from the power provider. These conditions mean that it is sometimes less expensive to pump more water than is demanded, and the

extremely dry conditions of late have set the stage for dramatic increases in pumping even when it's not required (Blaylock 2013).

Though TRWD is able to meet all of its customer demands currently, the utility has determined that additional water sources will be necessary to accommodate the projected demand growth. Several projects are in different stages of exploration, permitting, or design, but the first of these new supplies will not be available until at least 2020 (Tarrant Regional Water District 2013a). In the meantime, TRWD is actively pursuing methods of operating their current system to increase efficiency (TRC/Brandes 2009) while maintaining reliable service, an end to which this research contributes.

2.2 TRWD Infrastructure

TRWD owns four reservoirs: Bridgeport and Eagle Mountain at the extreme northwest end of the system, and Cedar Creek and Richland Chambers in the east. Note that while only Cedar Creek is technically within the East Texas Division, Richland Chambers is very close to the border and exhibits the stability of that climate zone. The utility has operational control over these reservoirs, meaning they make decisions about pool elevations and release times and rates. TRWD has contracts to store and use water from three reservoirs: Worth, Benbrook, and Arlington, all of which are owned by other entities and experience the variability associated with North Central Climate Division. These integrate into the TRWD system operations by way of specifications of how much they can store and release per year, how much supply they can access, and what pool elevations must be maintained.

To facilitate the east-to-west supplementation, TRWD currently has over 150 miles of pipeline connecting its reservoirs. The latest expansion of the pipeline added the Eagle Mountain Connection (EMC) in 2008, which linked the Eagle Mountain reservoir with the two eastern reservoirs, Cedar Creek and Richland Chambers. Refer to Figure 2.1 for a schematic of the major existing TRWD infrastructure.

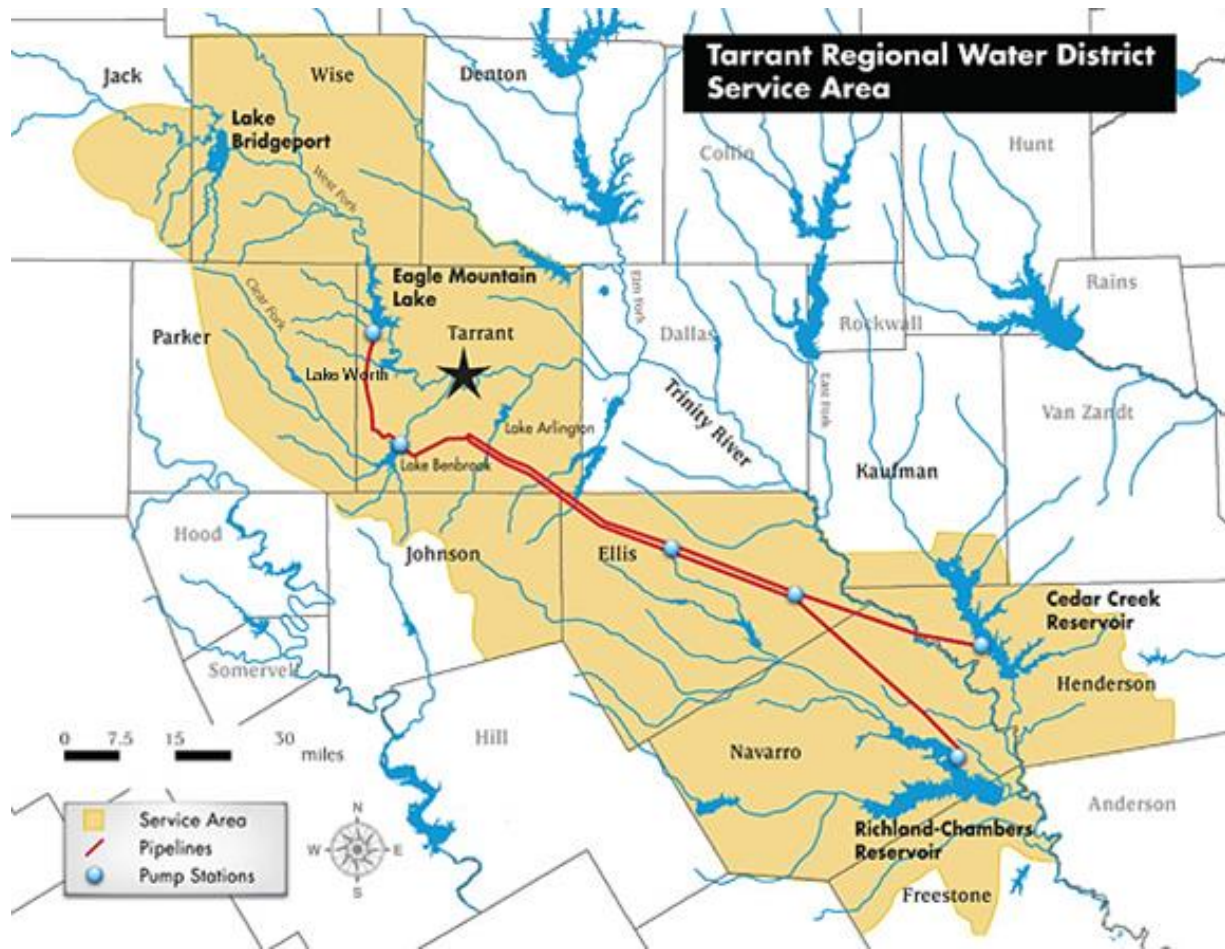


Figure 2.1 Map of the major existing infrastructure in TRWD's system (Tarrant Regional Water District 2013b).

2.3 Operations Overview

Bridgeport, Eagle Mountain, and Worth reservoirs are the only reservoirs in the TRWD system that are hydrologically connected by the West Fork Trinity River, and TRWD balances the levels in these lakes in the context of the fact that only Eagle Mountain is connected to the pipeline that distributes water from East Texas throughout the system. TRWD is the operators of Lake Worth, and is contractually obligated to keep its storage level at or above 590 feet above Mean Sea Level (MSL).

Use of Benbrook and Arlington reservoirs allows TRWD to meet the demands of customers in the vicinities of those lakes, and they are constantly supplied water from Richland Chambers and Cedar Creek to do so. Throughout the year, TRWD tries to meet specific target elevations for both Benbrook

and Arlington to ensure the maintenance of various contractual stipulations that accompany the rights to use the reservoirs. Because there is only one pipeline that moves water from the East Texas reservoirs to the rest of the system, and the EMC is a recent extension that connects Benbrook and Eagle Mountain, there are times when Eagle Mountain is cut off from the supplementation because Benbrook is pumping to the east. This occurs during approximately six months of the year when either one of the East Texas pumps is down for maintenance or the Rolling Hills water treatment plant requires extra water which is contractually required to come from Benbrook.

The pumps that are responsible for moving water from Richland Chambers and Cedar Creek are limited to six possible sets of pump rates that correspond to the most energy efficient way of operating them. Demand on the water from these reservoirs is calculated monthly and monthly pumping rates are correspondingly set to whichever of these six settings will completely meet that demand. This often means that more water than demanded is pumped west, and the “excess” is divided among three “terminal storage” reservoirs- Eagle Mountain, Benbrook, and Arlington based on the available room and the priority of each reservoir.

2.4 TRWD Model

To facilitate long- and short-term planning and exploration of new management policies, TRWD uses RiverWare to model its system. RiverWare is a generalized river basin modeling tool that is capable of modeling any river basin via features including object-orientation, in-model representation of a wide range of physical processes, multiple solvers, and a language to customize operating policy. RiverWare’s graphical interface facilitates modeling and communication of system properties, and because the software provides the framework for modeling any physical system, utilities and organizations can build upon their existing models over time instead of starting from scratch when there are significant infrastructure or operations changes (Zagona et al. 2001). TRWD’s RiverWare model is a detailed,

realistic depiction of their system implemented in RiverWare, hereafter referred to as the TRWD model.

A screenshot of the TRWD model is shown in Figure 2.2.

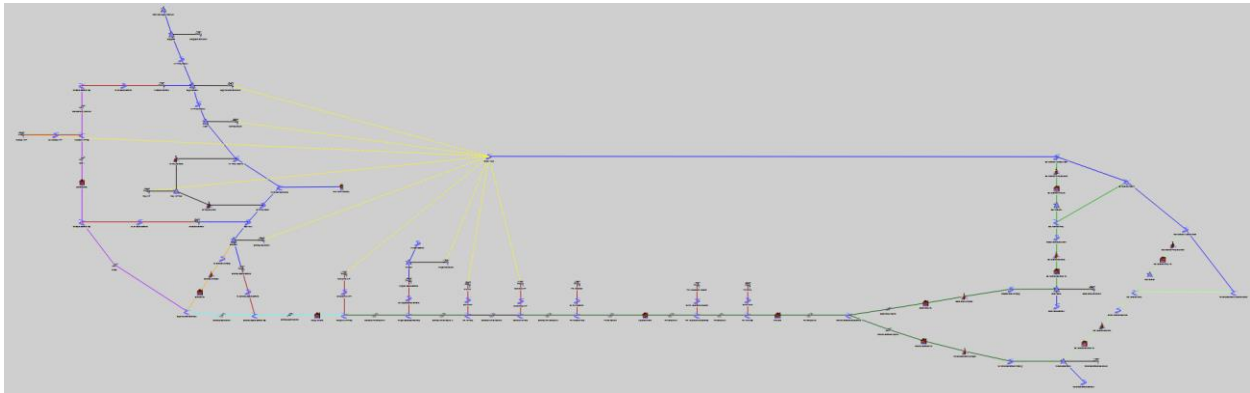


Figure 2.2 Screenshot of the TRWD RiverWare model in simulation view.

The TRWD model uses both Accounting and Rule Based Simulation at a daily timestep to mimic the utility’s operational policies, relying on over 400 custom rules and functions to precisely execute their system management. These rules iteratively evaluate reservoir levels, contractual requirements, pump settings, forecasted climate state, and individual accounts’ demands each month to determine the necessary daily releases and pump settings to meet 100% of the customer demands. Each reservoir has multiple accounts, or “pots” of water from which the various water treatment plants’ demands are met so that a detailed balance of water ownership and delivery can be maintained. This functionality and organization adds another layer of complexity to the model and system operations; for example, the Westside Water Treatment Plant demands can only legally be met by Eagle Mountain, *except* when the Eagle Mountain Connection is delivering water from East Texas to supplement Eagle Mountain, and then Westside’s demands may be met from Lake Benbrook, but only from a specific account that is contractually limited to delivering 72,500 acre-feet or less per year. The model has evolved over the past decade as new contracts, policies, and infrastructure have been added, with Hydros Engineering, Inc. providing modeling consulting and support.

One of the most substantial policy advancements within the model was the inclusion of the climate forecasting functionality, which the consulting firm AMEC studied and implemented in 2010 (AMEC 2010). It relies on rules which reference transition probabilities developed from historical data to forecast what the next month's climate state will be (dry, average, or wet), and references corresponding table values which attempt to accommodate the projected state. This functionality will bear on the policy variables, presented in Table 3.1 **Error! Reference source not found.** in the following chapter, and is described in detail in Section 4.5.

For planning purposes, TRWD runs its model from one to three years, and does so using stochastic hydrology and demands developed by the engineering firm Hydros. Note that the TRWD RiverWare model uses operational rules to run the water resources system, instead of prescribing narrowly-defined water releases as decision variables. Therefore, TRWD's operations are benefited by the use of a suite of stochastic reservoir inflows and customer demands in order to provide a type of "stress test" for their operational rules. The stochastic hydrologic traces created by Hydros are designed to capture a full range of possible hydrologic conditions from the historical record. The utility may run up to 100 traces comprised of hydrologic inflows and evaporation rates for each reservoir as well as a timeseries of ratios that alter the water treatment plants' demands.

The traces were produced using the aforementioned transition probabilities based on reservoir inflows from the historic period spanning 1941-2008. The inflows were derived by performing a mass balance using a combination of observed monthly streamflow data, inflow records maintained by TRWD, and precipitation and temperature readings from the National Oceanic and Atmospheric Administration. By referencing the transition probabilities, a Markov-chain method was used to produce 100 sequences of dry, average, and wet states. In other words, the stochastic hydrology assumes that hydrologic conditions are fully described by one of three climate states, and transition probabilities were also generated to model how the system changes from one state to another. After the sequences were

generated, they were used to resample daily data one month at a time from the historic period spanning 1980 to 2008 (AMEC 2010). For example, if the sequence of monthly states calls for a dry January that was preceded by an average December, the month of daily inflows from a dry January that followed an average December is randomly sampled from the historic data. These sequences and resampling can be regenerated depending on the chosen starting month and initial climate state of the simulation. For this study, the sequences are all generated using a dry October initialization and the simulation is run for one year.

Chapter 3 Methods

This chapter describes the methods and tools used in the study. First, there is a detailed description of the components of the problem formulation- decision variables (values that a decision maker can modify), objectives (performance metrics), and constraints (limits on acceptable performance). Next there is a section about the algorithm, the Borg MOEA. Finally, the software link, and Borg parameterization, and hydrologic ensemble selection are described.

3.1 Policy Variables

This section will begin discussing the many objective decision making framework that was created in this research for the TRWD system. The optimization problem formulation requires decision variables, or policy values that TRWD can change to alter system performance. It should be noted that in much optimization literature the term “decision variables” refers to physical actions taken by decision makers, such as specific volumes of water to transfer in a given scenario (Watkins Jr. and McKinney 1997) but for this study the term reflects values within the model that affect how model operational policies are carried out. For this reason, the term “decision variables” will be exchanged for the more accurate descriptor, “policy variables”. When choosing policy variables for the system, a significant goal of this research is to select variables, and therefore management policies, that are actually feasible for TRWD to implement. As such, the policy variables in Table 3.1 below correspond to operational variables within the model that TRWD confirmed as values that they are able to adjust unilaterally, accompanied by the upper and lower limits within which TRWD is comfortable operating¹. The Baseline values are those currently in use by TRWD, and the performance to which this study will compare its results. The variables are all values found in four different model tables that the custom-written rules

¹ Note that the upper and lower ranges of the variables in table Table 3.1 reflect ranges within which TRWD is comfortable operating. Decision variables within MOEA analysis also require upper and lower bounds, and these bounds are reflective of what the management policies result in after a series of decision variable transformations have been made.

refer to in order to carry out operating policy. The functional roles of these variables within the model are explained in the next section.

Table 3.1 List of 24 policy variables, their physical limits, what they correspond to in the system, and whether they are must be in ascending or descending order to properly operate.

Decision Variable	Lower Limit	Upper Limit	Baseline Values	Description	Ascending/Descending
Bridgeport to Eagle Mountain Balancing					
emzone1	644.1	648.1	644.1	Eagle Mtn pool elevation (ft)	Ascending
emzone2	644.1	648.1	644.1	Eagle Mtn pool elevation (ft)	Ascending
emzone3	644.1	648.1	648.1	Eagle Mtn pool elevation (ft)	Ascending
bpzone1	811	836	821	Bridgeport pool elevation (ft)	Ascending
bpzone2	811	836	826	Bridgeport pool elevation (ft)	Ascending
bpzone3	811	836	836	Bridgeport pool elevation (ft)	Ascending
East Texas to Eagle Mountain Supplementation					
emtrigdry1	641.1	648.1	643.1	Eagle Mtn pool elevation (ft)	Ascending
emtrigdry2	641.1	648.1	645.1	Eagle Mtn pool elevation (ft)	Ascending
emtrigdry3	641.1	648.1	647.1	Eagle Mtn pool elevation (ft)	Ascending
emtrigav1	641.1	648.1	641.1	Eagle Mtn pool elevation (ft)	Ascending
emtrigav2	641.1	648.1	643.1	Eagle Mtn pool elevation (ft)	Ascending
emtrigav3	641.1	648.1	645.1	Eagle Mtn pool elevation (ft)	Ascending
emtrigwet1	641.1	648.1	641.1	Eagle Mtn pool elevation (ft)	Ascending
emtrigwet2	641.1	648.1	643.1	Eagle Mtn pool elevation (ft)	Ascending
emtrigwet3	641.1	648.1	645.1	Eagle Mtn pool elevation (ft)	Ascending
emcrate1	0	200	150	EMC pumping rate (mgd)	Descending
emcrate2	0	200	100	EMC pumping rate (mgd)	Descending
emcrate3	0	200	75	EMC pumping rate (mgd)	Descending
Eagle Mountain to Worth Balancing					
worthlev1	590	594	590	Worth pool elevation (ft)	Ascending
worthlev2	590	594	591	Worth pool elevation (ft)	Ascending
worthlev3	590	594	591.5	Worth pool elevation (ft)	Ascending
worthlev4	590	594	592	Worth pool elevation (ft)	Ascending
worthlev5	590	594	593	Worth pool elevation (ft)	Ascending
worthlev6	590	594	593	Worth pool elevation (ft)	Ascending

3.1.1 Bridgeport to Eagle Mountain Balancing (emzones & bpzones)

Bridgeport and Eagle Mountain reservoirs balance using three elevation zones. Model rules governing Bridgeport to Eagle Mountain releases dictate that whatever zone Bridgeport is in, it must release enough water to either put Eagle Mountain in the same zone (if Eagle Mountain is in a lower zone) or the next zone up. Consider two scenarios: 1) if Bridgeport is in zone 2 and Eagle Mountain is in

zone 1, Bridgeport will release the volume necessary to put Eagle Mountain into zone 2; 2) if Bridgeport is in zone 2 and Eagle Mountain is also in zone 2, Bridgeport will release the volume necessary to put Eagle Mountain in zone 3. The policy variables used to optimize this dynamic are the elevations in both reservoirs that delimit the three zones.

