The Pennsylvania State University
The Graduate School
College of Engineering

OPTIMIZATION OF HYDROLOGICALLY LINKED HYDROPOWER SYSTEMS WITH MULTIPLE OWNERS THROUGH WATER PAYMENTS

A Dissertation in
Mechanical Engineering
by
David B. Beevers

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

December 2015

The dissertation of David B. Beevers was reviewed and approved* by the following:

Horacio Perez-Blanco
Professor of Mechanical Engineering
Dissertation Adviser
Chair of Committee

Hosam Fathy
Associate Professor of Mechanical Engineering

Anil Kulkarni
Professor of Mechanical Engineering

Uday Shanbhag
Associate Professor of Industrial and Manufacturing Engineering

Seth Blumsack
Associate Professor of Energy and Mineral Engineering

Karen Thole
Professor of Mechanical Engineering
Department Head of Mechanical and Nuclear Engineering

*Signatures on file in the Graduate School

ABSTRACT

Before the advent of deregulated electricity markets in the 1990's, hydroelectric power production schedules were typically coordinated with the schedules of other generating technologies by utility companies to meet expected regional demand. However, the growing adoption of deregulated markets requires the development of new methods for determining optimal utilization of hydroelectric resources. Much work has been performed regarding the optimization of cascaded hydropower systems owned and operated by single entities. However, these works fail to address the complexity involved in optimizing a cascaded system with owners along the same river. In those instances, each plant is limited to operate with whatever water flow the upstream plants choose to make available. In this situation, system optimal operation results in higher net system revenue than individual optimization of each plant with whatever water flow is available. However, such an optimal system solution must also result in lower net revenue for some owners compared to the value they would obtain through individual optimization. Hence, a novel methodology is proposed in which individual optimization naturally leads to system optimization through the use of iteratively determined water payments to direct the release of water toward a more optimal utilization. The proposed method alternates between dynamic programming for individual optimization and gradient analysis to determine and update hourly, location-specific water values for modified release. The sequence is repeated until there are no further additions or modifications to the payment schedule that results in increased total revenue. The water values are based upon local and downstream revenue gradients with respect to the flow rate released by the upstream facility during each time period. The method is applied in this thesis to a hydropower system consisting of three connected hydropower facilities, successfully increasing the total gross revenue of each installation. Although the algorithm it is not yet able to achieve the globally optimal solution, it comes quite close and offers unique insights into the value of water.

TABLE OF CONTENTS

List of Figures

List of Tables

Nomenclature

Sets	
i, I_j	Turbines is hydropower plant j.
j, J	The hydropower plants in the system.
J_{up}	The set of hydropower plants that are upstream of other plants.
k, K_j	A subset of J consisting of the plant(s) immediately upstream from plant j.
t, T	The time periods under consideration (in the planning horizon).
$\psi_j, \Psi_{\mathrm{j}}$	The unique combinations (operating sets) of the turbines in plant j.
Constants and Known Inputs	
g	Gravity constant.
k	Start-up cost ratio $\left(\frac{\$}{MW}\right)$
$P_{nom,i,j}$	Nominal (rated) power output of turbine i in plant j. (MW)
Q_j^0	Flow through plant j during the time period immediately before the optimization begins. $\left(\frac{m^3}{s}\right)$
$Q_{in,j,t}$	Uncontrolled flow into plant j's reservoir during time period t. $\left(\frac{m^3}{s}\right)$
$Q_{min,i,j}$ $Q_{max,i,j}$	Minimum and maximum flow through turbine i in plant j. $\left(\frac{m^3}{s}\right)$
$Q_{min,j}$ $Q_{max,j}$	Minimum and maximum net flow through plant j. $\left(\frac{m^3}{s}\right)$
$Q_{min,\psi,j}$ $Q_{max,\psi,j}$	Minimum and maximum net flow through the turbines in operating set ψ in plant j. $\left(\frac{m^3}{s}\right)$
ΔQ_{max}	Maximum allowed change in flow between time periods. $\left(\frac{m^3}{s}\right)$
V_j^1	The volume of water in plant j's reservoir before the initial time period. $(m^3)(m^3)$
V_j^{T+1}	The volume of water in plant j's reservoir after the final time period. (m^3)
$V_{min,j}$ $V_{max,j}$	Minimum and maximum volume limits for plant j's reservoir. (m^3)
$\pi_{j,t}$	Power price at plant j during time period t. $\left(\frac{\$}{MW}\right)$
ρ_{H_2O}	The density of water. $\left(\frac{kg}{m^3}\right)$

$\rho_{j,t}$	Regulation power price at plant j during time period t. $\left(\frac{\$}{MW}\right)$
$\phi_{i,j}^0$	The status of turbine i in plant j before the optimization begins.
ψ_j^0	The operating set active in plant j before the optimization begins.
Independent Variables	
$Q_{i,j,t}$	Flow through turbine i in plant j during time period t. $\left(\frac{m^3}{s}\right)$
$Q_{j,t}$	Net flow through plant j during time period t. (Independent after section 2.1.2) $\left(\frac{m^3}{s}\right)$
Q	Used in figures to refer to the net flow rate through a facility. $\left(\frac{m^3}{s}\right)$
$\phi_{i,j,t}$	On/Off status of turbine i in plant j during time period t.
$\psi_{j,t}$	Active set of turbines in plant j during time period t.
Dependent Variables	
a, b, c, d	Coefficients used in interpolation function.
$H_{j,t}$	The operating net head of plant j during time period t. (m)
$P_{i,j,t}$	Power produced by turbine i in plant j during time period t. (MW)
$P_{j,t}$	Net power produced by plant j during time period t. (MW)
$P_{reg,j,t}$	Regulation power capability of plant j during time period t. (MW)
$P_{\psi,j,t}$	Power produced by turbine set ψ in plant j during time period t. (MW)
$Q_{j,t}$	Net flow through turbines in plant j during time period t. $\left(\frac{m^3}{s}\right)$
ΔQ_{grid}	Flow rate discretization step size. $\left(\frac{m^3}{s}\right)$
R	Total revenue produced by the system. ($)
$SU_{i,j,t}$	Start-up costs incurred by turbine i in plant j at the beginning of time period t. ($)
$SU_{j,t}$	Total start-up costs incurred by station j during time period t. ($)
$\boldsymbol{u_{j,t}}$	Vector of input variables used in dynamic programming.
$V_{j,t}$	Volume of water in plant j at the beginning of time period t. (m^3)
ΔV_{grid}	Volume discretization step size. (m^3)
VtG	Value-to-go, used in dynamic programming optimization. ($)