3.1.2 East Texas to Eagle Mountain Supplementation (emtrigs & emcrates)

Eagle Mountain reservoir gets supplemented by the East Texas reservoirs via the EMC. The way the EMC pumping rate is determined is by either a prediction and distribution of above-average user demands (the model rules allow the system to foresee whether the combined demand of Holly and Eagle Mountain water treatment plants will exceed 100,000 acre-feet), or by a table which associates a demand pumping rate with three different Eagle Mountain reservoir elevation triggers. The model rules compare the current Eagle Mountain elevation with values in one of three columns in the *Configuration.Eagle Mountain Trigger Levels* table². The column referenced corresponds to the model's predicted climate state- dry, average, or wet. For any current elevation, the policy sets the demand to be the pumping rate associated with the *next highest* elevation trigger level. Consider this scenario: the elevation of Eagle Mountain is 642 ft, which falls between two trigger levels, 641 and 643; the model will set the EMC demand as the pumping rate associated with the 643 trigger level, and this demand is lower than the demand associated with 641, because, logically, the higher the elevation of Eagle Mountain reservoir, the less supplementation it needs from East Texas reservoirs. The policy variables used to optimize this dynamic are the Eagle Mountain elevation triggers in all three climate states and the corresponding EMC pumping rates.

3.1.3 Eagle Mountain to Worth Balancing (worthlevs)

The releases from Eagle Mountain to Worth are dependent on the combined percent full of Bridgeport and Eagle Mountain, or West Fork percent full (WF%Full). The model contains a table,

² Italicized table names here refer to names used within the TRWD RiverWare model.

Configuration.Worth Maintenance Level, that associates a Worth elevation with a WF%Full. As an example, if the combined elevations of Bridgeport and Eagle Mountain produce a WF%Full value of 73%, Eagle Mountain will release the volume of water necessary to bring Worth from its current elevation (if it is lower) to the elevation associated with the WF%Full equal to 70%. The policy variables used to optimize the releases from Eagle Mountain to Worth are the Worth elevations in the table.

3.2 Model Variables vs. Algorithm Variables

As demonstrated in the previous section, values for elevation zones and triggers within this formulation often need to be generated in an ascending or descending fashion for consistency. This requirement provides a challenge to typical MOEA implementation, since decision variables within the optimization are often considered independent. To ensure that the format of the zone, trigger level, and pump rate variables maintains the necessary ascending or descending structure of the corresponding model tables, it was necessary to use a range of (0, 1) for the algorithm variables and transform those values into model input values, as was demonstrated in previous work (Zeff et al. 2014). The relationships are described below:

Equation 3.1

$$\mathbf{model\ variable} = (\mathbf{upper\ limit} - \mathbf{lower\ limit}) * \mathbf{algorithm\ variable}(s) + \mathbf{lower\ limit}$$

Using emzones as an example, which are ascending in the *Configuration.Eagle Mountain Zone Delineators* model table, the following transformations result if the algorithm variables for emzone1, emzone2, emzone3 are 0.3, 0.5, and 0.7, respectively:

$$emzone1 = (648.1 - 644.1) * (0.3 * 0.5 * 0.7) + 644.1 = 644.5\ ft$$

$$emzone2 = (648.1 - 644.1) * (0.5 * 0.7) + 644.1 = 645.5\ ft$$

$$emzone3 = (648.1 - 644.1) * (0.7) + 644.1 = 646.9\ ft$$

3.3 Objectives, Problem Formulation 1

The problem formulation for the MOEA requires defined objectives that quantify the system performance. In this section, we discuss the set of objectives chosen for the initial problem formulation. These initial four objectives were developed through multiple consultations with both TRWD and Hydros, who between them have expertise in the management of the system as well as the development and structure of the model. Based on these conversations, it was clear that TRWD’s most prominent concerns for the operation of their existing infrastructure centered on timing and volume of pumping from East Texas, and the initial objectives reflect those broad concerns. Though the water from East Texas meets demands throughout the system, this study focuses on the question “Can the interactions between the three westernmost reservoirs be optimized to improve the efficiency and stability of pumped supplementation?” The goal of multiobjective optimization is to optimize a vector-valued objective function, $\mathbf{F}(\mathbf{x})$, with multiple constituent objective functions. \mathbf{x} is the vector of policy variables (see Section 3.1) and Ω is the space of feasible policy variables, as defined in the following equations:

Equation 3.2

$$\mathbf{F}(\mathbf{x}) = (emc50, pump211, spill, pumpvar)$$

$$\forall \mathbf{x} \in \Omega$$

Equation 3.3

$$\mathbf{x} = (emzone1, emzone2, emzone3, \\ bpzone1, bpzone2, bpzone3, \\ emtrigdry1, emtrigdry2, emtrigdry3, \\ emtrigav1, emtrigav2, emtrigav3, \\ emtrigwet1, emtrigwet2, emtrigwet3, \\ emcrate1, emcrate2, emcrate3, \\ worthlev1, worthlev2, worthlev3, worthlev4, worthlev5, worthlev6)$$

The first objective, *emc50*, tracks the percentage of time over the simulated year that the pumping rate in the EMC is greater than 50 million gallons per day (mgd). Below, the expectation notation, $E[\]_t$, is used to denote average objective value over the ensemble of t hydrologic traces, and d indicates number of days simulated. *EM Outlet Demand Tap.Inflow* is the slot, or value, on the pipeline object within the model that corresponds to the rate of supplementation requested by the Eagle Mountain reservoir.

Equation 3.4

$$\text{Minimize: } emc50(\mathbf{x}) = E \left[\sum_{i=1}^d \left(\frac{EM\ Outlet\ Demand\ Tap.Inflow\ 1 > 50\ mgd}{d} \right) \right]_t$$

Utilizing thresholds to quantify performance over time is common in water resources management, where information about frequency of failure (or reliability) is often more meaningful than the information that is conveyed by the mean or variance of performance (Hashimoto, Stedinger, and Loucks 1982). For this system, flow through the EMC is the most energy-intensive supplementation since the water has to travel a further distance than when it is delivered to the other reservoirs, and it is desirable to minimize the volume of water pumped that far. Additionally, this objective would be highly responsive to the effects of the re-balancing of the three westernmost reservoirs through alterations in the policy variables.

The second objective, *pump211*, is another reliability objective that tracks the percentage of time that the total East Texas pumping rate is greater than 211 mgd. The 211 mgd threshold reflects the demarcation between low and high volume pumping recognized by TRWD, and the utility would like to minimize the frequency of high volume pumping. Below, *Richland Chambers Cedar Creek.Outflow* is the slot within the model that corresponds to the total volume pumped from the two East Texas reservoirs and the notation is the same as in Equation 3.4.

Equation 3.5

$$\text{Minimize: } pump211(\mathbf{x}) = E \left[\sum_{i=1}^d \left(\frac{Richland\ Chambers\ Cedar\ Creek.Outflow > 211\ mgd}{d} \right) \right]_t$$

While the first two objectives are direct measures of pumping, the third objective seeks to penalize unnecessary pumping by monitoring reservoir spills that might occur if reservoirs receive excess water. Specifically, objective three is to minimize the total annual rate of spill from four reservoirs: Bridgeport, Eagle Mountain, Worth, and Benbrook.

Equation 3.6

$$\text{Minimize: } spill(\mathbf{x}) = E\left[\sum_{i=1}^d (Bridgeport.Spill + Eagle Mountain.Spill + Worth.Spill + Benbrook.Spill)\right]_t$$

The final objective in the initial problem formulation is *pumpvar*, which seeks to minimize the pumping variance, or the squared deviation from mean pumping rate. As mentioned in Section 2.1, TRWD is required by their power provider to maintain stable energy use from day to day and month to month so that the electricity company can plan for their load, and failure by TRWD to do so incurs price penalties that are more expensive and longer lasting than the choice to pump more water (use more energy) than demanded in any given month. *Policy Analysis.ET Pumping Daily* is the model slot that keeps track of the daily East Texas pumping rate.

Equation 3.7

$$\text{Minimize: } pumpvar(\mathbf{x}) = E[\text{var}(Policy Analysis.ET Pumping Daily)]_t$$

Each evaluation of the TRWD system uses the full complexity RiverWare model. The custom rules within the TRWD model ensure that 100% of demands are met and physical limitations of the TRWD system are respected, while the internal operations of the RiverWare software are designed to maintain mass balances. Therefore, it was not necessary to specify any explicit constraints within the problem formulation³.

³ In classical operations research optimization modeling, such as linear programming, constraints are designed to maintain water balances and represent physical properties within systems. However, MOEA decision support is a simulation-optimization technique, where constraints such as mass balance are handled within the simulation model's internal logic, and not optimization constraints. Constraints in MOEA formulations are typically only used to set acceptable limits on performance, so it is reasonable to have a problem without constraints.

3.4 Objectives, Problem Formulation 2

The four objectives in the first problem formulation reflect performance metrics of the system as a whole. In analyzing results from this formulation, it was determined that while the four objectives were meaningful system performance measurements, they did not capture all of the dynamics that TRWD would be concerned about. The particulars of this outcome are discussed further in Section 4.1, but the ensuing additional objectives are described here.

All three of the additional objectives are reliability objectives that keep track of the pool elevations of three western reservoirs: Bridgeport, Eagle Mountain, and Worth. Thresholds were set for each reservoir, and the frequency of meeting their respective thresholds is calculated. For these three objectives the average or expected value of the traces is not used, but instead the value of the maximum failure percentage (or worst performing trace) is evaluated by the algorithm. The Bridgeport threshold is 811 ft MSL based on the level below which TRWD experiences release difficulties. The Eagle Mountain threshold is 644.1 ft MSL, an elevation necessary for recreation considerations such as boat ramps and docks. The Worth threshold is 590 ft MSL which is the elevation that TRWD is legally required to maintain per the contract with the City of Fort Worth that allows them to use the reservoir.

Equation 3.8

$$\text{Maximize: } \textit{bridgeport} - \textit{rel}(\mathbf{x}) = 1 - \max_t \left[\sum_{i=1}^d \left(\frac{\textit{Bridgeport.Pool Elevation} < 811 \textit{ ft}}{d} \right) \right]$$

Equation 3.9

$$\text{Maximize: } \textit{eaglemtn} - \textit{rel}(\mathbf{x}) = 1 - \max_t \left[\sum_{i=1}^d \left(\frac{\textit{Eagle Mountain.Pool Elevation} < 644.1 \textit{ ft}}{d} \right) \right]$$

Equation 3.10

$$\text{Maximize: } \textit{worth} - \textit{rel}(\mathbf{x}) = 1 - \max_t \left[\sum_{i=1}^d \left(\frac{\textit{Worth.Pool Elevation} < 590 \textit{ ft}}{d} \right) \right]$$

Table 3.2 summarizes all seven objectives with short descriptions and what formulation it was used in. Additionally, it highlights what factors affect each objective’s performance. Note the distinction between the first four objectives, which are the result of several contributing factors, and the three reservoir reliability objectives, which measure the performance of specific objects within the system.

Table 3.2 List of all objectives, a short description of their purpose, the problem formulation they were used in, and what system elements affect them.

Objective	Description	Formulation	Affected By
emc50	Minimize % of time pump rate through EMC is > 50 mgd	1 & 2	EM elevations, Holly WTP + EM WTP demands, demands of other objects to the east, East Texas pumping capacity
pump211	Minimize % of time ET pumping is “high volume”, i.e. rate > 211 mgd	1 & 2	All demands from all objects in between East Texas and Benbrook, pool elevations of Benbrook & Arlington, EMC demand
spill	Minimize annual rate of spill from BP, EM, Worth, Benbrook	1 & 2	Elevations of the 4 reservoirs
pumpvar	Minimize the variance of the ET pumping volume	1 & 2	All demands from all objects in between East Texas and Benbrook, pool elevations of Benbrook & Arlington, EMC demand
bridgeport-rel	Maximize % of time BP elevation is at least 811 ft	2	Bridgeport elevation
eaglemtn-rel	Maximize % of time EM elevation is at least 644.1 ft	2	Eagle Mountain elevation
worth-rel	Maximize % of time Worth elevation is at least 590 ft	2	Worth elevation

3.5 Multiobjective Evolutionary Algorithms

Because water resources systems serve many purposes and have to consider many aspects of performance that often conflict, multiobjective evolutionary algorithms (MOEAs) are useful tools for quantifying the tradeoffs between objectives. The algorithms employ the concept of Pareto optimality (or non-domination) to compare the performance of solutions across multiple objectives. A solution is Pareto optimal if no other solution exhibits improvement in any objective without sacrificing performance in another objective. The algorithms are “evolutionary” because the process of defining an approximation to the Pareto-optimal set produces new generations of solutions to evaluate based on

the traits of previous well-performing solutions (Deb 2001; Coello, Lamont, and Veldhuizen 2007; Nicklow et al. 2010; Reed et al. 2013).

The MOEA used for this research is a state-of-the-art algorithm called Borg (Hadka and Reed 2013). It is actually a combination of six separate search operators in a framework that uniquely incorporates several features introduced by other MOEAs to improve the search of complex solution spaces associated with many-objective (more than 3 objectives) problems. Descriptions of the key features follow.

Borg evaluates objective performance based on a concept called epsilon-dominance, first developed by Laumanns et al (2002), that allows users of the algorithm to select a resolution for each objective that defines significant performance improvement. This is especially useful in multiobjective problems since the units and sensitivities of each objective are often very different. The value chosen for epsilon tells the algorithm how much improvement must be achieved in any given objective before a solution is considered better than another. Intelligent definition of epsilons helps ensure diversity of solutions, and effectively limits the size of the archive. Building on this concept, Borg employs epsilon-progress, a condition that requires the algorithm to maintain a minimum amount of search progress to prevent a restart, which triggers an overhaul of the solution population to revive search if stagnation occurs. Restarts can also be triggered if at any time in the search process the population-to-archive ratio is not within an acceptable range, a modification of a feature called adaptive population sizing introduced by Kollat and Reed in their epsilon Non-Dominated Sorting Algorithm II (ϵ -NSGAII) (2006).

As mentioned above, Borg uses six different recombination operators, to vary the traits of good solutions in the formation of new alternatives. With each successive generation, Borg learns which operators perform well on a particular problem, and auto-adapts its focus to those operators, giving them more opportunity to produce offspring (though Borg does not ever stop using any operators completely). This adaptive discovery of key recombination operators is extremely important in that it

makes Borg applicable to many different kinds of problems without any prior knowledge about the characteristics of the solution space (Hadka and Reed 2013). In tests against other algorithms Borg performed well on challenging problems (Hadka, Reed, and Simpson 2012), and its use in several recent studies (Kasprzyk et al. 2009; Kasprzyk et al. 2012; Piscopo, Kasprzyk, and Neupauer 2014; Zeff et al. 2014; Giuliani et al. 2014) recommends it for use in this research.

3.6 Computational Experiment

The sophistication of RiverWare and the complexity of the TRWD model have motivated the need to find ways to minimize computational load in order to speed up optimization time. In the past, multiobjective optimization research with MOEAS has been performed using tens of thousands of simulation iterations, made possible thanks to models that evaluated in a matter of seconds or less, e.g. Kasprzyk et al. (2009). The time necessary to complete the optimization- that is, the number of evaluations multiplied by the time required for each simulation- was further reduced by the use of massively parallel supercomputing, allowing the user to develop tradeoff sets in short amounts of time (often in several hours). The sections below describe the computational challenges and the steps taken to address them.

3.6.1 Simulation Time

The TRWD model is set up to very intricately reflect the actual dynamics of the TRWD infrastructure and operations, a fact that encourages confidence in the very precise system information it provides. This introduces a challenge, however, for MOEA optimization, which relies on several thousand iterations of a simulation that takes about 80 seconds to complete.

The time necessary to run a one year simulation at a daily timestep was diagnosed by task: a full cycle of opening the model, simulating the year, and closing the model takes 79 seconds on a computer with 32 gb of RAM and 12 cores operating at 2.6 ghz, 77 of which are actual simulation run time. This means that the vast majority of the time is consumed simply due to the approximately 200 rules that

fire every month, and to substantially decrease the model simulation time would require a modeling overhaul (both infeasible and contrary to the goal of using their existing, “legacy” model). Because TRWD makes its short-term operational plans (the type this study aims to address) based on the model using a daily timestep, and runs each simulation for one, two, or three years, this research would be less valuable and less illustrative if it did not use the model comparatively.

3.6.2 Hydrology

To ensure that the results of the optimization are not tailored to one specific timeseries of hydrologic inflows, but rather are optimal in a range of plausible scenarios, the evaluation of each set of policy variables is based on their performance over an ensemble of hydrologic traces. A set of 100 stochastic traces was provided by Hydros (the development of which is described in Section 2.4), and from those, three separate ensembles were chosen.

Without the option to significantly reduce model simulation time, it became necessary to limit the amount of hydrologic variability included in the study. Note that this document refers to stochastic hydrologic traces, but each “trace” is actually made up of hydrologic inflows for seven reservoirs, evaporation rates for nine surface water bodies (the seven reservoirs plus two wetlands areas), and a demand ratio that adjusts the historical average demand for every water customer to match the simulated hydrologic conditions. Per Laura Blaylock (Blaylock 2014b), TRWD operates its system to manage under dry conditions, but the utility does want to be able to take advantage of large inflows when they occur. Furthermore, examining the results of optimization in the context of dry or wet hydrology facilitates analysis of what variables are most important in a given climate. To this end, it was determined that separate optimizations would be performed for two scenarios— one that tended to have more dry hydrology and one that had a larger amount of wet hydrology. However, recall that the climate states within the model’s internal forecasting change from season to season. Therefore it is difficult to classify a long-term “wet” hydrological signal from the climate forecasting alone. To alleviate

this issue, we created a new classification system for the long-term hydrology that is named “stressed” for hydrologic inflows that tend to be drier, and “surplus” for when there is more water available in the system. The method of classifying the traces follows below.

Classifying the stochastic hydrologic traces into dry, average, or wet based on inflows is not straightforward; total annual inflows for the entire system do not necessarily characterize the aspects of climate that most significantly impact the operation of this system. Considering the total annual inflows for just the West Fork (Bridgeport and Eagle Mountain) reservoirs is more pertinent to the controlling factors, but does not account for the timing of the inflows- a large event can skew an otherwise very dry year. In light of these challenges, the idea of characterizing the stochastic flows by how much they stressed the system *as it is now operated* emerged. In other words, we would use the state variables (storages) within the system instead of statistics of the hydrologic inflows. Per TRWD hydrologist Laura Blaylock (Blaylock 2014a), the fullness of the West Fork (WF%Full) is an important metric for the operation of their system: the model judges whether the system is in drought based on this value, and the releases from Eagle Mountain to Worth are based on this value. So, rather than make assumptions about how significant any particular hydrologic statistic is for the management of the system, characterizing hydrologic traces based on how little water is in the West Fork at the end of the year reflects the system’s ability to maintain pool elevations in Bridgeport and Eagle Mountain in *response* to hydrology. Thus three ensembles of traces were chosen based on system response: stressed, surplus, and random.

A one year daily simulation was run for each of the 100 traces using the baseline solution (TRWD’s current management scheme) and the timeseries of WF%Full output for each simulation. The plot of these 100 timeseries is represented by the solid lines in Figure 3.1. This model output was compared to the actual historical WF%Full for each year from 1980 to 2012, represented by the black dotted lines. The WF%Full at the end of each year spans the same range for model output and historic

data. Based on this evaluation, the 100 traces were classified by the final simulated WF%Full value, since it is reasonable to assume that a set of hydrologic conditions that yielded low storage values for the West Fork reservoirs would be considered “stressful” to the TRWD system.

An important consideration when performing water resources modeling is the condition of the state variables (e.g. system storage) at the beginning of the model run. Our work is based on the state of the system at the time the model was provided to us in October of 2013, so every simulation begins with the West Fork at about 53% full (informally, this means the system is very stressed at the beginning of the model run). Though the historic data show that the system starts in a wide range of WF%Full states, the range of modeled final states is similar to the range of historically observed final states. In other words, even with a low initial value of West Fork storage, it is possible for the West Fork to fill over the course of the modeled year. The fact that the simulation begins with depleted West Fork reservoir levels means that the system will be required to send significant East Texas supplementation in order to ensure that the western reservoirs do not drop below acceptable storage thresholds very early into the simulated year. This makes it more “difficult” for the system to perform well in the objectives used for this optimization, but such challenging conditions are a major concern for TRWD and occur frequently.

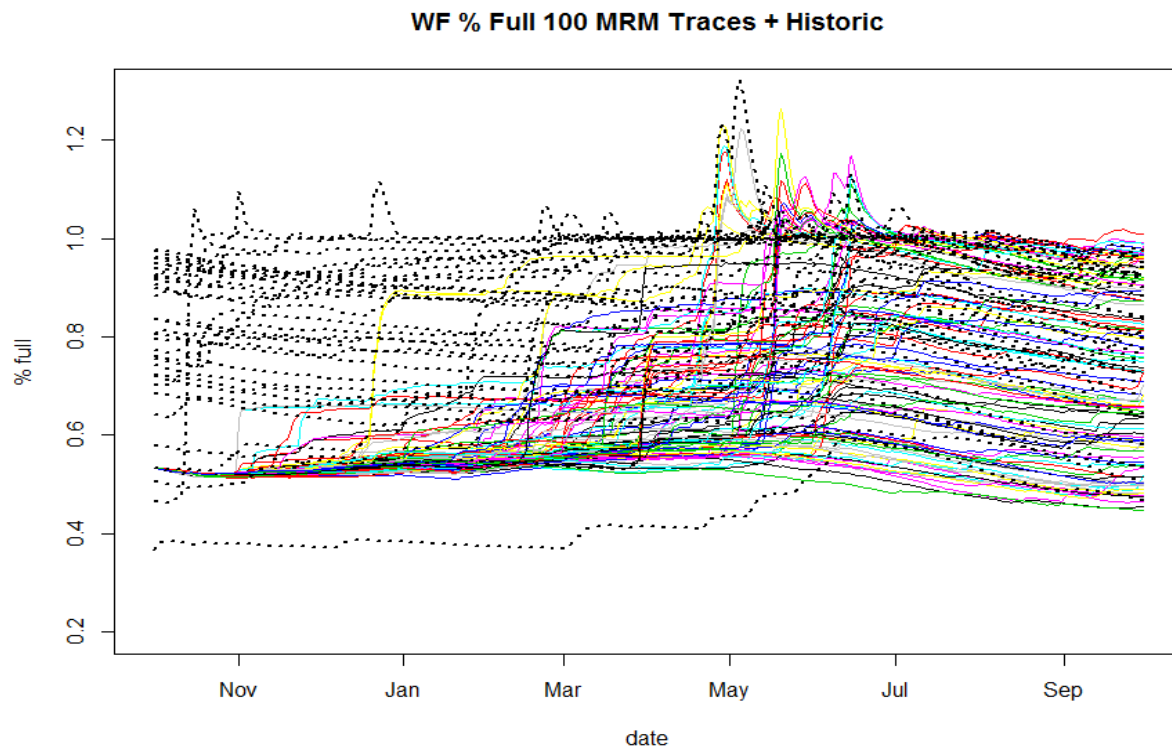


Figure 3.1 Plot of the combined percentage full of the Bridgeport and Eagle Mountain reservoirs (also known as West Fork Percent Full or WF%Full) for 100 stochastic simulations (color) and the observed from 1980-2013 (black dotted).

The definitions of stressed and surplus traces are based on determining that an average ending system state is anything within approximately one standard deviation, or 17.4%, of the mean (73.2%) of the 100 traces' final WF%Full values, resulting in a range of 55% to 90% full on the last simulated day. Thus, a stressed trace is any trace which has a final WF%Full value below 55% and a surplus trace is one that has a final WF%Full value above 90%. From these subsets of traces, 10 stressed and 10 surplus were randomly chosen and used for all stressed and surplus optimization simulations, respectively. The random ensemble is made up of 10 traces chosen randomly from the full 100. The resulting ensembles are plotted in Figure 3.2. Plots (a) and (b) are the timeseries of WF%Full for each trace, and below (c and d) are the West Fork inflows for each trace. The WF%Full stressed and surplus distinctions appear to be a reasonable way to characterize the dry conditions that are a high priority for TRWD and the large inflows that also need to be accounted for.

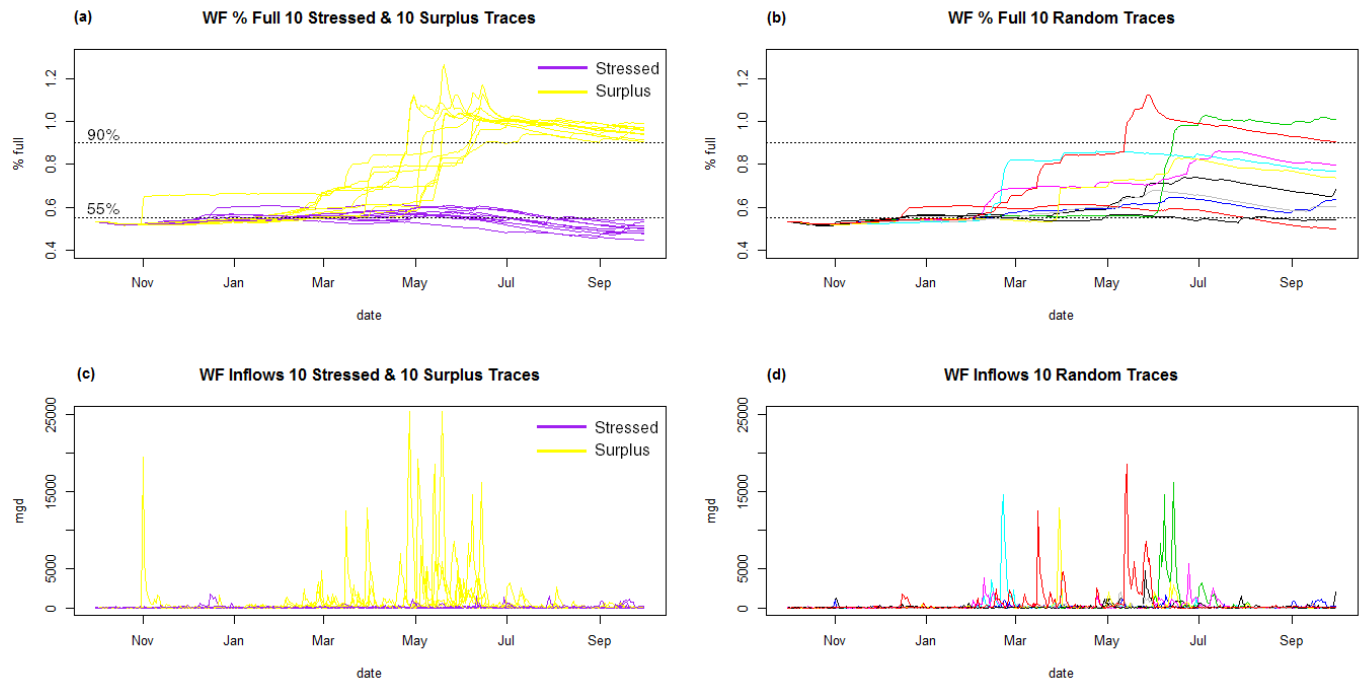


Figure 3.2 (a) West Fork % Full for the 10 traces in the stressed and surplus ensembles; (b) West Fork % Full for the 10 traces in the random ensemble; (c) West Fork inflows for the 10 traces in the stressed and surplus ensembles; (d) West Fork inflows for the 10 traces in the random ensemble.

3.6.3 Embedding RiverWare

Borg can be implemented in two ways: by embedding a simulation model within the search loop or by a loose coupling wherein the simulation and the algorithm execute separately but communicate via a standard input/output pipe. In this study, the TRWD RiverWare model is embedded within the Borg search loop via a “wrapper” executable written in C. The wrapper feeds solutions generated by Borg to RiverWare and objective output from RiverWare back to Borg (see Figure 3.3). This process takes advantage of several RiverWare features, perhaps the most important of which is concurrent Multiple Run Management (MRM); RiverWare is able to direct multiple simulations, or “runs”, of the same model concurrently and distribute them across multiple processors through the MRM. This allowed simulations for all 10 hydrologic traces to be run simultaneously from just one instance of RiverWare. By using a 12-core Windows machine, the time required for 10 simulations was approximately equal to the time required for a single run since each run was completed by its own processor. Each trace was loaded

from a set of text files called by an input DMI (data management interface) because, though RiverWare users commonly take advantage of the Excel input and output compatibility, the ability to reference individual text files instead of one (or 10) Excel files prevented competition between processors for file access.

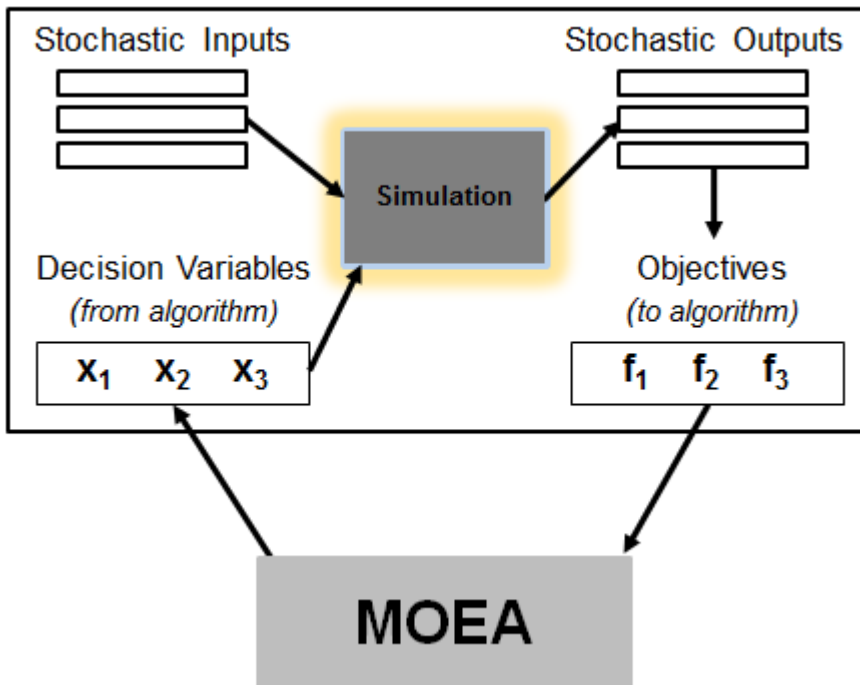


Figure 3.3 Depiction of the MOEA search loop with embedded simulation model.

Another convenient feature of RiverWare is the option to run it in “Batch Mode”, wherein the full version of a model is accessed from the command line without opening the user interface. The concurrent MRM and all input and output DMIs can be called from a batch script, so full simulations are executed as computationally efficiently as possible.

3.6.4 Borg Parameterization

For every problem formulation, decision variable ranges and objective epsilons (as well as information about constraints, if applicable) must be specified for Borg to operate. Decision variable ranges put bounds on the values Borg will suggest and reflect the feasible limits of operational policy variables within the model. All policy variables in this study were given a range of [0, 1] to account for

ascending and descending requirements, and the variables are transformed before being imported into RiverWare. The transformations are described in Section 3.2.