$\boldsymbol{x_{j,t}}$	Vector of state variables used in dynamic programming.
$\eta_{i,j,t}$	The mechanical efficiency of turbine i in plant j during time period t.
μ, λ	Lagrange multipliers in the KKT analysis.
Optimization Results	
$\mathrm{PS}_j^{\mathrm{n}}$	The data vector necessary for calculating payments made to an upstream facility j, based on the results of the gradient analysis in the n^{th} iteration.
$Q_{opt,j,t}^{n}$	The optimal net flow rate of plant j during time period t from the n^{th} iteration. $\left(\frac{m^3}{s}\right)$
$Q_{opt,up,t}^{n,s+}$	The optimal net flow rate of an upstream plant during time period t due to requiring that its net flow be incremented for time period s during the n^{th} iteration. $\left(\frac{m^3}{s}\right)$
R_{flow}	The payment an upstream facility receives for modifying its flow based on the results of the gradient analysis. ($)
$R_{opt,j}^{n}$	The optimal revenue of plant j from the n^{th} iteration. ($)
$\Delta R_{opt}^{1,s+}$	The effect on the total revenue of a pair of neighboring facilities due to requiring that the net flow of the upstream plant be incremented for time period s during the n^{th} iteration. ($)
$\Delta R_{opt,down}^{n,s+}$	The effect on the optimal revenue of a downstream plant due to requiring that the net flow of the upstream plant be incremented for time period s during the n^{th} iteration. ($)
$\Delta R_{opt,up}^{n,s+}$	The effect on the optimal revenue of an upstream plant of a pair due to requiring its net flow to be incremented for time period s during the n^{th} iteration. ($)
$V_{opt,j,t}^{n}$	The optimal volume of plant j during time period t from the n^{th} iteration. (m^3)
$\Psi_{\mathrm{opt,j,t}}^{n}$	The optimal operating set of plant j during time period t from the n^{th} iteration.

1 Introduction

1.1 The History of Hydropower

Hydropower has been a prominent source of power for thousands of years. In ancient times, the energy of falling and flowing water was harnessed to power grain mills and sawmills. This application remained the primary use of hydropower for over a thousand years until people began to understand electricity in the 19th and 20th centuries. The first recorded usage of hydropower to produce electricity occurred in 1880 when a hydro turbine in a chair factory was connected to a dynamo and used to power theater and storefront lights in Grand Rapids, Michigan [1] [2]. Thereafter, the utilization of hydropower to generate electricity exploded. The next year a flour mill turbine in Niagara Falls, New York was connected to a dynamo to generate electricity for city street lights [1]. Then, in 1882 the world's first hydroelectric power plant began operation in Appleton, Wisconsin [3]. After that, hydropower rapidly became a primary element of the nation's power production portfolio. By 1886 there were over 40 hydroelectric plants either in operation or under construction in the United States. By 1889, two hundred different electric companies were reported to use hydropower for some (or even all) of their generation [1]. The development of hydroelectric power resources often went hand-in-hand with the development of irrigation in the West. The U.S. Bureau of Reclamation was established in 1902 for the primary purpose of encouraging increased settlement in the West by developing the water resources to provide irrigation. The Bureau's first major project was the Theodore Roosevelt Dam on the Salt River, which was constructed to provide irrigation and flood control to the region northeast of Phoenix, Arizona. However, in 1906 the Bureau was further authorized by Congress to also develop and sell hydroelectric power [4]. By 1916 the Bureau's installations powered nine irrigation pumps delivering water to over 10,000 acres of land while also supplying all of the power for the city of Phoenix [1]. By the 1940's approximately 30% of the nation's electricity was supplied by hydroelectric power. Since then the percentage of electricity supplied by hydropower has gradually decreased due to rapidly increasing demand as well as the decreasing number of sites suitable and available for hydropower development. Today hydroelectric power accounts for approximately 8% of the nation's power usage, which has increased by nearly a factor of 10 from the 1940's. However, throughout all of this time, hydropower has continued to develop as improvements have been made to the efficiencies of turbines and existing sites have been upgraded. Furthermore, in the 1960's, as the power of the digital computer was becoming more evident, its computational abilities were increasingly applied to hydroelectric plant operation optimization within the constraints of regulated environments.

1.2 Motivation for Studying Optimization

Before the advent of the deregulated power market, the value of hydropower production was typically considered to be equal to the marginal cost of producing the same amount of power using thermal generation sources [5]. However, in recent years many regions of the world, including parts of the United States, have adopted a deregulated (pool-based) market structure [6]. In this market structure, power production facilities submit hourly bids (sometimes even more frequently) for the provision of power production and ancillary services to the market. There has thus been a shift in the goals of hydropower operators away from simply relieving thermal power plants of peak load production, to simply operating in a manner that generates the highest revenue for the hydropower company [7]. This can be achieved with relative ease for an isolated hydropower facility, but the problem becomes much more complex when it is extended to multiple hydrologically connected hydropower stations [8].

Many different optimization methods have been used in the past to solve the short-term hydro scheduling optimization problem (short-term referring to the hourly scheduling for a single day) for such an interconnected system. Methods used include dynamic programming, evolutionary and genetic algorithms, linear programming, mixed-integer linear programming, and mixed-integer quadratic programming [9]. Each of these methods has its own advantages and disadvantages for finding an optimal system commitment, but one crucial aspect that is still lacking from every reviewed optimization method is consideration of the fact that the ensuing commitment may not be optimal for all of the individual facilities. Furthermore, much of the work done in the past has overly simplified the relationship between flow rate, net head, and power production. The effect of flow rate on net head has often been neglected, as has the effect of variations in the net head on the power production in order to make the problem more tractable. Finally, hydropower facilities are typically modeled as either consisting of a number of independent turbines, each with its own power curve, or as a single net power curve representing total power production as a function of net flow rate, with predefined distributions of flow for any given operating point. Both of these approaches have difficulties.

The first issue, regarding the fact that system optimization will typically result in an operation schedule that is sub-optimal for some of the hydropower stations is the main point addressed by this work. In a system comprised of independently owned hydropower stations, each owner is primarily interested in maximizing the revenue produced by his own hydropower facilities. Unfortunately, this tends to result in suboptimal operation of the whole system. There are two possible ways that such a system can be optimized: a regulating entity can force the various owners to operate according to the system optimal schedule (possibly dividing the resulting revenue among the various owners in some way), or the owners can cooperate financially in such a way that any increased revenue of one station due to a change in the operations of another station results in a division of the resulting increase in total revenue. The second option is the one that is explored herein. This work also attempts to address some of the issues that arise from use of the following methods.