Each objective must have a corresponding epsilon precision so that algorithm can evaluate when significant improvement in objective function values has been achieved. To ensure meaningful optimization results, an epsilon value should be chosen with an understanding of how the values within an objective can vary (Laumanns 2002; Kasprzyk et al. 2009). For three of the four reliability objectives, emc50, bridgeport-rel, eaglemtn-rel, and worth-rel, the epsilon was set to approximately the change that would occur for each additional day of surpassing or failing to meet the respective thresholds, or 0.003. The fourth reliability objective, pump211, had an epsilon of 0.08, an approximation of 1/12. Because the East Texas pumping rates are only set once per month, any performance improvement can only be achieved at this increment due to TRWD operational policies.

Choosing the epsilons for spill and pumpvar required a more creative approach. To determine the range and variation of possible values for the objectives, objective outputs for all 100 of the stochastic traces (using the baseline values for policy variables) were examined. Because each of the optimization runs for this research were using the average over 10 traces to calculate these two objectives, the output from the 100 traces was randomly split into 10 sets of 10 values, and the mean of each set was taken. Characterizing these 10 means for each objective provided the basis for choosing a value of 5000 for both epsilons. This approach is similar to that of Kasprzyk et al. (2012), which used multiple evaluations of a noisy objective function to choose suitable epsilon values.

For all Borg parameters, default values were used (Hadka, Reed, and Simpson 2012; Reed et al. 2013), with two exceptions. Both changes were made to account for the long simulation time. The first change was to initial population size, which was reduced from 100 to 10 so that the evolutionary search would commence more quickly. The second change addressed the time between stagnation checks. Recall from Section 3.5 that Borg periodically checks to make sure that a reasonable amount of search

progress is being made, and if stagnation is detected, a restart is triggered. The frequency of the checks is determined by a variable called “windowSize”, with a default value of every 200 function evaluations, or simulation runs. Because of the extreme increase in model simulation time for this study it was deemed impractical to complete 200 function evaluations between assessments of search progress. In order to increase algorithm responsiveness to potential stagnation, a windowSize of 50 was chosen.

3.6.5 Choosing an Appropriate Search Duration

In order to find a high-quality approximation of the Pareto optimal set, MOEA search typically takes hundreds, or thousands of function evaluations. The TRWD model’s long simulation time necessitated conservation of computing time through an efficient (but sufficient) number of function evaluations on top of the selection of the number of stochastic traces used in each evaluation. After implementing all of the time-saving RiverWare options available, the total time required for each iteration of 1) generating values of policy variables, 2) invoking RiverWare, 3) running 10 concurrent simulations, and 4) returning objective output to Borg, was 105 seconds, which translates to about 800 function evaluations per day. Because prior diagnostic studies have shown Borg to efficiently and consistently converge (Reed et al. 2013; Hadka, Reed, and Simpson 2012), limiting function evaluations was not seen as very risky to the optimization outcomes.

To determine just how few function evaluations could be used, the optimization archive (the current set of non-dominated solutions) output was evaluated periodically throughout initial optimization runs via an animation function available in AeroVis (visualization software used in this and several other MOEA studies) (Kollat and Reed 2007). The ability to watch the search evolve essentially in real-time (Kollat and Reed 2006) allowed for a confident decision to limit each round of optimization to 3000 function evaluations. There were situations where the algorithm happened to be repopulating after a restart and not writing the archive file for several dozen to several hundred evaluations past 3000; in such cases, the final archive is the result of more evaluations, but 3000 (taking about four days

on a 12 core computer, running at 2.6 ghz with 32 gb of RAM) was determined to be the minimum. To be sure that no regions of the solution space were being missed due to limited function evaluations, several rounds were allowed to run for over 7000 evaluations, and while there was some added diversity for the extra evaluations, no substantial improvements were achieved.

Chapter 4 Results

This chapter presents analysis of the results of two different multiobjective optimization problem formulations as well as the results from an exploration of robustness between dry, wet, and hybrid optimization solutions in varied hydrologic conditions. Multiple ways of visualizing results are shown, including using RiverWare to provide in depth information about particular solutions, and analysis of the impacts of various policy variables are explored. The organization follows the learning process experienced by the researchers: initial results; determination of necessary adjustments to the set of objectives; results from the reformulation; exploration of how the policy variables impacted system performance; and finally, analysis of interactions between policy variables, model factors, and hydrologic ensembles.

4.1 Results: Problem Formulation 1

Problem formulation 1 evaluated solution performance based on four objectives: frequency of large supplementation to Eagle Mountain (emc50), frequency of high-volume East Texas pumping (pump211), spill, and East Texas pumping variance (pumpvar). The two plots in Figure 4.1 show the results of the optimization in three dimensions, or objectives; because the hydrology was so dry, spill was negligible to zero for all solutions. For both plots, pumpvar is represented by the x axis, emc50 is on the y axis, and pump211 is in color.

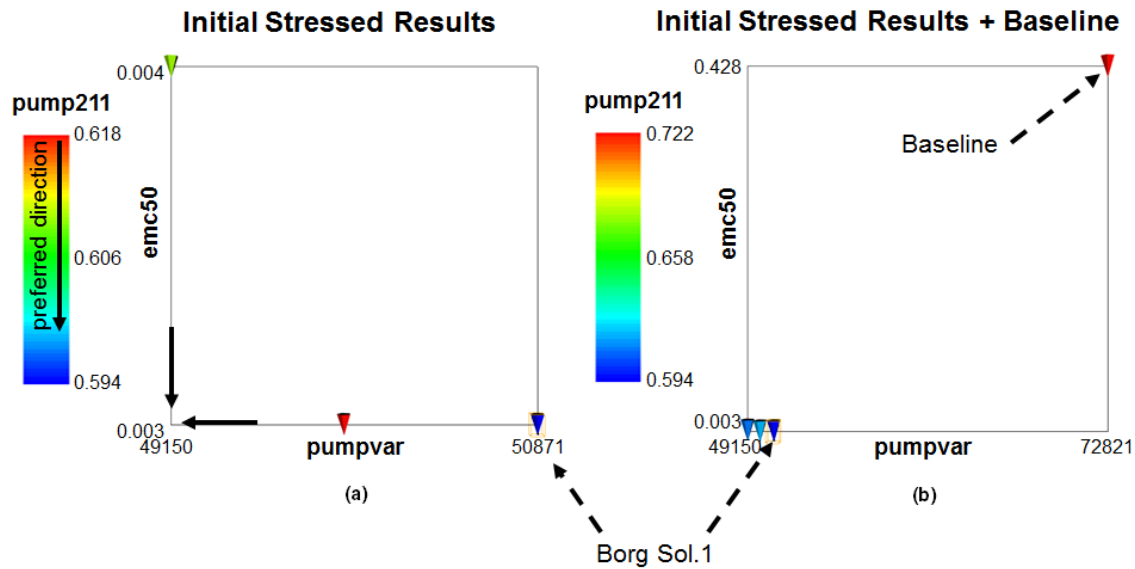


Figure 4.1 (a) The three solutions resulting from the initial stressed optimization plotted in three dimensions: emc50 on the y-axis, pumpvar on the x-axis, pump211 in color; (b) the three stressed solutions plotted along with the performance of TRWD’s current operations using the same orientation as (a). Borg Sol.1 and Baseline labeled.

Close inspection of plot (a) reveals that the ranges for all three objectives are very small: emc50 total range is 0.001, pumpvar total range is about 1721 (very small for squared values), and the range for pump211 is 0.024. The ideal solution would be dark blue, located in the bottom left corner, and the absence of this solution indicates that there are tradeoffs between objectives. Recall that for a solution to be nondominated, its performance in one objective cannot improve without degradation in another objective. While each of the three solutions generated by the Borg MOEA is nondominated in one of the three objectives, a central goal of multiobjective optimization is to gain information from a maximally diverse set of alternatives (Brill et al. 1990), and three solutions with qualitatively equivalent performance does not fulfill this goal. In other words, each of these solutions has low demand on the Eagle Mountain Connection (emc50), with a moderate amount of pumping variance, low spill, and a moderate amount of high volume pumping (pump211). They represent, qualitatively, only one way to manage the TRWD system.

The goal of formulation 1 was to improve the management of the system relative to current management practices, hereafter referred to as the baseline strategy. To highlight just how similar the three solutions from the optimization run are, refer to plot (b) of Figure 4.1, which plots the three solutions from the optimization run on the same axes as the objectives from the baseline solution. Note that pump211 still plots color; because the baseline has a much higher pump211 than the optimization solutions, it is shown in red in the plot. The plotting ranges have increased drastically to accommodate the baseline solution. Plot b illustrates that though the algorithm did in fact find solutions that outperformed current operations based on the objectives it was given, the solutions are not diverse. This outcome suggests in several ways that the problem formulation was insufficient: lack of diversity and volume of solutions is indicative that the chosen objectives did not conflict very strongly, and such extreme improvement over baseline implied that either TRWD had a baseline management strategy that was far from optimal in all objectives, or the objectives were not representative of some important system metrics that their operating policy implicitly considers.

To get more information about the solutions produced by the four objective stressed optimization run, they were loaded back into the TRWD model. Here, the value of optimizing using a complex RiverWare model is very clear: analysts and decision makers can run the simulation using the solutions suggested by Borg and explore detailed system performance beyond what was monitored via objectives. For decision makers, especially, it is incredibly informative to be able to visualize the implications of various alternatives using model output that they are familiar with (John W. Labadie 2004). Additionally, this capability greatly facilitates iterative problem formulation, which has been shown to increase the effectiveness of optimization in multiple studies (Kasprzyk et al. 2009; Kasprzyk et al. 2012; Woodruff, Reed, and Simpson 2013; Piscopo, Kasprzyk, and Neupauer 2014). Indeed, the complex, legacy model and the iterative problem formulation framework are crucial to this research; analysis of the solutions from the first problem formulation using RiverWare revealed that the

improvements in pumping-related objectives were gained by sacrificing storage in the western reservoirs, as demonstrated by Figure 4.2.

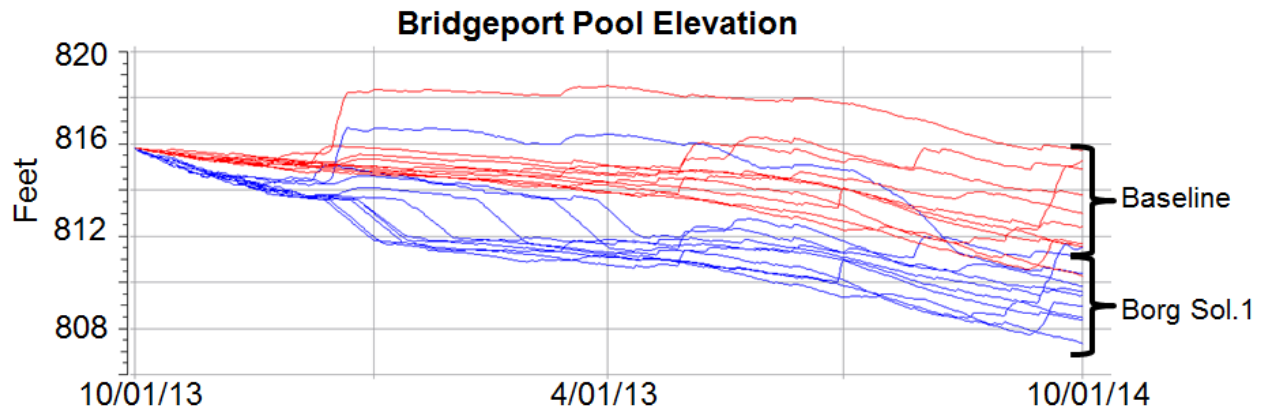


Figure 4.2 Plot of the Bridgeport reservoir’s pool elevation throughout the year for all 10 stressed traces under baseline operations (red) and using Borg Sol.1 (see Figure 4.1) policy variables (blue).

The pool elevation in the westernmost reservoir, Bridgeport, is plotted for all 10 stressed traces, with the baseline solution in red and the Borg solution in blue. By the end of the simulated year, the Borg solution has drained Bridgeport by four to seven feet compared to the baseline management strategy, depending on the trace, which places the entire western end of the system in a critical position. If Bridgeport is struggling to increase its pool elevation, it will not have water available to release to Eagle Mountain and causing repercussions for both the elevations of and balancing between Eagle Mountain and Worth. Adding reliability objectives for each of these three reservoirs would make the problem formulation more representative of TRWD’s management considerations (the pool elevations of these reservoirs are legally and operationally critical) and quantify the tradeoffs inherent in balancing the reservoirs as well as provide a more nuanced set of criteria for the algorithm to optimize.

4.2 Results: Problem Formulation 2, stressed & surplus

A fundamental aspect of successful multiobjective optimization is an intelligent problem formulation which can require several iterations (Piscopo, Kasprzyk, and Neupauer 2014). Indeed, the iterative process results from the learning inherent in multiobjective optimization, through which

decision makers can identify exactly what their performance priorities have been in the past and evaluate their importance going forward (Hitch 1960; Liebman 1976; Zeleny 1989). Problem Formulation 2 adds three reliability objectives to track the pool elevations of reservoirs Bridgeport, Eagle Mountain, and Worth, to the four objectives from the initial formulation. The volume and diversity of results increased dramatically, providing valuable information about performance tradeoffs inherent in managing the system.

The plot in Figure 4.3 (a) shows the results from both stressed and surplus optimization runs in six dimensions (spill was again omitted, though it does factor into the results of the surplus optimization). The reservoir reliabilities are on the x-, y-, and z-axes; pump211 is again on color; the orientation of the cones reflects the performance in emc50 (pointing straight down is best, straight up is worst); large cone size means large pumpvar and small size means less variance. The ideal solution would be small, blue, pointed down, in the left corner. The most obvious information conveyed by plot (a) is how dramatically the hydrology affects system performance: the solutions from the 10 stressed traces are all in high volume pumping more than 60% of the time (yellow to red) and spatially shifted away from the ideal reservoir reliability corner by 20-50% as compared with the surplus solutions. Note that the stressed results have many solutions with low pumping variance (i.e. they are plotted with a small size); because of the scarcity of water in the system, there is not much opportunity for reducing the East Texas pumping rate, so though the pumping rate is high, the variability is low.

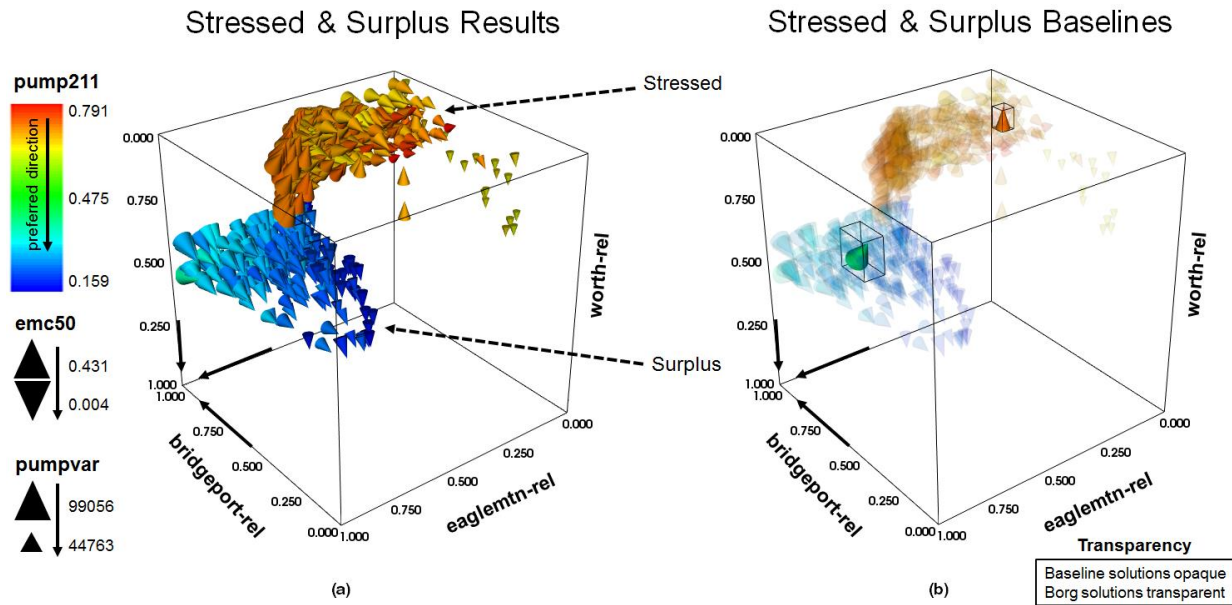


Figure 4.3 (a) plot of results of stressed and surplus solutions in six dimensions: bridgeport-rel, eaglemtn-rel, and worth-rel on spatial axes, pump211 in color, emc50 conveyed by cone orientation, and pumpvar conveyed by cone size; (b) plot of baseline performances for both stressed and surplus ensembles (opaque, boxed cones) against transparent optimized solutions.