When turbines are modeled as individual units, there is a very large number of possible operating conditions, because each turbine has to be either on or off, and the flow rate through each turbine that is on must be determined. Furthermore, there is calculation loop that occurs when any change is made to an operating point: a change in the flow rate through a single turbine will change the net flow, which will affect the net head, which will affect the efficiency of every turbine, which will affect the changes needed to approach the optimal operating conditions.

When a station is modeled using a single power curve, the optimization becomes much easier, but a very important part of the analysis can become obscured. Start-up costs for individual turbines are estimated to be approximately three dollars per rated megawatt, according to a survey made of various hydropower companies in 1997 [10]. Though the current value is likely a bit larger than this due to inflation, it is obvious that start-up costs are a significant factor in the determination of the optimal commitment strategy. Simply operating at maximum efficiency would result in generating the highest possible amount of total power for any given flow rate, but the losses incurred due to the increased number of required start-ups and shut-downs often outweigh the increased revenue for the additional power generated.

This work attempts to address many of these issues by developing an optimization method that utilizes the structured nature of a hydrologically connected hydropower system (on a system-wide level as well as at a local station level), and the fact that optimal operation scheduling cannot occur without cooperation among the various owners.

In summary, the goal of this project is create to a quantitative method whereby financial cooperation between the various owners of various hydrologically connected hydropower facilities results in the optimization of the entire interconnected system of hydropower facilities as well as the revenue of each individual facility, with no exceptions. This approach serves the dual purpose of both maximizing each installation's revenue and maximizing the utilization of the available water in the system by delivering the generated power when it is most needed.

1.3 Hydropower Optimization Literature Review

Some of the earliest literature relating to the application of computational methods to optimizing hydropower dates back to the 1960's, but it wasn't until the 70's that much attention was paid to the development of computational models and methods for use in the optimal scheduling of hydropower. Perhaps the earliest and most widely used method is dynamic programming, which has been applied to the determination of optimal hydropower resource scheduling dating at least as far back as 1961 [11]. In 1970 several researchers at the University of California published a description of a method that utilized dynamic programming techniques to optimize the utilization of the available water in a single reservoir for providing on-peak power and reliable water supply [12]. The next year, several University of Illinois researchers published a paper describing the use of dynamic programming in long-term resource optimization [13]. Computers of the time were not yet powerful enough to consider the full range of possibilities present in a dynamic programming analysis without significant computing cost, so the researchers developed an iterative procedure in which several trial trajectories were used as bases for seeking the optimal strategy. Around the same time computational optimization techniques also began to be employed to assist in planning the development of new reservoir systems. In 1972 researchers described their utilization of mixed integer programming techniques to determine the optimal locations and sizes of reservoirs on a river system, even including considerations for the impacts on water quality control and recreation [14].

Over the next few decades the analyses grew gradually more and more complex as the memory and speed capabilities of computers multiplied. By the early 1980's significantly more complex models and analysis methods began to show up. In 1982 researchers in the University of California developed and published a method using linear programming and dynamic programming to optimize a weighted function of hydropower production, fish protection, water quality maintenance, water supply, and recreational use [15]. In the mid 1980's more attention was devoted to the hourly coordination of hydrothermal power systems. Several different approaches were formulated to optimize the scheduling of a hydrothermal power system, including a mixture of dynamic programming, Lagrangian relaxation, and gradient methods used by Bertsekas [16], as well as a method utilizing stochastic dynamic programming developed by Pereira [17]. An excellent review of these methods, as well as optimization methods based on branch-and-bound techniques, nonlinear mixed-integer programming, and Benders decomposition was given in an invited paper in the Proceedings of the IEEE in 1987 [18].

In the 1990's the attention devoted to operations optimization on an hourly basis continued to grow, but the techniques being used remained largely the same. The only novel approach was the introduction of genetic algorithms by Chen [19], though such methods seem to never have gained much prominence. However, when various regions of the world began operating under a deregulated energy market, a shift in the goals of hydropower optimization occurred. In a region with a deregulated energy market, rather than one of coordinated hydro-thermal generation, hydroelectric facilities have the ability to independently determine their operating schedule. In some ways this made the problem of determining the optimal scheduling of hydropower more straightforward since the hydropower generation no longer had to be coordinated with thermal generation. However, the removal of this constraint also made the problem significantly more complex due to the uncertainties of pricing in a deregulated market.

Hydroelectric power producers also have to consider the effect of selling not only power, but also various ancillary services such as voltage regulation, frequency regulation, and spinning reserves. These, and other, issues began to appear in the literature in the late 1990’s. Much of the work done relates primarily to hydroelectric facilities with little to no market power which are thus often referred to as “price-takers”. In 2002 Conejo published an article on scheduling a hydropower producer in a pool-based electricity market [6]. Using mixed integer linear programming techniques, they formulated piecewise-linear models of nonconcave, head dependent unit performance curves, which were then used to maximize the power revenue of a hydrologically connected system of hydroelectric power plants. Although in some respects they used a significantly more advanced model than many before them, they still neglected to address the effect of future price uncertainty. This matter of future prices was addressed by García-González in a 2007 paper that focused on minimizing the risk of scheduling connected hydropower resources in the day-ahead market [20]. Using an Input/Output Hidden Markov Model they generated possible future price scenarios along with probabilities for each scenario. Two separate methods of limiting the risk were addressed. The first was a simple minimum profit requirement, which they believed to be overly simplistic, and so proposed a second method. This second method was the utilization of a Value-at-Risk analysis, which can be used to limit the risk of a negative deviation from expected profits due to low-probability scenarios. Finally, a recent work relating to optimal scheduling was published in 2012 by Pousinho [9], and uses a similar approach to that of [20]. The main differences were the use of Conditional Value-at-Risk methods and the use of mixed integer quadratic programming rather than mixed integer linear programming.

Although many different methods have been used to optimize the utilization of hydropower resources, none of the reviewed methods addressed the main focuses of this work, namely, optimizing multi-owner systems, and valuing released water based on future downstream revenue potential.

1.4 Complete Mathematical Problem Definition

The optimization of a complete hydropower system is necessarily a highly complex undertaking due to the large number of variables that come into play. The operators of every facility in a system must not only determine which turbines should be active during each period in the planning horizon, but they must also determine the flow rate through each active turbine during each period. The combination of these two parameters has many far-reaching effects. The first and most obvious is that the active turbines and their flow rates will determine the power output of the facility. However, there are several other effects that must be considered, including, but not necessarily limited to: available power for regulation, water level in the tailrace affecting the net head, and start-up costs for turbine selection decisions. Furthermore, an operator may want to take into consideration the uncertainty of predicted electricity prices when planning daily operations, which introduces yet another level of complexity into the optimization. Below is a complete mathematical statement of the optimization problem in which the goal of the optimization is to maximize the revenue of a system of connected hydropower facilities.