Plot (a) also reveals objective tradeoffs, which result from conflicting priorities, most clearly illustrated in the dry results. The curved front located at the “back” of the dry results populates a tradeoff between eaglemtn-rel and Bridgeport-rel: improved performance in Eagle Mountain reliability cannot be achieved without lowering Bridgeport reliability. Similarly, the separate region of small yellow cones in the dry results illustrate that attaining greater than 20% reliability in worth-rel is only achieved in solutions where eaglemtn-rel is 0% (meaning Eagle Mountain reservoir never meets the 644.1 ft threshold).

Plot (b) is the same results oriented in the same way as plot (a), but with the baseline solutions shown with opaque cones and the Borg solutions plotted transparently. This was done to easily facilitate evaluation of the baseline solutions. The two boxed, opaque cones are the six-objective representations of how current operational policies compare to the optimization results. In the case of the surplus optimization, Borg found solutions that exhibit performance improvements over the baseline in every objective: cones that are smaller (less pumping variance), have colors that are toward

blue on the color spectrum (lower frequency of high-volume pumping) and pointed down (infrequent large EMC supplementation) positioned closer to the ideal left corner of the reservoir reliabilities exist within the optimization results.

Evaluating the baseline performance in the stressed results in Figure 4.3 plot (b), the only objective that did not get vastly improved upon was bridgeport-rel; the baseline achieves almost 100% reliability for Bridgeport pool elevation. However, within the optimization results, comparable bridgeport-rel performance *can* be achieved with better performance in emc50, pump211, and pumpvar with no degradation in eaglemtn-rel or worth-rel. Below, the results of the results from the stressed and surplus optimizations will be discussed separately.

4.3 Results: Problem Formulation 2, surplus

In Figure 4.4, the results of just the surplus optimization are presented in the same format as the above combined results (note different objective ranges); the optimal point would be small, dark blue, and located in the leftmost corner. The surplus optimization was performed to determine how the system performance and policy variables change when the West Fork has abundant water supply, which is guaranteed through the use of the surplus ensemble traces (refer to Section 3.6.2). The large red and green cones at the left edge plotted below show that surplus hydrology makes it *possible* to achieve very high reliability in Bridgeport and Eagle Mountain (impossible in the stressed ensemble results), but at the cost of increased high-volume pumping, greater pumping variability, and poor Worth reliability. Better performance in worth-rel requires sacrifices in both bridgeport-rel and eaglemtn-rel, but can combine with improvements in emc50, pump211, and pumpvar.

Surplus Results

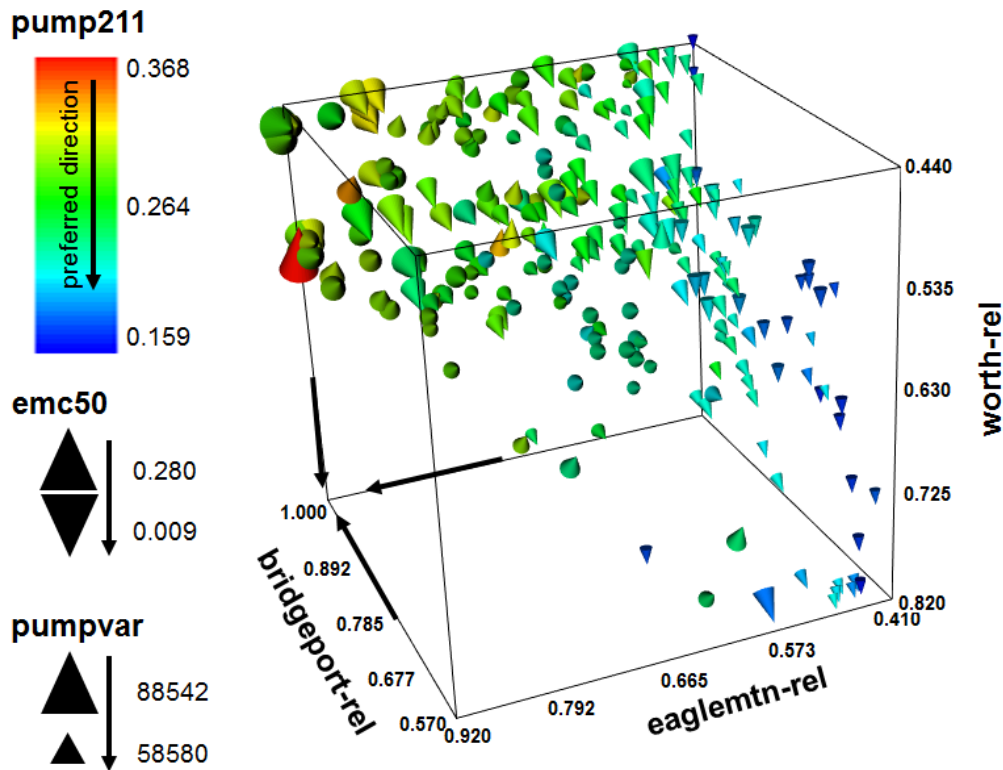


Figure 4.4 Plot of the results of the optimization using the surplus hydrologic ensemble in six dimensions: bridgeport-rel, eaglemtn-rel, and worth-rel on spatial axes, pump211 in color, emc50 conveyed by cone orientation, and pumpvar conveyed by cone size (spill is not plotted).

4.4 Results: Problem Formulation 2, stressed

Here, more detailed analysis of the stressed results is presented, including alternate visualization and comparison of specific solutions. The stressed results were obtained by optimizing with the stressed ensemble traces, or the traces which proved very dry for the West Fork (refer to Section 3.6.2). The stressed ensemble corresponds to the climate conditions that are of growing concern to TRWD, and for which it is focused on improving its system management.

4.4.1 Objective Performance

Figure 4.5, below, presents AeroVis plots of the stressed solutions in exactly the same scheme as the combination stressed and surplus plots, but the ranges for the objectives have changed. A small, dark blue cone pointed down in the left corner is still the ideal solution. Note that the stressed and

surplus results exhibit basically the same shape because the reliability conflicts between the three western reservoirs are present regardless of hydrology. This is intuitive because they are hydrologically linked; that is, the three reservoirs essentially must share their limited natural inflows, and the shortfall between their supplies and demands can only be met by supplementing Eagle Mountain, the middle reservoir in the series. The value in visualizing the tradeoffs present in these diverse results, however, is that TRWD can make an informed choice about how much performance to sacrifice in one reservoir to gain reliability in another. Other conflicts between objectives that were seen in the surplus results are present in the stressed results as well: good performance in the emc50 is associated with worse eaglemtn-rel; higher frequency of pump211 corresponds to better bridgeport-rel.

Stressed Results

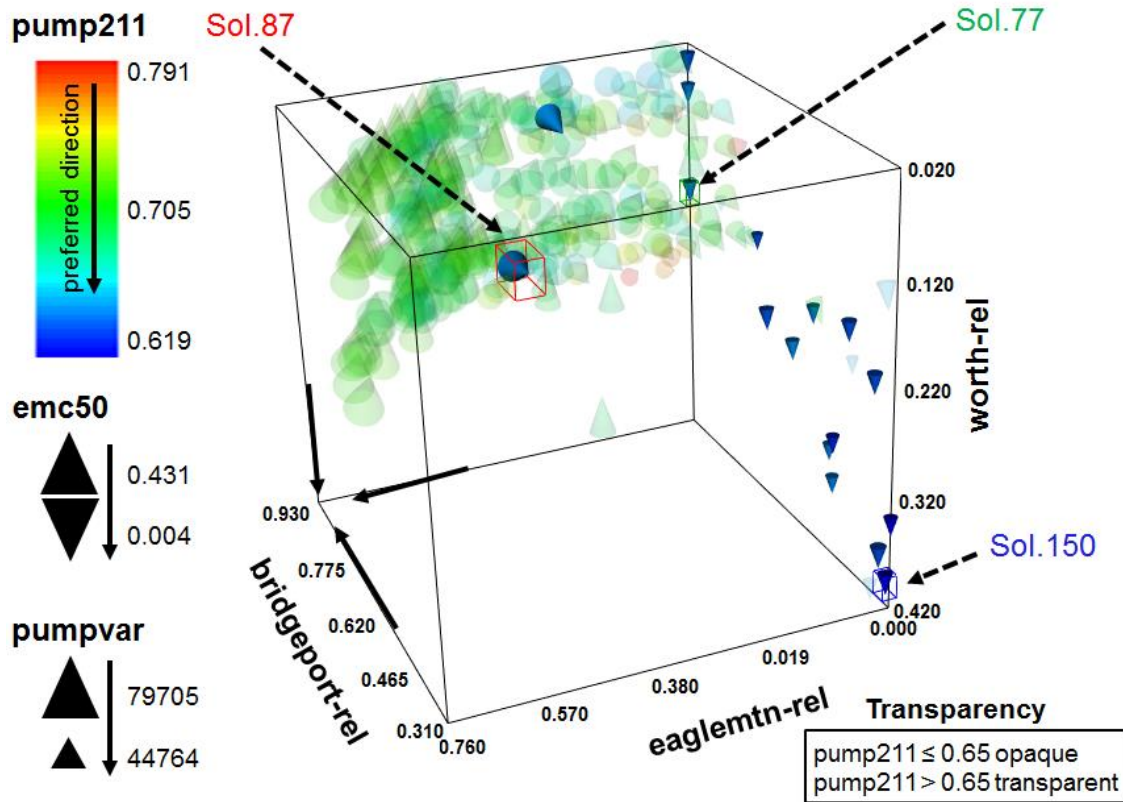


Figure 4.5 Plot of the results of the optimization using the stressed hydrologic ensemble in six dimensions: bridgeport-rel, eaglemtn-rel, and worth-rel on spatial axes, pump211 in color, emc50 conveyed by cone orientation, and pumpvar conveyed by cone size (spill was not relevant for this hydrology). Transparent cones are solutions that violate the pump211 threshold of 65%, while opaque cones exhibit high-volume pumping less than 65% of the year. Three labeled solutions are representative of different reservoir priorities: Sol.150 has the best worth-rel but poor bridgeport-rel and eaglemtn-rel; Sol.87 is the solution that has the highest eaglemtn-rel while meeting the pump211 condition but poor bridgeport-rel and worth-rel performance; Sol.77 has the highest possible bridgeport-rel, poor eaglemtn-rel, and medium worth-rel.

In Figure 4.5, transparency is used to visualize a threshold on the pump211 objective. If TRWD was interested in focusing on results that required high-volume pumping less than 65% of the year, the opaque cones meet that condition. Within the opaque solutions, there are those which have good performance in each of the three reservoir reliabilities: solution 150⁴, marked in blue, achieves good

⁴ Here, the number of the solution refers to the index of the solution in AeroVis's database. The ability to reference specific solutions facilitates visualizing the results in a variety of ways, including importing them into RiverWare to examine state variables. This approach to naming solutions will be used throughout the rest of this document.

worth-rel at the expense of both eaglemtn-rel and bridgeport-rel; solution 87, with the red marker, performs better in eaglmtn-rel but poorly in both bridgeport-rel and worth-rel; solution 77, in green, has the best possible bridgeport-rel performance but mediocre worth-rel and the worst possible eaglemtn-rel performance. The precise values associated with each of these three results' performance are shown in Table 4.1 below.

Another way to directly compare the objective performance of these three marked solutions is with a parallel plot, a technique shown to be useful for viewing data in many dimensions (Inselberg 1985; Wegman 1990). Figure 4.6 shows the performance of solutions 150, 87, 150, and 77 in blue, red, and green, respectively, along with the baseline (dashed black) against the full archive of solutions in grey. As compared to the previous plots, each solution is not a point but rather a line, where the vertical position on each of the columns represents the objective function value. Because of the way the objectives are arranged, the optimal solution would be a line straight across the bottom of the plot.

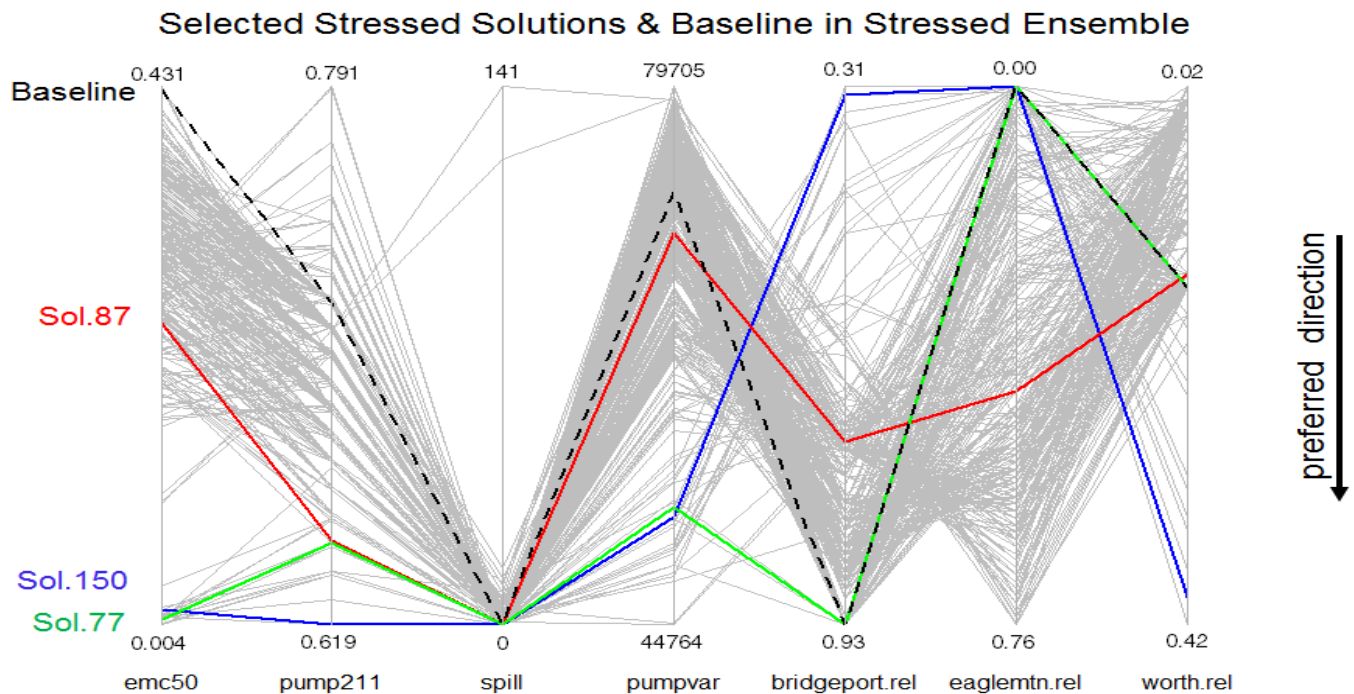


Figure 4.6 Parallel plot showing 7-objective performance of all stressed solutions (grey) with prominent comparison of the baseline, Sol.150, Sol.87, and Sol.77 solutions.

Parallel coordinate plots facilitate the comparison of all objectives simultaneously, to show the relative strengths and weaknesses of these solutions. Solution 150 outperforms the other three in all but two objectives- Bridgeport and Eagle Mountain reliabilities. Solution 87 performs very well in every objective except bridgeport-rel and eaglemtn-rel, notably outperforming the other selected solutions in worth-rel. Solution 77 performs equivalently to the baseline in all of the reservoir reliabilities, but with far better performance in the three other relevant objectives. In other words, solution 77 had better quantitative performance on the pumping objectives (emc50, pump211, pumpvar) suggesting that the optimization approach could improve upon the current baseline management. However, further study would have to be undertaken in order to determine, qualitatively, whether solution 77's management strategies were appropriate for all aspects of the TRWD performance.

More light can be shed on these solutions by examining system components in RiverWare. Figure 4.7 is the Eagle Mountain pool elevation over the course of the year simulation. Blue, red, green, and black lines represent solutions 150, 87, 77, and baseline, as in the parallel plot above. When the storage reliability objectives were added in Formulation 2, an elevation threshold needed to be set in order to transform the timeseries of storage into a single quantitative objective for optimization. However, it is difficult to do this because there are usually multiple targets of interest to managers. For Eagle Mountain, we optimized using the recreation threshold of 644.1 ft, marked by a solid line in the plot. The dashed line at 641.1 ft is the lowest elevation TRWD would tolerate for Eagle Mountain.