Maximize:

$$Revenue = R = \sum_{t}^{T}\sum_{j}^{J}\left(\rho_{j,t}P_{reg,j,t} + \pi_{j,t}\sum_{i}^{I_j} P_{i,j,t}\phi_{i,j,t} - \sum_{i}^{I_j} SU_{i,j,t}\right) \qquad \textbf{1.1}$$

This equation specifies that the object of the optimization is to maximize the revenue produced by all of the hydropower facilities in a system. The summations are performed over the sets $T, J, and\ I_j$, which refer to the time period of the planning horizon, the set of hydropower facilities in the system, and the sets of turbines in each hydropower facility, respectively. The values being summed are, in order, the product of the regulation price $\left(\rho_{j,t}\ \left(\frac{\$}{MWhr}\right)\right)$ and the regulation power $\left(P_{reg,j,t}\ (MWhr)\right)$ being sold by plant j during time period t, the product of the power price $\left(\pi_{j,t}\ \left(\frac{\$}{MWhr}\right)\right)$, power $\left(P_{i,j,t}\ (MWhr)\right)$, and unit status $\left(\phi_{i,j,t}\right)$ of turbine i in plant j, during time period t, and lastly, the start-up costs $\left(SU_{i,j,t}\ (\$)\right)$ (which are on the order of three dollars per rated megawatt [10] incurred by each turbine for each time period. Note that the power produced by a turbine becomes irrelevant if it is inactive.

Subject To:

$$\phi_{i,j,t} \in \{0, 1\} \qquad \textbf{1.2}$$

This term refers to the status of an individual turbine at any time. A value of 0 indicates that a turbine is inactive, while a value of 1 indicates that it is active.

$$SU_{i,j,t} = \phi_{i,j,t}(\phi_{i,j,t} - \phi_{i,j,t-1}) \times P_{nom,i,j} \times k \qquad \textbf{1.3}$$

The start-up costs are calculated based upon the current and previous on/off status (ϕ) of a turbine. This cost is linearly related to the nominal capacity of the turbine $(P_{nom,i,j})$.

$$Q_{i,j,t} \in [Q_{min,i,j}, Q_{max,i,j}] \qquad \textbf{1.4}$$

The flow rate through a turbine must be within given acceptable operating limits.

$$Q_{j,t} = \sum_{i \in I_j} Q_{i,j,t}\phi_{i,j,t} \qquad \textbf{1.5}$$

The net flow rate through plant j during time period t is the summation of the flow through each of its active turbines.

$$Q_{j,t} \in [Q_{min,j}, Q_{max,j}] \cap [Q_{j,t-1} - \Delta Q_{max,j}, Q_{j,t-1} + \Delta Q_{max,j}] \qquad \textbf{1.6}$$

The net flow rate though an installation must not only be within the release limits for the river, it must also be such that rate of change of the flow rate is within prescribed boundaries, preventing sudden changes in the downstream flow.

$$P_{i,j,t} = \rho_{H_2O} \times g \times H_{j,t}(Q_{j,t}, V_{j,t}) \times Q_{i,j,t} \times \eta_{i,j,t}(Q_{i,j,t}, H_{j,t}) \qquad \textbf{1.7}$$

The power produced by an individual turbine is the product of the water density, the acceleration due to gravity, the net available head $(H_{(j,t)})$ (determined by the amount of water in the reservoir and the net flow rate through the station), the flow rate of water through the turbine $(Q_{i,j,t})$, and the efficiency $(\eta_{i,j,t})$ of the turbine-generator (itself a function of the flow through the turbine and the net head, which is assumed to be a convex function). It is generally assumed that these values are constant over a single time period.

$$P_{j,t} = \sum_{i \in I_j} P_{i,j,t}\phi_{i,j,t} \qquad \textbf{1.8}$$

The net power produced by plant j during time period t is the sum of the power produced by each turbine in plant j that is active during time period t.

$$P_{reg,j,t} = \min\left(\sum_{i \in I_j} P_{max,i,j}\phi_{i,j,t} - P_{j,t}, \sum_{i \in I_j} P_{j,t} - P_{min,i,j}\phi_{i,j,t}\right) \qquad \textbf{1.9}$$

The regulation power capability of plant j during time period t $(P_{reg,j,t})$ is determined for the entire station at once. Regulation power is the amount by which the station operators can vary the station's net power output within a 5-minute window. It is assumed here that any active turbine can vary its power within its full operating range in that time. The regulation power available is considered to be the smaller of the available net increase and the available net decrease in power.

$$V_{j,t} = V_{j,t-1} - Q_{j,t-1} + \sum_{k \in K_j} Q_{k,t-1-\tau_{k,j}} + Q_{in,j,t} \qquad \textbf{1.10}$$

The volume of water in each plant's reservoir $\left(V_{j,t}\right)$ at the beginning of each planning period is updated based upon the reservoir contents at the beginning of the previous period $\left(V_{j,t-1}\right)$, the water released from the reservoir during the previous time period $\left(Q_{j,t-1}\right)$, the summation of water flow into the reservoir which was previously released from neighboring upstream reservoirs, and the water that flows in from uncontrolled natural sources $\left(Q_{in,j,t}\right)$. A new set, K_j, is introduced here. The set K_j consists of all stations immediately upstream from station j. The new term in the subscript of the flow rate in the summation, $\tau_{k,j}$, is the time that it takes for water to flow from station k to station j.

$$V_{j,t} \in \left[V_{min,j}, V_{max,j}\right] \qquad \textbf{1.11}$$

The volume of water in each reservoir must, at all times, adhere to appropriate bounds.

$$V_{j,1} = \mathrm{V}_{\mathrm{j}}^{1} \qquad \textbf{1.12}$$

$$V_{j,T+1} = \mathrm{V}_{\mathrm{j}}^{\mathrm{T+1}} \qquad \textbf{1.13}$$

$$Q_{j,0} = Q_j^0 \qquad \textbf{1.14}$$

$$\phi_{i,j,0} = \phi_{i,j}^0 \qquad \textbf{1.15}$$

Finally, the initial conditions and final target reservoir volume must be given.

The full set of optimization equations for the above constraints are compiled on the following page, with set notations replaced by inequalities.

Maximize:

$$Revenue = R = \sum_{t}^{T}\sum_{j}^{J}\left(\rho_{j,t}P_{reg,j,t} + \pi_{j,t}\sum_{i}^{I_j}P_{i,j,t}\phi_{i,j,t} - \sum_{i}^{I_j}SU_{i,j,t}\right) \qquad \textbf{1.16}$$

Subject To:

$$\phi_{i,j,t} \in \{0,1\} \qquad \textbf{1.17}$$

$$SU_{i,j,t} = \phi_{i,j,t}\left(\phi_{i,j,t} - \phi_{i,j,t-1}\right) \times P_{nom,i,j} \times k \qquad \textbf{1.18}$$

$$Q_{min,i,j} - Q_{i,j,t} \leq 0 \qquad \textbf{1.19}$$

$$Q_{i,j,t} - Q_{max,i,j} \leq 0 \qquad \textbf{1.20}$$

$$Q_{j,t} = \sum_{i \in I_j} Q_{i,j,t}\phi_{i,j,t} \qquad \textbf{1.21}$$

$$Q_{j,t} - Q_{max,j} \leq 0 \qquad \textbf{1.22}$$

$$Q_{j,t} - \left(Q_{j,t-1} + \Delta Q_{max}\right) \leq 0 \qquad \textbf{1.23}$$

$$Q_{min,j} - Q_{j,t} \leq 0 \qquad \textbf{1.24}$$

$$\left(Q_{j,t-1} - \Delta Q_{max}\right) - Q_{j,t} \leq 0 \qquad \textbf{1.25}$$

$$P_{i,j,t} = \rho_{H_2O} \times g \times H_{j,t}\left(Q_{j,t}, V_{j,t}\right) \times Q_{i,j,t} \times \eta_{i,j,t}\left(Q_{i,j,t}, H_{j,t}\right) \qquad \textbf{1.26}$$

$$P_{j,t} = \sum_{i \in I_j} P_{i,j,t}\phi_{i,j,t} \qquad \textbf{1.27}$$

$$P_{reg,j,t} = \min\left(\sum_{i \in I_j} P_{max,i,j}\phi_{i,j,t} - P_{j,t}\,, \sum_{i \in I_j} P_{j,t} - P_{min,i,j}\phi_{i,j,t}\right) \qquad \textbf{1.28}$$

$$V_{j,t} = V_{j,t-1} - Q_{j,t-1} + \sum_{k \in K_j} Q_{k,t-1-\tau_{k,j}} + Q_{in,j,t} \qquad \textbf{1.29}$$

$$V_{j,t} - V_{max,j} \leq 0 \qquad \textbf{1.30}$$

$$V_{min,j} - V_{j,t} \leq 0 \qquad \textbf{1.31}$$

$$V_{j,1} - \mathrm{V}_{\mathrm{j}}^{1} = 0 \qquad \textbf{1.32}$$

$$V_{j,T+1} - \mathrm{V}_{\mathrm{j}}^{\mathrm{T+1}} = 0 \qquad \textbf{1.33}$$

$$Q_{j,0} - Q_{j}^{0} = 0 \qquad \textbf{1.34}$$

$$\phi_{i,j,0} - \phi_{i,j}^{0} = 0 \qquad \textbf{1.35}$$

2 Approach

2.1 Preliminary Analysis

2.1.1 Preprocessing of Data

Before performing any optimization it is necessary to first ensure that the data is organized into a useful format. There are two primary sets of variables in this problem for which reorganization is necessary. The first one is the net head at each station, which depends on the reservoir volume and the net flow rate through the particular station. The second set consists of the turbine efficiencies, which are typically assumed to be a convex function of flow rate and head (or near enough to convex that this approximation does not introduce significant error). The relationship between head, reservoir volume, and net flow rate is straightforward, as is shown in the following description of possible methods for determining this relationship. First, it must be recognized that the effects of the reservoir volume and the net flow rate on the net head are independent of each other, and can be treated separately. Then there are several options to determine the effect of reservoir volume on the net head. One option would be a direct calculation using topographical data and numerical integration of the surface area of the reservoir with respect to the height of the surface. This would result in the most accurate, but also quite complex, relationship between head and reservoir volume. Another option would be to empirically determine the relationship by tracking the flow rates into and out of the reservoir and the height of the surface of the reservoir, then generating a curve fit for this data. For the purpose of this work, it was assumed that the volume of water in the reservoir is a polynomial function of the head, measured from the bed of the tailrace. The dependence of net head on the net flow rate can also be determine through either analytical or empirical means. Again, for the purpose of this work it is assumed that the relationship between the net flow rate and net head can be approximated by treating the net flow rate as a polynomial function of the net head, as measured from the bed of the tailrace.

The second set of variables in the problem requiring a preliminary analysis consists of the turbine efficiencies. Even if high quality, accurate data were available, it likely would not be in a form conducive to use in optimization, and must be preprocessed to transform it into a useable form. In this instance, the data available identifies iso-efficiency curves for a hill chart (which shows the efficiency of a hydro turbine as a function of nondimensionalized head and flow rate). Figure 1 shows an example of how these iso-efficiency curves appear when plotted. The numerical optimization calls for the data be discretized over the operating range such that, when given values for head and flow rate, the turbine power can be directly determined. Finally, it is expected that output power is a concave function of flow rate if the head is held constant [21].