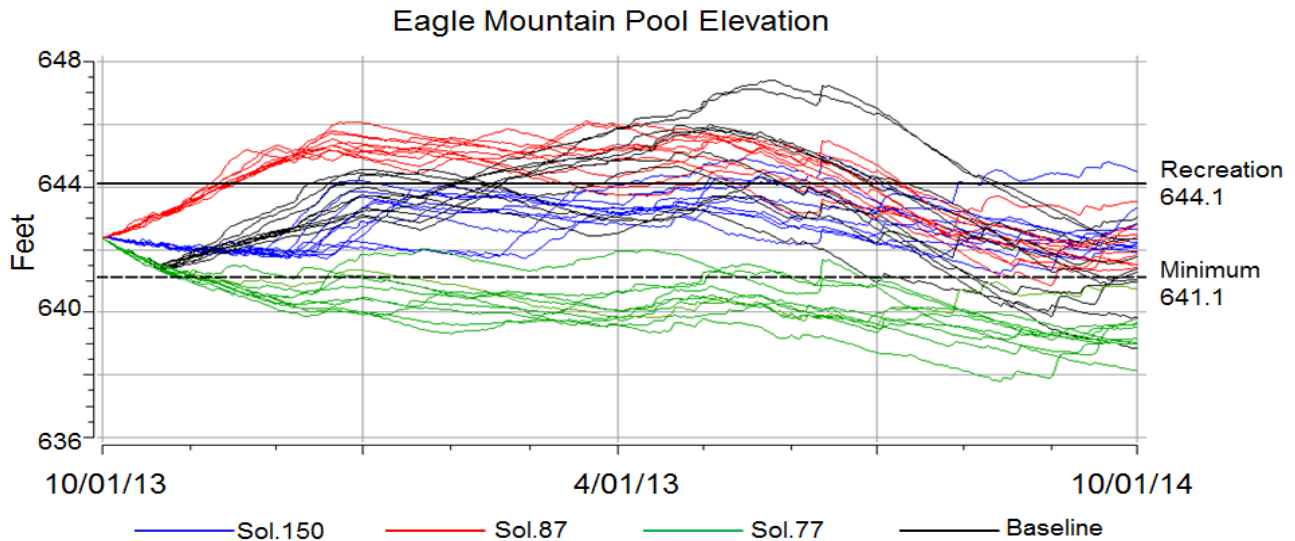


Figure 4.7 Plot of Eagle Mountain reservoir’s pool elevation under baseline, Sol.150, Sol.87, and Sol.77 management for all 10 stressed traces relative to recreation and minimum thresholds.

Solution 77 has Eagle Mountain reliability (with respect to the recreation objective) of 0%, but observing the timeseries, the solution puts Eagle Mountain in an unacceptable state very quickly for almost all of the 10 traces. Examining the other solutions, we find that solutions’ performance in a reliability objective is not a straightforward predictor of its performance in other metrics; solution 150, like 77, has 0% reliability for Eagle Mountain, but never fails to meet the 641.1 ft minimum. In fact, solution 150 outperforms solution 87 (which has an eaglemtn-rel of 43%) which falls below the 641.1 ft threshold in one of the traces, as does the baseline in five traces.

4.4.2 Connecting Policy Variables and Objective Performance

Because each decision variable is in fact a management decision, analyzing the values of policy variables in the context of objective performance can offer insights into how the system responds to the optimized decisions. Table 4.1 presents the 24 decision variable values that comprise each of the three solutions marked in Figure 4.5, along with the Baseline values, and objective values for each. The following paragraphs will go through the logic behind some insights that were obtained from examining how policy variables affect objectives.

Table 4.1 Values for all objectives and policy variables for Sol.150, Sol.87, Sol.77, and baseline.

	Sol.150	Sol.87	Sol.77	Baseline
Objectives				
emc50	0.016	0.243	0.008	0.428
pump211	0.619	0.646	0.645	0.722
spill	0.000	0.000	0.000	0.000
pumpvar	51738.585	70175.718	52401.700	72821.251
bridgeport-rel	0.320	0.720	0.930	0.930
eaglemtn-rel	0.000	0.430	0.000	0.000
worth-rel	0.400	0.160	0.170	0.170
Decision Variables				
emzone1	645.860	644.814	644.179	644.100
emzone2	647.164	645.362	644.425	644.100
emzone3	647.518	647.949	647.044	648.100
bpzone1	813.691	815.747	820.036	821.000
bpzone2	817.985	819.614	827.152	826.000
bpzone3	830.056	832.790	828.873	836.000
emtrigdry1	641.425	641.390	641.155	643.100
emtrigdry2	641.457	641.809	641.195	645.100
emtrigdry3	645.994	644.666	642.058	647.100
emtrigav1	642.076	642.675	641.281	641.100
emtrigav2	642.780	642.965	641.839	643.100
emtrigav3	644.380	643.871	642.674	645.100
emtrigwet1	641.100	641.115	641.201	641.100
emtrigwet2	641.137	641.194	641.258	643.100
emtrigwet3	641.821	642.669	647.266	645.100
emcrate1	48.076	200.000	28.399	150.000
emcrate2	46.898	197.008	2.836	100.000
emcrate3	18.813	194.292	1.202	75.000
worthlev1	590.410	590.236	590.366	590.000
worthlev2	590.783	590.411	590.492	591.000
worthlev3	590.985	590.442	590.761	591.500
worthlev4	591.108	590.516	591.345	592.000
worthlev5	591.256	590.677	592.836	593.000
worthlev6	591.475	590.744	592.836	593.000

Solution 77 has the best performance in bridgeport-rel (it meets the threshold 93% of the time) but unacceptable performance in eaglemtn-rel (it never meets the threshold). In comparison, solution 87 has a much higher Eagle Mountain reliability of 43% and a relatively high Bridgeport reliability of

72%. The values for the three emzones, which tell Bridgeport how much water to release relative to the bpzones, are not vastly different, but the values for the bptime1 and bptime2 are much higher in solution 77. This means that the balance of water between the two reservoirs heavily favors Bridgeport. It seems the three bpzones have a much greater impact on both bridgeport-rel and eaglemtn-rel than the emzones in stressful hydrology.

Something else learned from this detailed comparison of policy variables and objectives concerns the system-wide impact of Eagle Mountain supplementation. Solution 87 has three very high emcrates (the maximum is 200 mgd), which contribute to the high Eagle Mountain reliability compared with the other three solutions. They also factor into the higher frequency of the EMC pumping rate being over 50 mgd compared with solution 150 and 77 (worse performance in emc50). However, the minor increase in frequency of high-volume pumping (increase in pump211) relative to the extreme increase in emcrates suggests that the EMC demand rates do not contribute much to the necessity of high-volume pumping, a useful insight for TRWD.

In Figure 4.6, it was shown that the solutions identified by Borg had significant improvements in the system wide pumping objectives compared to the Baseline solution. The decision variable values in Table 4.1 illuminate why these objective function differences occur. The emtrigdry values from the baseline solution are all around two to five feet higher than any of those in the Borg solutions. The emtrigdry solutions also figure prominently in the stressed optimization results, as discussed in Section 4.5. Higher values for emtrigdry mean that the three different supplementation levels (emcrates) are triggered sooner in the Borg solutions than in the Baseline. Earlier triggering has the effect of maintaining a much higher level of storage in Eagle Mountain during dry seasons. The earlier triggering does not supplement enough to maintain the recreation elevation in the most stressful traces (refer to Figure 4.7) but does, on average of the stressed traces, increase the high-volume pumping, emc50, and pumpvar quite substantially.

Table 4.1 allows us to determine the sensitivity of objective function outputs to changes in the decision variable values. However, RiverWare’s internal rule structure allows users to set up models in which triggers are turned on and off depending on climatic and hydrologic conditions. Therefore, it is critically important to diagnose which climate states occurred in our hydrologic traces, to shed light on which policy variables are most important in times of stress versus surplus. To address this, the following section discusses the model’s climate forecasting procedures as well as the implications of how some policy variables are used by the rules.

4.5 Impact of Climate State Forecast on Policy Variables

The TRWD RiverWare model has climate forecasting functionality through which it tries to predict the climate three months at a time and plan for scarcity or surplus by adjusting some management variables based on the predicted climate state, which can be dry, average, or wet. This variation in management is reflected in the policy variables both explicitly and implicitly, and thus the predicted climate state affects which policy variables are being used in the stressed and surplus ensembles. Recall that the selection of hydrologic ensembles (stressed vs. surplus) relied on the logic that certain traces caused reductions in total water availability in the West Fork reservoirs- stressful traces were more water scarce than either average or surplus traces (see Section 3.6.2 for more in depth discussion). This classification scheme based on longer term system response is different than the model’s internal climate state, which is based on a comparison of each month’s inflows to the dry, average, and wet categories developed from historic data (the same categories used to produce the stochastic traces described in Section 2.4).

The forecasting is carried out quarterly, so predictions are made every October, January, April, and July that set the climate state for the next three months. Two custom rules are responsible for the predictions. The first rule determines the actual quarterly climate state from the previous three months by evaluating the hydrologic inflows to the West Fork and classifying them as dry, average, or wet

compared with historic observations. The second rule predicts the climate state of the next three months based on the actual climate state of the previous three (this eliminates any incorrect predictions propagating indefinitely).

To make the prediction, the rule references a table of transition probabilities (also developed from historic data, see Table 4.2) that assigns a probability of staying in the same climate state or changing to another depending on which month (quarter) it is currently. Note that in this comparison, the current state for comparison is the “actual” state (so the calculation works even if the forecast is incorrect). Due to the particular way the rules are written, there are only two options for how a prediction can turn out: if the probability of remaining in the same state is greater than a 50%, it predicts the same state as the actual state of the previous quarter; if there is a less than 50% chance of remaining in the previous state, the predicted state defaults to “average” for the next three months. Additionally, for a previous state of “average” or wet, the rule references the transition probability from the *previous* quarter to make a prediction (though if the previous state is dry, the current quarter’s transition probability is used).

Table 4.2 Table of quarterly transition probabilities developed from historic data and used by the TRWD model to predict climate states.

	Dry to Dry	Dry to Avg	Dry to Wet	Avg to Dry	Avg to Avg	Avg to Wet	Wet to Dry	Wet to Avg	Wet to Wet
January	54.5	40.9	4.5	21.7	34.8	43.5	21.7	30.4	47.8
April	54.5	31.8	13.6	41.7	25.0	33.3	4.5	45.5	50.0
July	43.5	34.8	21.7	26.1	43.5	30.4	31.8	22.7	45.5
October	54.5	22.7	22.7	21.7	34.8	43.5	22.7	45.5	31.8

Because the model is set up to always start in a dry state and the probability of “Dry to Dry” in October is 54.5%, the first three months will always be predicted as dry. Even if the actual climate state was “wet” for the October through December quarter, the probability of “Wet to Wet” in October (because the rule references the previous quarter’s probabilities) is only 31.8%, so the predicted climate state defaults to average. By inspecting the last column of Table 4.2, it shows that the only way it is possible for the model to predict a wet state is if the actual climate state from April through June was

wet, when the July prediction would reference the April “Wet to Wet” transition probability of 50%. An inspection of the forecasted climate states for all three sets of hydrology bears this out- even for the surplus traces a wet state is only predicted for the last three months of the year simulation. The comparison of how frequently each climate state is predicted for each set of traces is shown in Figure 4.8. The model never predicts a wet state in the stressed traces, and the distribution of climate states in the random traces is a compromise between the stressed and surplus frequencies.

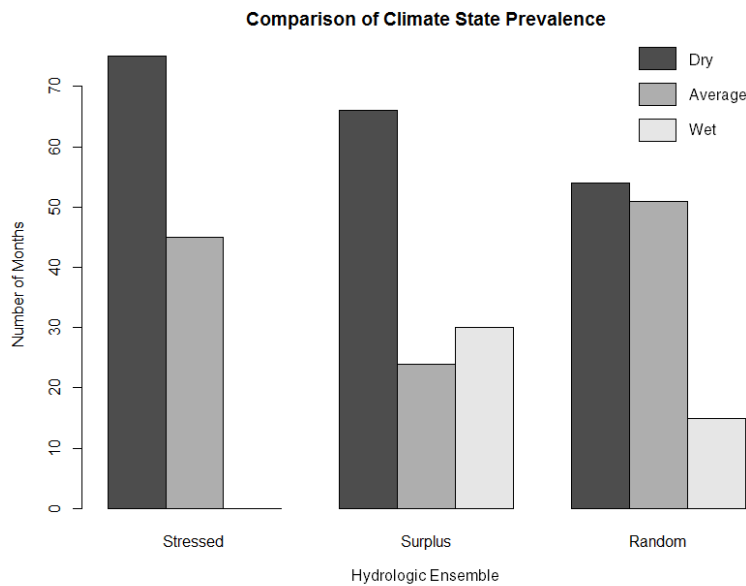


Figure 4.8 Bar plot showing relative prominence of dry, average, and wet climate forecast states in stressed, surplus, and random ensembles.

Within the 24 policy variables, nine are emtrigs that are directly related to the forecasted climate state. If the state is dry, the emtrigdry values are used; if the state is average, the emtrigav values are used; if the state is predicted by be wet, the emtrigwet values are used. If the climate is never forecasted to be wet, as is true for the stressed ensemble, the values of the three emtrigwet policy variables are never actually used by the model and therefore have no relevance to the performance.

Another set of policy variables is implicitly affected by the hydrologic traces used for optimization: the optimization: the six worthlev variables are associated with particular values of West Fork fullness (see

Table 4.3). In the stressed ensemble, the WF%Full never went above 60% under Baseline management, making it unlikely to ever reference the worthwhile values associated with 80-100% West Fork fullness. Similarly, the surplus traces never dipped below 50% in Baseline management, so worthwhile1 is unlikely to impact the optimization of decisions under surplus conditions.

4.6 Hybrid Solutions

The previous sections showed optimization results for both the stressed and surplus scenarios considered separately. The resulting solutions may or may not perform well in a more varied climate, though. Therefore, a final investigation in this thesis is to test the robustness of selected stressed and surplus solutions on an ensemble of 10 traces that were chosen randomly from the set of 100 provided by Hydros. Further refinement of the optimization results may be possible by combining some aspects of the stressed solution set with decisions from the surplus solution set to account for climate variability. To provide one method for doing this, we created three “hybrid” solutions. The goal of the hybrid solutions is to determine whether intelligent combination of decision variable values found to be nondominated in the stressed and surplus conditions can produce performance improvements over their component solutions in the random trace ensemble.

4.6.1 Creation of Hybrids

To create the hybrids, solutions from the stressed and surplus optimizations were chosen and combined. They were chosen based on exhibiting similar performance characteristics within their respective solution spaces. The three stressed solutions are 150, 87, and 77, which have already been analyzed in depth in this chapter, and the corresponding surplus solutions are 172, 224, and 67. Side by side plots of the solution spaces with the marked hybrid contributors are shown in Figure 4.9.

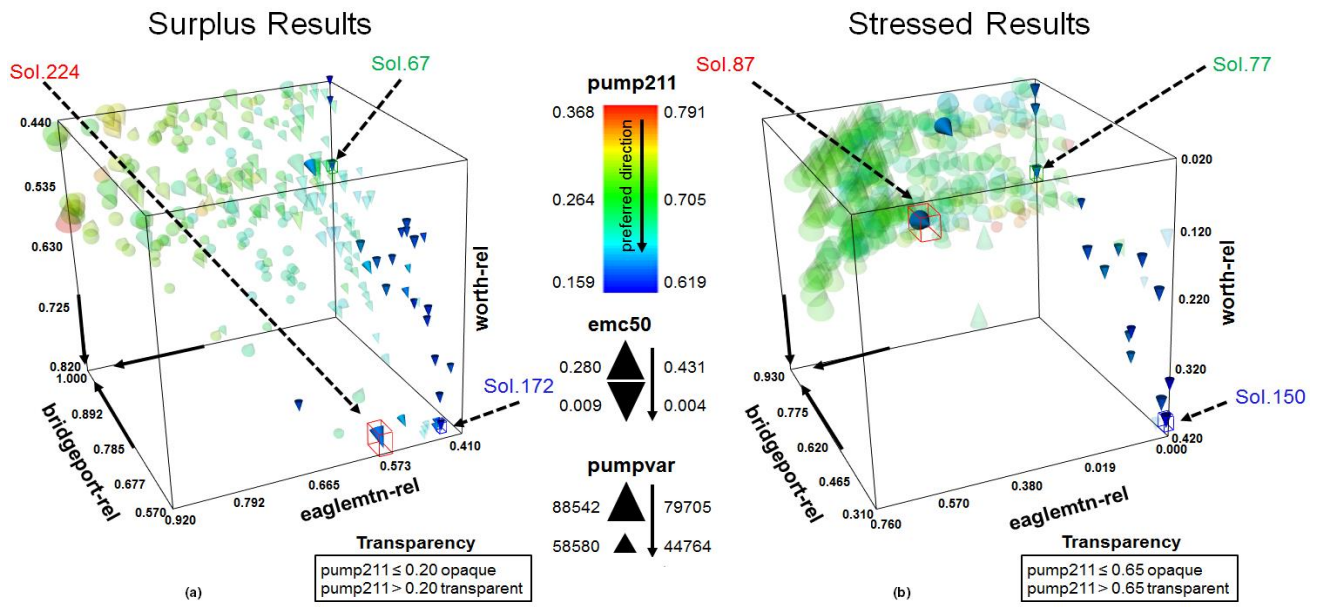


Figure 4.9 (a) plot of surplus optimization results in same format as Figure 4.4 with transparent solutions violating a 20% *pump211* threshold and opaque solutions which do fall below 20% high-volume pumping. Marked and labeled solutions are the surplus solutions chosen as having corresponding performance traits to the three marked stressed solutions which were explained in Figure 4.5; **(b)** the same plot as Figure 4.5, shown for comparison of the three solutions selected from surplus and stressed to become hybrids.

To choose the three surplus solutions, a process similar to the one that produced the three dry solutions was carried out. First, a threshold of 20% was placed on the *pump211* objective to rule out solutions that exhibited frequent high-volume pumping. Second, from the remaining solutions, three that mimicked the reservoir reliability compromises discussed in the stressed results (Section 4.4.1) were chosen. The opaque solutions in plot (a) are the surplus solutions that meet the *pump211* condition. The solutions marked in blue (stressed 150 and surplus 172) will become hybrid 1, both having the best possible *worth-rel* performance and the worst *eaglemtn-rel* performance. Hybrid 2 is comprised of the solutions marked in red (stressed 87 and surplus 224), which both meet the threshold with improved *eaglemtn-rel* performance. Hybrid 3 is marked in green (stressed 77 and surplus 67), with both contributing solutions performing well in *bridgeport-rel* but poorly in *eaglemtn-rel* and *worth-rel*. To illustrate the logic behind creation of hybrid solutions,

Table 4.3 shows the values for every decision variable along with the variable's relevance (impact on system operations) in either stressed, surplus, or both ensembles. The first two sets of variables, emzones and bpzones, dictate the releases from Bridgeport to Eagle Mountain. They are not subject to any alternative functioning as a result of either the model's climate forecasting policies or system response to water scarcity, thus they are relevant in both dry and wet conditions. The second set of variables- emtrigdry, emtrigav, emtrigwet- are columns in a single table that are referenced depending on whether the model forecasts a dry, average, or wet climate state. The worthlev policy variables are the policy variables that determine how much water Eagle Mountain releases to Worth. In the model table, each of the elevations corresponds to a West Fork % Full value which has the effect of implicitly limiting the relevance of policy variables: in very dry hydrology, the West Fork is unlikely to be above 70% full, and when the hydrology is wet, the West Fork will be unlikely to fall below 50% full. Because the average climate state is more prevalent in the stressed ensemble, and the dryer, scarcer hydrology is of greater management concern for TRWD, the values of policy variables that are relevant in both ensembles are taken from the stressed solutions.

Table 4.3 Values for policy variables of a selected stressed solution and a selected surplus solution that will be combined to create Hybrid 1. Hydrologic relevance denotes whether the policy variable is important in stressed hydrology, surplus hydrology, or both (meaning the decision variable is important regardless of hydrology).

Decision Variable		Stressed 150	Surplus 172	Hydrologic Relevance	Hybrid 1
emzone1		645.860	648.025	both	645.860
emzone2		647.164	648.025	both	647.164
emzone3		647.518	648.025	both	647.518
bpzone1		813.691	814.948	both	813.691
bpzone2		817.985	819.479	both	817.985
bpzone3		830.056	833.123	both	830.056
emtrigdry1		641.425	641.100	both	641.425
emtrigdry2		641.457	642.472	both	641.457
emtrigdry3		645.994	643.612	both	645.994
emtrigav1		642.076	641.216	both	642.076
emtrigav2		642.780	643.458	both	642.780
emtrigav3		644.380	645.583	both	644.380
emtrigwet1		641.100	641.176	surplus	641.176
emtrigwet2		641.137	641.194	surplus	641.194
emtrigwet3		641.821	641.313	surplus	641.313
emcrate1		48.076	33.469	both	48.076
emcrate2		46.898	26.121	both	46.898
emcrate3		18.813	4.633	both	18.813
WF % Full					
worthlev1	0	590.410	590.000	stressed	590.410
worthlev2	50	590.783	590.066	both	590.783
worthlev3	70	590.985	590.433	both	590.985
worthlev4	80	591.108	590.504	surplus	590.504
worthlev5	90	591.256	593.693	surplus	593.693
worthlev6	100	591.475	593.693	surplus	593.693

4.6.2 Hybrid Results in Varied Hydrology

A parallel plot facilitates straightforward comparisons to show the relative performance of the hybrids, their component solutions, and the baseline solution in the random ensemble (the 10 stochastic traces chose randomly from the 100 provided by Hydros). Below, Figure 4.10 plots each of the three hybrids as solid lines, where blue is hybrid 1 (stressed 150 and surplus 172), red is hybrid 2 (stressed 87 and surplus 224), and green is hybrid 3 (stressed 77 and surplus 67). The stressed component solutions

are plotted in dots and the surplus component solutions are plotted in dash-dots. The baseline solution is the black dashed line. The grey lines are solutions optimized specifically to the random ensemble. An ideal solution would be a straight line across the bottom.

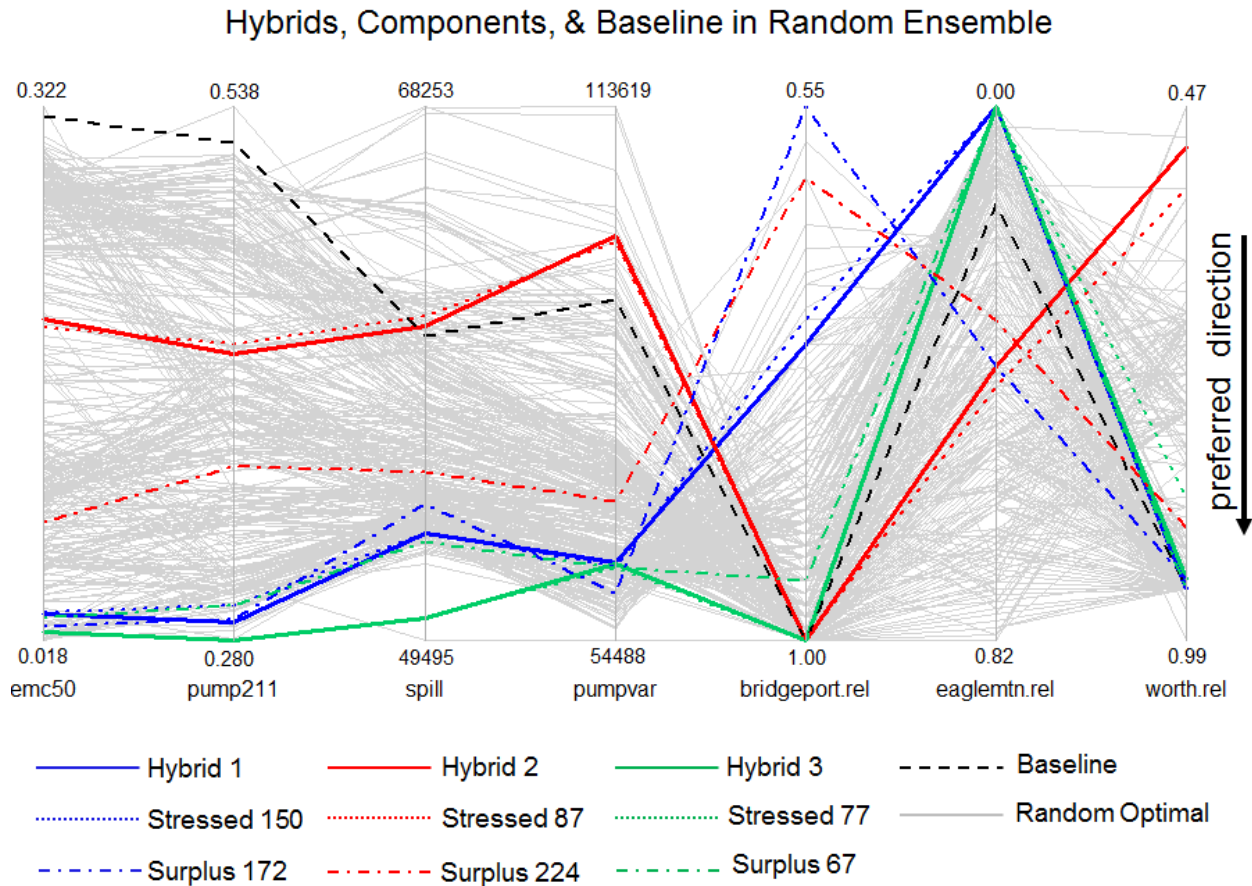


Figure 4.10 The random ensemble performance of three hybrids, their component stressed and surplus solutions, the baseline, and the solution set optimized using the random traces.

One observation from this plot is that the stressed component solutions and the hybrid solutions perform very similarly across all objectives. Since only six of the 24 decision variable values differ between the two solutions, this makes sense. It can be concluded, then, that the six policy variables that were determined to matter mostly for wetter conditions do not have very much impact on the performance of solutions. Referring back to Figure 4.8 that shows the prevalence of climate states in each of the three ensembles, and considering that every simulation begins with a WF%Full

value of 53% (quite far from the 80% WF%Full at which the last three worthlevs are in use), it seems unlikely that the surplus policy variables were used much.

Both Worth and Eagle Mountain reliabilities show mixed results for hybrids: for eaglemtn-rel, hybrids 1 and 3 are outperformed by the surplus solutions, but both hybrids perform equivalently in worth-rel; hybrid 2 has almost equivalent eaglemtn-rel performance as its surplus component and then outperforms both of its components in worth-rel. There is no clear evidence that hand-picked combinations are especially robust in variable conditions. Though hybrid 3 does exhibit either improved or equivalent performance to its components in all objectives, it does not exhibit improvement in eaglemtn-rel, the objective that saw significant draining of the Eagle Mountain reservoir in the stressed ensemble (refer to Figure 4.7), so it is still gaining performance improvement at an unacceptable cost. Furthermore, the hybrids did not significantly outperform the grey solutions optimized to the random ensemble.

In Bridgeport reliability, however, the hybrids uniformly outperform the surplus components. For the surplus-specific optimization, the increased hydrologic inflows to Bridgeport and Eagle Mountain resulted in relatively high emzones (that determine releases from Bridgeport and in turn necessitate less East Texas supplementation) which end up draining Bridgeport in a more varied climate (refer to Table 4.1). While it was the bpzones that had the most impact on Bridgeport and Eagle Mountain reliability in the stressed ensemble, in a more varied climate, emzones seem to be larger factors. This insight makes a case for developing climate-specific operations to balance Bridgeport and Eagle Mountain, a task that TRWD is considering (Blaylock 2014c).

Chapter 5 Concluding Remarks

This chapter highlights some main prominent points of discussion, including ways to expand the scope of this study and address open questions in the future, and then ends with a conclusion to summarize the undertaking described in the previous chapters.

5.1 Discussion and Future Work

Water basin planning in the United States has traditionally been done using cost-benefit analysis, in which there is a single utility function for the whole system or basin that quantifies the basin's benefits (Harou et al. 2009). In contrast, this thesis research uses a many objective approach that evaluates performance based on multiple, separate objectives. Initially, the four objectives used to optimize were defined through system-wide performance goals- reduce pumping, reduce pumping variability, and monitor spill so that pumping is not done wastefully. The initial problem formulation was constructed in such a manner to reflect TRWD's goal to improve the efficiency of their entire system. However, the initial results (Section 4.1), demonstrated that the initial formulation only yielded a small number of solutions that were qualitatively very similar. In other words, there was essentially only one management solution in light of the chosen objectives. While a lack of diverse alternatives does not inherently mean optimization failure, it does suggest that the objectives do not conflict strongly. Furthermore, by performing an in-depth investigation of the three initial stressed results in RiverWare, we discovered that the solutions penalized individual reservoirs in ways that would have been difficult to predict a priori. Similar to a theoretical discussion of the pitfalls of aggregated objective functions expressed in (Franssen 2005), optimizing only on system-wide objectives did not consider the distinct system components that could suffer if not included in performance evaluation.

The results from the second problem formulation, which included reservoir reliability objectives, expanded the set of management alternatives and revealed performance conflicts not only between individual reservoirs, but between the individual reservoirs and the system-wide goals as well. The new

problem formulation greatly expanded the number of objectives considered, from four to seven. Our successful use of such a high dimensional formulation was enabled by advances in search technology, such as the auto-adaptive methodology of the Borg MOEA. Recent algorithm comparisons have shown that the Borg algorithm performs well for up to 10 objectives (Reed et al. 2013). Therefore, the results demonstrate an opportunity for water managers to expand their problem formulations in order to reflect the realization that priorities of individual components (reservoirs, water users, etc.) are not necessarily accounted for when trying to improve the performance of a system as a whole. Innovative many objective formulations, such as used in this thesis, can help shed light on such issues.

This study's iterative process of defining objectives is an example of how MOEA-assisted multiobjective optimization enables constructive decision aiding (Hitch 1960; Liebman 1976; Zeleny 1989). Constructive decision aiding is a concept in which problem formulations are discovered as part of the decision support process itself. Such an approach may mean that "intermediate" problem formulations (such as Problem Formulation 1 in this research) are presented during the process. While the results of intermediate problem formulations may be infeasible to implement as management strategies, they uncover implicit objectives that managers may not have been aware of, or important aspects of problem definition that had not been considered, and offer opportunity for revision. The examination of the initial results in objective space and then in model space through RiverWare plots revealed the importance of including reservoir levels in performance evaluation. While TRWD has inherently accounted for the reservoir elevations in its decisions, the need to explicitly measure their performance was information gained through the specific context provided by optimization results.

Using the results of the first optimization to suggest additional objectives for the second problem formulation speaks to the iterative and complex nature of human decision making (Zeleny 1989). Similar to many other technology-assisted decision support techniques, MOEA-based decision support is only as good as the problem formulation that is presented to the algorithm. Because the

algorithm needs to be told explicitly what to measure, while humans consciously or subconsciously consider many goals at once, multiple iterations to discern managers' goals may be necessary before the algorithm is truly able to produce solutions that accurately reflect the system's potential for operational improvement. This was echoed in a recent editorial about a real-world MOEA decision support exercise, in which the authors discussed that there is no substitute for human judgment within this process (Basdekas 2014). After adding three reservoir reliability objectives, comparison of the optimization results with the baseline showed, for example, that TRWD's current operations prioritize the elevation of Bridgeport at the expense of Worth and Eagle Mountain. Using an MOEA to quantify this tradeoff gives TRWD information about the ramifications of their current decisions and allows them to make an educated determination about their priorities in the future.

Another realization aided by examination of the second round of optimization results in objective and model space was that measures of reservoir reliability were not sufficient to characterize the reservoirs' performances. One strategy to aid this issue is the use of reliability, resiliency, and vulnerability measures (Hashimoto, Stedinger, and Loucks 1982). Reliability measures how often a quantity falls above or below a threshold, so choice of the threshold is critically important. Our results demonstrate that solutions that had equivalent reservoir reliability could have very different degrees or timing of failures, and vastly different storage trajectories. Accounting for resilience, or the time required to recover after a failure, would be an important inclusion in a third problem formulation. Alternatively, vulnerability objectives that measure the severity of the failure could also contribute to a new formulation. Another possible way to incorporate thresholds deemed inviolable by TRWD would be to add elevation constraints to a future problem formulation.