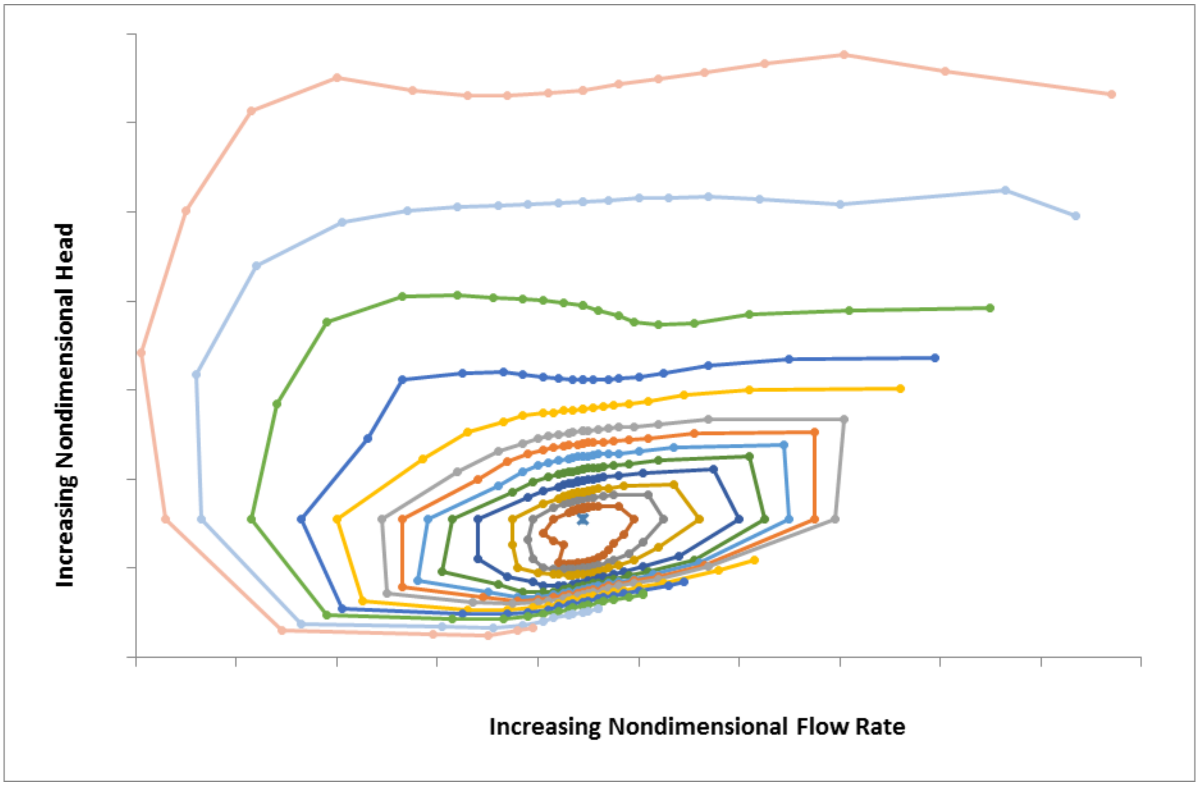

Figure 1: Iso-efficiency lines of a Hill chart. Highest efficiency occurs in the location marked with an 'x'.

In order to transform the given data into this more useful form, several steps are required. First, the power output at each operating point is determined using Eq. 1.26. Next, a triangulation of the given flow rate and head values is generated using MATLAB's built-in functionality. In order for the surface interpolation function to accurately represent the expected shape of the power surface, it is necessary to impose some constraints on the triangulation. The applied constraints are selected so that points on the flow rate – head plane are not be connected to other points associated with non-sequential levels of efficiency (i.e. if iso-efficiency curves were available for 0.89, 0.90, and 0.91, then an operating point associated with an efficiency of 0.89 could not be connected to an operating point with an efficiency of 0.91 or with any operating points of the same efficiency (0.89) except its immediate neighbors). Finally, the triangulation is used in conjunction with the previously calculated power values and an interpolation function to create convex power curves as functions of flow rate for various values of head. The values of head used are determined based upon requirements detailed in the next preliminary analysis stage. The particular results of this stage of the analysis are further detailed in Chapter 3.

2.1.2 Preliminary Analysis of Unit Combinations

2.1.2.1 Utility

The optimization of even a single hydropower facility is not a trivial task. The computational resources required to optimize the facility's operations increase drastically with the number of turbines in the facility and with finer levels of discretization. The equation for the total number of possible ways to operate a facility for a single day is given by Eq. 2.1.

$$N = ((N_{disc} + 1)^{N_{turb}})^{24} \qquad \textbf{2.1}$$

N is the total number of possible operating strategies for a single day, N_{disc} is the number of discrete flow rates each turbine's operation has been divided into, and N_{turb} is the number of turbines in the facility. For a single day a facility with 4 turbines, each of which has been discretized to allow for the analysis of 9 different flow rates (10 including a flow rate of zero indicating that the turbine is inactive) would have 10^{96} possible operating strategies. Obviously, it would be impossible, even with modern computers, to analyze this number of possible operating strategies. Fortunately, there are ways of approaching the problem that make analysis of every single possibility unnecessary. However, even using dynamic programming (the application of which will be detailed later, in Section 2.3), the number of possible operating strategies that must be analyzed is given by Eq. 2.2.

$$N = 24 \times (N_{disc} + 1)^{N_{turb}} \qquad \textbf{2.2}$$

Obviously, the number 24 cannot be reduced as it is the number of hours in a day. However, if $(N_{disc} + 1)^{N_{turb}}$ can be decreased by a preliminary analysis of each possible combination of operating turbines, then any further analysis will be much faster. In fact, it turns out that this *can* be pre-analyzed, resulting in large potential savings in the required computational time.

2.1.2.2 Justification

The second stage of the preliminary analysis could also be considered the first stage of the optimization. This pre-optimization of the problem definition makes the subsequent analysis less cumbersome by reducing the number of independent variables. To illustrate the utility of this step, consider the following: Using the full problem definition from Section 1.4, the efficiency, and thus the power production, of any given turbine depends on the net head, and thus the net flow rate. Therefore, in order to determine the effect of any adjustment to the flow rate through a single turbine, it is necessary to recalculate the net flow rate and then net head. This change in net head will necessarily also simultaneously affect every other operating turbine. Furthermore, during the optimization, adjustment of the flow rate through a single turbine necessitates adjustment of the flow rates through all of the other active turbines such that the gradient of the revenue with respect to the flow through each active turbine remains the same. Finally, it is also possible when adjusting the flow through a single turbine to arrive at an operating state that would be better served by a different combination of turbines, requiring some turbines to be activated or deactivated, and the whole process to start over again. Therefore any time a change is made to the operating point of a single turbine in the search for the optimal operating points, a sub-problem must be solved to determine the corresponding required changes in the other active turbines, and it is reasonably projected that the same sub-problems would have to be solved multiple times throughout the course of each optimization.

Consider next an operating point that is known to be the optimal operating point for the entire system.

$$\exists\{\phi_{i,j,t}, Q_{i,j,t}\}_{opt}\ \forall\ \{i,j,t\} \in \{I_j, J, T\}\ \Big|\ R\left(\{\phi_{i,j,t}, Q_{i,j,t}\}_{opt}\right) = R_{max} \qquad \textbf{2.3}$$

This operating point defines which turbines should be active at each station for each time period in the planning horizon. Further, it defines the specific flow rates through each of those turbines. Thus, for each station j during each time period t in the planning horizon, the optimal operating point results in a known net flow $\left(Q_{j,t,opt}\right)$ and net power production $\left(P_{j,t,opt}\right)$ for each station in the system. Consider the independent variables to be grouped by station and time period $\left(\{\phi_i, Q_i\}_{j,t,opt}, \ldots\right)$. Thus, each group of independent variables completely defines the status of a single station during a single time period of the optimization. Each of these subsets of independent variables must themselves result in the optimal division of each station's net flow among its active turbines (where optimal division of flow among turbines is that division which results in the highest power production for a given net flow rate).