In this thesis, the issues mentioned above were explored in the context of the TRWD planning problem, a water resources case study that has never before been optimized in the literature. It is a large, multi-reservoir network with supply and demand challenges that are compounded by complex

electricity purchasing considerations. Balancing reservoirs, in general, and especially in light of the costs of pumped supplementation, is difficult; the best performance for one reservoir means depletion of other reservoirs and/or increased pumping. Outsourcing the production of management alternatives to an algorithm is useful for discovering creative solutions (Kasprzyk et al. 2009; Zeff et al. 2014) but the results are contingent on modeling assumptions that, in their attempts to address uncertainty, demand further assumptions by users and analysts.

The complexity of the operating policy in the model translates into a relatively long simulation time. While this may not matter for the usual way TRWD uses the model, it limited the opportunity for this research to address hydrologic variability and explore the effects of initial conditions. The results showed that hydrology has a significant impact on objective performance, and the implications of varying degrees of stress or surplus on optimization results could not be adequately characterized in the time allotted for study. Past research suggests that evaluating management alternatives in scenarios that stress the system is a valuable strategy (Cui and Kuczera 2009; Kasprzyk et al. 2012), and further study using this model could both support those claims and provide information to TRWD about what conditions cause vulnerability for their system.

Another concern for the robustness of the results is the initial conditions of the reservoirs- every simulation began with the West Fork at 53% full, which is a relatively low storage level to start from given the varying starting percentages of the historical record (see Figure 3.1). The depleted state of the Bridgeport and Eagle Mountain reservoirs at the beginning of the simulation would cause them to be more likely to fall below the thresholds set in the objectives and necessitate more East Texas pumping. It's possible that this is not a major factor in the performance of the pumping-related objectives (emc50, pump211, pumpvar). If the system has been drawn down to this degree by the previous year, it is likely that just maintaining the reservoirs up to 53% required significant pumped supplementation. In light of the electricity purchasing scheme, TRWD would choose to maintain a high

level of pumping. Essentially, the initial conditions used by this study correspond to the conditions that would control the objective function performance for our simulation year no matter what hydrology was experienced within the simulated year. Regardless, it can reasonably be stated that TRWD operates from both more and slightly less advantageous starting points, so it would be interesting to determine in future work whether operating policy should change depending on the state of the West Fork reservoirs at the beginning of the water year. Additional information about how adjusting the policy variables can improve the outlook for the longer term given stressful initial conditions could be gained by performing the optimization on a multi-year simulation.

A challenge presented by complicated operating policies, like those in the TRWD model,, is that it creates some ambiguity about how policy variables are affecting the system. For example, the climate forecasting capability adds a degree of uncertainty to the already uncertain hydrology; it may incorrectly predict the climate state and create a disconnect between what *would have been* a sound operational decision (value for decision variable) and what the model rules dictate. This intervention suggests that in future work, optimizing using certain subsets of hydrology but no forecasting would be a promising way of populating the policy variables that are subject to change when forecasting is employed.

There are several advantages that come with the use of this complex model, though. Labadie (2004) highlighted that evolutionary algorithms facilitate the ability to use trusted models within the search process. Our decision support framework enables TRWD to examine its solutions using detailed visualizations within RiverWare. For example, plots of reservoir elevation, views of multiple water accounts, time series of spill, and a suite of many output objectives (including objectives not including in search) facilitate iterative problem formulation. TRWD can evaluate solutions on its own terms and easily articulate deficiencies. Additionally, the model represents a high degree of infrastructure and policy complexity - no gross simplifications create doubt about how the system would really respond to any given alternative. The trust that TRWD has in the model to aid its short- and long-term planning

translates to trust in the optimization results. Finally, and perhaps most applicable to the outcomes in this study, because this model is a faithful representation of the actual system, all information gained from every step of the optimization process is useful to TRWD and will aid the utility in future decisions.

5.2 Conclusions

This research builds upon past studies using MOEAs to solve complex water management problems by embedding a sophisticated RiverWare model in the algorithm search loop and incorporating decision maker collaboration to assist in iterative problem formulation. The challenges associated with using the RiverWare model included linking the model to the algorithm and the model's long simulation time. Addressing the model simulation time necessitated several creative approaches to ensuring efficient but sufficient search as well as intelligent incorporation of hydrologic variability. In addressing these issues, this research confirms that advanced, GUI-based water management models can be used in MOEA-driven multiobjective optimization.

The complexity of the TRWD system, and therefore the model, offered additional challenges through which this study was able to gain further insight into the MOEA-assisted multiobjective optimization methodology. The initial set of objectives centered on system-wide reduction in pumping based on concerns expressed by TRWD. Through the failure of this initial set of objectives to account for the performance of individual reservoirs, this research confirmed that conflicts exist between objectives both at a sub-system scale as well as between system components and the broader system-wide objectives. Furthermore, the incorporation of this new information into a second problem formulation, which provided further system insights, exemplified the utility of iterative problem definition as crucial to the decision making process.

The results obtained with this complex model suggest the need for further refinement of problem formulation, but also provide valuable information to TRWD. By comparing the performance of the baseline to that of the optimization results, this research provides context for the current

operations. For example, from Figure 4.6 we see that in stressed hydrologic conditions, TRWD currently prioritizes the elevation of Bridgeport over the elevations of both Eagle Mountain and Worth, and also over the pumping-related objectives. This may or may not be intentional or desirable, but with this evidence the utility can determine whether this reflects the true management priorities. As a result of this Bridgeport preference, there are several instances of Eagle Mountain's elevation falling below the minimum under baseline management (per Figure 4.7), which TRWD has expressed concern about. By engaging with these results, TRWD is able to evaluate exactly what their operating priorities are and what performance metrics are inviolable at any cost.

Finally, the implications of climate forecasting and initial conditions within their model could have a significant impact on the performance of suggested management alternatives, and may contribute to some ambiguity in the relationships between the decisions made to balance and supplement reservoirs and the performance outcomes. This knowledge may inform TRWD's approach to optimization and decision making in the future, as well as prove the value of the intermediate outcomes in multiobjective optimization.

References

- AMEC. 2010. Tarrant Regional Water District Forecasting and Operational Efficiency Study.
- Basdekas, L. 2014. "Is Multiobjective Optimization Ready for Water Resources Practitioners? Utility's Drought Policy Investigation." *Journal of Water Resources Planning and Management* 140 (3): 275–76. doi:10.1061/(ASCE)WR.1943-5452.0000415.
- Blaylock, Laura. 2013.
- . 2014a.
- . 2014b.
- . 2014c.
- Brill, E. Downey, Jr., John M. Flach, Lewis D. Hopkins, and S. Ranjithan. 1990. "MGA: A Decision Support System for Complex, Incompletely Defined Problems." *IEEE Transactions on Systems, Man, Cybernetics* 20 (4): 745–57.
- Coello, Carlos A. Coello, Gary B. Lamont, and David A. Van Veldhuizen, eds. 2007. *Evolutionary Algorithms for Solving Multi-Objective Problems*. 2nd ed. Genetic and Evolutionary Computation. New York: Springer.
- Cui, L, and G Kuczera. 2009. "Assessment of the Replicate Compression Heuristic to Improve Efficiency of Urban Water Supply Headworks Optimization." *Journal of Water Resources Planning and Management* 135 (6): 451–57.
- Deb, K. 2001. *Multi-Objective Optimization Using Evolutionary Algorithms*. New York: John Wiley.
- Deloitte & Touche, LLP. 2013. Tarrant Regional Water District Annual Financial Report As of and For the Year Ended September 30, 2012. Fort Worth.
- Franssen, Maarten. 2005. "Arrow's Theorem, Multi-Criteria Decision Problems and Multi-Attribute Preferences in Engineering Design." *Research in Engineering Design* 16: 42–56.

- Giuliani, M., J. D. Herman, A. Castelletti, and P. Reed. 2014. "Many-Objective Reservoir Policy Identification and Refinement to Reduce Policy Inertia and Myopia in Water Management." *Water Resources Research*. doi:10.1002/2013WR014700.
- Gleick, P. H. 2003. "Global Freshwater Resources: Soft-Path Solutions for the 21st Century." *Science* 302 (5650): 1524–28. doi:10.1126/science.1089967.
- Gober, P. 2013. "Getting Outside the Water Box: The Need for New Approaches to Water Planning and Policy." *Water Resources Management* 27 (4): 955–57. doi:10.1007/s11269-012-0222-y.
- Hadka, David, and Patrick Reed. 2013. "Borg: An Auto-Adaptive Many-Objective Evolutionary Computing Framework." *Evolutionary Computation* 21 (2): 231–59. doi:doi:10.1162/EVCO_a_00075.
- Hadka, David, Patrick M. Reed, and Timothy W. Simpson. 2012. "Diagnostic Assessment of the Borg MOEA for Many-Objective Product Family Design Problems." In *Evolutionary Computation (CEC), 2012 IEEE Congress on*. Brisbane. doi:10.1109/CEC.2012.6256466.
- Harou, Julien J., Manuel Pulido-Velazquez, David E. Rosenberg, Josue Medellin-Azuara, Jay R. Lund, and Richard E. Howitt. 2009. "Hydro-Economic Models: Concepts, Design, Applications, and Future Prospects." *Journal of Hydrology* 375: 627–43.
- Hashimoto, Tsuyoshi, Jery R. Stedinger, and Daniel P. Loucks. 1982. "Reliability, Resiliency and Vulnerability Criteria for Water Resource System Performance Evaluation." *Water Resources Research* 18 (1): 14–20.
- Hitch, Charles J. 1960. *On the Choice of Objectives in Systems Studies*. P-1955. The RAND Corporation.
- Inselberg, Alfred. 1985. "The Plane with Parallel Coordinates." *The Visual Computer* 1: 69–91.
- Jacobs, J. 2011. "The Sustainability of Water Resources in the Colorado River Basin." *The Bridge: Linking Engineering and Society* 41 (4): 6–12.

- Kasprzyk, Joseph Robert, Shanthi Nataraj, Patrick Michael Reed, and Robert J. Lempert. 2013. "Many Objective Robust Decision Making for Complex Environmental Systems Undergoing Change." *Environmental Modelling and Software* 42: 55–71. doi:10.1016/j.envsoft.2012.12.007.
- Kasprzyk, Joseph Robert, Patrick Michael Reed, Gregory W. Characklis, and Brian R. Kirsch. 2012. "Many-Objective de Novo Water Supply Portfolio Planning under Deep Uncertainty." *Environmental Modelling and Software* 34: 87–104. doi:10.1016/j.envsoft.2011.04.003.
- Kasprzyk, Joseph Robert, Patrick Michael Reed, Brian R. Kirsch, and Gregory W. Characklis. 2009. "Managing Population and Drought Risks Using Many-Objective Water Portfolio Planning under Uncertainty." *Water Resources Research* 45. doi:10.1029/2009WR008121.
- Kollat, J. B., and P. M. Reed. 2006. "Comparing State-of-the-Art Evolutionary Multi-Objective Algorithms for Long-Term Groundwater Monitoring Design." *Advances in Water Resources* 29 (6): 792–807.
- . 2007. "A Framework for Visually Interactive Decision-Making and Design Using Evolutionary Multiobjective Optimization (VIDEO)." *Environmental Modelling & Software* 22 (12): 1691–1704.
- Labadie, J. W., and M. L. Baldo. 2000. "MODSIM: Decision Support System for River Basin Management: Documentation and User Manual". Dept. of Civil Engineering, Colorado State University.
- Labadie, John W. 2004. "Optimal Operation of Multireservoir Systems: State-of-the-Art Review." *Journal of Water Resources Planning and Management* 130 (2): 93–111.
- Laumanns, M. 2002. "Combining Convergence and Diversity in Evolutionary Multiobjective Optimization." *Evolutionary Computation* 10 (3): 263–82.
- Liebman, Jon C. 1976. "Some Simple-Minded Observations on the Role of Optimization in Public Systems Decision-Making." *Interfaces* 6 (4): 102–8.
- Lund, J. 2012. "Expanding Water Storage Capacity in California." *California Water Blog*.
<http://californiawaterblog.com/2012/02/22/expanding-water-storage-capacity-in-california/>.

- Matrosov, Evgenii S., Julien J. Harou, and Daniel P. Loucks. 2011. "A Computationally Efficient Open-Source Water Resource System Simulator – Application to London and the Thames Basin." *Environmental Modelling and Software* 26 (12): 1599–1610. doi:10.1015/j.envsoft.2011.07.013.
- Milly, P. C. D., J. Betancourt, M. Falkenmark, R. Hirsch, Z. W. Kundzewicz, D. P. Lettenmaier, and R. J. Stouffer. 2008. "Stationarity Is Dead: Whither Water Management?" *Science* 319: 573–74.
- Mortazavi, Mohammad, George Kuczera, and Lijie Cui. 2012. "Multiobjective Optimization of Urban Water Resources: Moving toward More Practical Solutions." *Water Resources Research* 48. doi:10.1029/2011WR010866.
- Nicklow, John, Patrick Reed, Dragan Savic, Tibebe Dessalegne, Laura Harrell, Amy Chan-Hilton, Mohammad Karamouz, et al. 2010. "State of the Art for Genetic Algorithms and Beyond in Water Resources Planning and Management." *Journal of Water Resources Planning and Management* 136 (4): 412–32.
- Piscopo, A. N., J. R. Kasprzyk, and R. M. Neupauer. 2014. "An Iterative Approach to Multi-Objective Engineering Design: Optimization of Engineered Injection and Extraction for Enhanced Groundwater Remediation (In-Review)." *Environmental Modelling & Software*.
- Reed, P. M., D. Hadka, J. D. Herman, J. R. Kasprzyk, and J. B. Kollat. 2013. "Evolutionary Multiobjective Optimization in Water Resources: The Past, Present and Future." *Advances in Water Resources* 51 (January): 438–56.
- Smith, Rebecca, Joseph Kasprzyk, Edith Zagona, Laura Blaylock, John Carron, and Nick Mander. 2014. "Many Objective Analysis to Optimize Pumping & Releases in a Multi-Reservoir Water Supply Network." In Portland, OR.
- Smith, Rebecca, Joseph Kasprzyk, Edith Zagona, John Carron, and Nick Mander. 2013. "Combining Interactive Infrastructure Modeling and Evolutionary Algorithm Optimization for Sustainable Water Resources Design." In Fall Meeting 2013. San Francisco.

- Tarrant Regional Water District. 2013a. "Integrated Pipeline: Regional Partnership for Water Supply."
iplproject.com.
- . 2013b. "TRWD." www.trwd.com.
- Texas Water Development Board. 2014a. "TWDB." www.twdb.texas.gov.
- . 2014b. "Water Data For Texas." www.waterdatafortexas.org.
- TRC/Brandes. 2009. Assessment of TRWD System Operating Policies for Different Climatic Conditions
Using the RiverWare Models. Austin.
- U.S. Census Bureau. 2012. Growth in Urban Population Outpaces Rest of Nation, Census Bureau Reports.
- Watkins Jr., David W., and Daene C. McKinney. 1997. "Finding Robust Solutions to Water Resources
Problems." *Journal of Water Resources Planning and Management* 123 (1): 49–58.
- Wegman, Edward J. 1990. "Hyperdimensional Data Analysis Using Parallel Coordinates." *Journal of the
American Statistical Association* 85 (41): 664–75.
- Western States Water Council. 2012. 47th Annual Report. Murray.
- Woodruff, Matthew J., Patrick M. Reed, and Timothy Simpson. 2013. "Many Objective Visual Analytics:
Rethinking the Design of Complex Engineered Systems." *Structural and Multidisciplinary
Optimization* 48: 201–19.
- World Water Assessment Programme. 2009. The United Nations World Water Development Report 3:
Water in a Changing World. Paris: UNESCO.
- Yates, D, J Sieber, D Purkey, and A Huber-Lee. 2005. "WEAP21 - A Demand-, Priority-, and Preference-
Driven Water Planning Model Part 1: Model Characteristics." *Water International* 30 (4): 487–
500.
- Zagona, E, T. J. Fulp, R Shane, Y Magee, and H. M. Goranflo. 2001. "RiverWare: A Generalized Tool for
Complex Reservoir System Modeling." *Journal of the American Water Resources Association* 37
(4): 913–29.

Zeff, Harrison B., Joseph R. Kasprzyk, Jonathan D. Herman, Patrick M. Reed, and Gregory W. Characklis.

2014. "Navigating Financial and Supply Reliability Tradeoffs in Regional Drought Management Portfolios." *Water Resources Research*, June, n/a–n/a. doi:10.1002/2013WR015126.

Zeleny, Milan. 1989. "Cognitive Equilibrium: A New Paradigm of Decision Making?" *Human Systems Management* 8: 185–88.