$$(R = R_{max}) \rightarrow \left(P_{j,t,opt} = P_{j,t,max},\ \text{for given}\ Q_{j,t} = Q_{j,t,opt}\ \&\ \{\phi_\mathrm{i}\}_{\mathrm{j,t}} = \{\phi_i\}_{j,t,opt}\right) \qquad \textbf{2.4}$$

To prove this, one must only consider the consequences if this were not the case. If the division of flow among the active turbines were not optimal, then there would exist another division of flow among the active turbines for which the net flow would be equal to the net flow of the optimal combination (shown below in Eq. 2.5), and the station's power production would be higher than the power production resulting from the optimal system operating point.

$$\text{i.e. assume}\ \exists\{Q_i\}_{j,t} \neq \{Q_i\}_{j,t,opt}\ |\ Q_{j,t} = Q_{j,t,opt}\ \&\ P_{j,t} > P_{j,t,opt} \qquad \textbf{2.5}$$

Were this true, more power could be produced using the same amount of water and the only effect on the system would be to increase the power output of the station under consideration. This would in turn result in a higher total revenue, contradicting the original assumption that the known optimal operating point is, in fact, optimal. Thus, in order to uniquely identify the optimal operating point for a hydropower system, it is only required that the net flow rate and active turbines be known for each station in the system, rather than knowing the flow through each turbine.

In order to use the net flow and active set of turbines as the independent variables, rather than the flow through and status of each turbine, it is necessary to know the power output of each station as a function of the net flow rate and the active set of turbines. Fortunately this is a sub-problem that can be predicted and preemptively solved [21] over the entire operating range through application of the Karush–Kuhn–Tucker (KKT) conditions. It is also straightforward to use this transformation, the details of which are developed below) to do away with the difficulty of constantly recalculating the net head for each updated value of flow rate, and instead express power as a function of the volume of water in the reservoir, the net flow rate through the station, and the active set of turbines.

2.1.2.3 Methodology

Some new definitions must be made before continuing with the optimization of this preliminary sub-problem. Previously, in Section1.4, the set of turbines in a station was identified as I_j, in which each turbine was identified by a unique number "i" within its station. Before continuing, in order to reduce the total number of active sets of turbines that must be analyzed, each turbine's numeric identifier is changed such that identical turbines have the same identifier. This creates set I'_j. Next, set Ψ_j is defined as the set of unique elements of the power set of I'_j, excluding the null set.

With this new definition, the application of the KKT conditions to the preliminary optimization can begin. The objective is to determine the optimal power surfaces as functions of only reservoir volume (V_j), net flow rate (Q_j), and active operating set (ψ), thus generating power surfaces for each operating set in each facility.

The formal definition of the sub-problem as required for applying the KKT conditions is as follows:

Maximize:

$$P_{\psi,j}(Q_j, V_j) = \max \sum_{i \in \psi} P_{i,j}\left(Q_{i,j}, H_j(Q_j, V_j)\right) \qquad \textbf{2.6}$$

Subject To:

$$\sum_{i \in \psi} Q_{i,j} - Q_j = 0 \qquad \textbf{2.7}$$

$$Q_{i,j} - Q_{max,i,j} \leq 0 \ \forall \ i \in \psi \qquad \textbf{2.8}$$

$$Q_{min,i,j} - Q_{i,j} \leq 0 \ \forall \ i \in \psi \qquad \textbf{2.9}$$

The only additional equation needed to fully specify the problem is the equation of the power output of a turbine as a function of flow rate and head, which is shown in Eq. 1.26.

This problem must be solved for all discretized points in the following permissible operating domain for each operating set in each station:

$$Q_j \in [Q_{min,j}, Q_{max,j}] \cap [Q_{min,\psi}, Q_{max,\psi}] \quad \textbf{2.10}$$

$$V_j \in [V_{min,j}, V_{max,j}] \quad \textbf{2.11}$$

These equations show the domain of the independent variables. The first equation indicates that the net flow rate (Q_j) can vary from the greater of the minimum station flow rate $(Q_{min,j})$ and the minimum operating set flow rate $(Q_{min,\psi})$ to the lesser of the maximum station flow rate $(Q_{max,j})$ and the maximum operating set flow rate $(Q_{max,\psi})$. The second (Eq. 2.11) simply indicates that the domain of the reservoir volume includes all points between the minimum and maximum allowable volume.

Application of the KKT conditions results in the following set of equations (which are both necessary and sufficient conditions for optimality because the power equation is assumed to be concave and the constraints are all linear):

$$\frac{\partial P_{i,j}\left(Q_{i,j}, H_j(Q_j, V_j)\right)}{\partial \mathrm{Q_{i,j}}} = \mu_{1,i} - \mu_{2,i} + \lambda \quad \textbf{2.12}$$

$$\mu_{1,i} \geq 0, \qquad \mu_{2,i} \geq 0 \quad \textbf{2.13}$$

$$\mu_{1,i}(Q_{i,j} - Q_{max,i,j}) = 0 \quad \textbf{2.14}$$

$$\mu_{2,i}(Q_{min,i,j} - Q_{i,j}) = 0 \quad \textbf{2.15}$$

Thus, if the turbines are working within their operating domains $(Q_{min,i,j} < Q_{i,j} < Q_{max,i,j})$, then $\mu_{1,i} = \mu_{2,i} = 0$ in order to satisfy Eq. 2.14 through Eq. 2.15, and the slopes of the turbine power curves (as functions of turbine flow rate) are all identical and equal to λ. However, if any turbine is operating at an extreme of the allowable range, either $\mu_{1,i}$ or $\mu_{2,i}$ will be non-zero. The analysis of the effects of each possibility are similar at each extreme, so only the effect of a turbine operating at its minimum allowed value will be examined herein. Eq. 2.12 must be valid for each turbine under consideration, so it is known that λ must be equal to the slope of the power vs. flow curve for each turbine that is operating within its allowed domain. However, if a turbine is operating at the lower limit of its permissible range, then $(Q_{min,i,j} - Q_{i,j})$ in Eq. 2.15 will be equal to zero, allowing $\mu_{2,i}$ to take a positive non-zero value. This further indicates that the slope of the power curve of the turbine can take a value other than λ. In the case of a turbine operating at the lower limit of its domain, the slope of its power curve at the optimal operating point will be less than that of other turbines that are not operating at their lower limits since $\mu_{2,i}$ is constrained to be positive. This aligns with the requirement that the objective function be convex.

Thus, the basic condition for optimality is that the turbines all operate such that the derivatives of their power curves with respect to their flow rates are the same, unless operating at a limiting flow rate. This is a relatively straightforward problem to address numerically. The easiest way to do so begins by calculating the derivatives of power with respect to flow rate at each value of flow rate for which the power is known (using numerical differencing methods), for each turbine in the active set. Then for a known the target net flow rate, the flow through each turbine can be determined by iteratively incrementing the flow rate through the turbine with the highest derivative until the desired net flow is achieved. However, a more accurate method is desired in which operating points for the turbines can be found between the discretized points of flow rate for which the power function is known. Such an approach is especially useful when the available discretization of the flow rate is not uniform, as is the case with the data available for this analysis. In such a case it becomes desirable to accurately interpolate the power for values of flow rate between the available discretization. It is also desirable for the net power curve resulting from the adopted interpolation method to not only be continuous, but to also have continuous derivatives. Several options were investigated to accomplish this. First, the slopes that any interpolation function would need to match were calculated using 2nd order methods (central differencing for interior points, and 2nd order forward and backward differencing for flow rates at the edges of the domain). Thus, any interpolation function must have four unknowns so that all of the conditions can be matched. Since a 3rd order polynomial is the simplest such function, its use was the first to be examined. However, it was discovered that this approach would not work as the polynomial which matched the values and slopes already known was often not convex over the entire region between the known values for flow rate. With this in mind, and lacking any previously known method that would produce the desired result, the original interpolation function shown in Eq. 2.16 was created.

$$P = \mathrm{sgn}\sqrt{a \times (Q-b)^2 + c} + d \qquad \textbf{2.16}$$

This equation can take several forms, as shown in Figure 2 – Figure 4. It exhibits several characteristics that make it ideal for interpolating the power as a function of flow rate. First, it is concave over its entire domain. Second, the coefficients can be solved (though the solution is admittedly complex) for known values for power and slope at two values of flow rate as shown below in Eq. 2.19 – Eq. 2.23 (Mathematica was used to solve for the coefficients). Finally, the slope is a reversible function of the flow rate, allowing the flow rate corresponding to any possible value of the slope to be calculated, as shown in Eq. 2.24. The astute reader will notice that these are simply different conic sections. Figure 2 and Figure 4 show sections of hyperbolas, while Figure 3 shows a section of an ellipse.

Given that $q_1, q_2, p_1 = p(q_1), p_2 = p(q_2), p_1' = p'(q_1), \text{and}\, p_2' = p'(q_2)$ are known: **2.17**

$$dq = q_2 - q_1, \&\, dp = p_2 - p_1 \qquad \textbf{2.18}$$

$$sgn = -\frac{a \times c}{|a \times c|} \qquad \textbf{2.19}$$

$$a = \frac{dp\big(2 \times dq \times p_1' \times p_2' - dp(p_1' + p_2')\big)}{dq(dq(p_1' + p_2') - 2dp)} \tag{2.20}$$

$$b = q_1 - \frac{dq \times p_1' \times (dp - dq \times p_2')}{2dq \times p_1' \times p_2' - dp(p_1' + p_2')} \tag{2.21}$$

$$c = \frac{dp(dp - dq \times p_1')^2 (p_1' + p_2')(dp - dq \times y_2')^2}{(dq(p_1' + p_2') - 2dp)^2 \times (dp(p_1' + p_2') - 2dq \times p_1' \times p_2')} \tag{2.22}$$

$$d = p_1 + \frac{a \times c}{|a \times c|} \times \sqrt{\frac{dp^2(dp - dq \times p_2')^2}{(dq(p_1' + p_2') - 2dp)^2}} \tag{2.23}$$

$$q = -\frac{c}{|c|} \times p' \times \sqrt{\frac{c}{a^2 - a \times p'^2}} + b \tag{2.24}$$

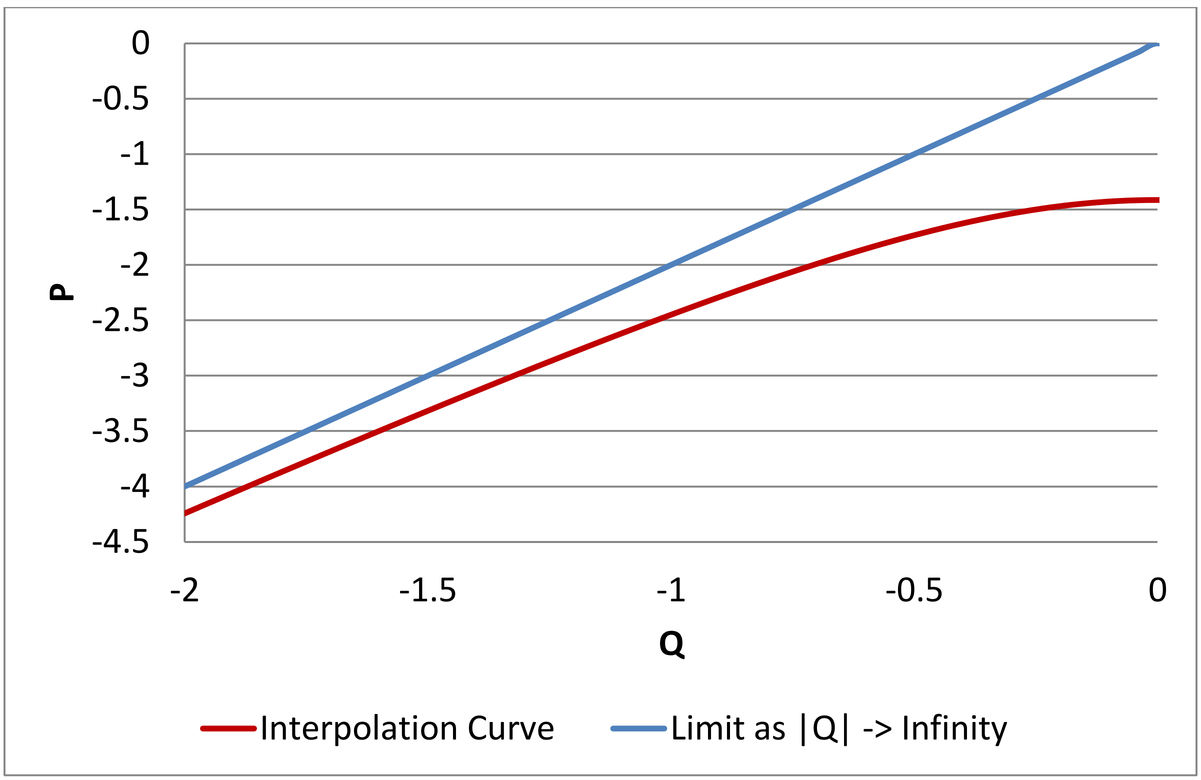

Figure 2: Interpolation curve shape when coefficients 'a' and 'c' are both positive.

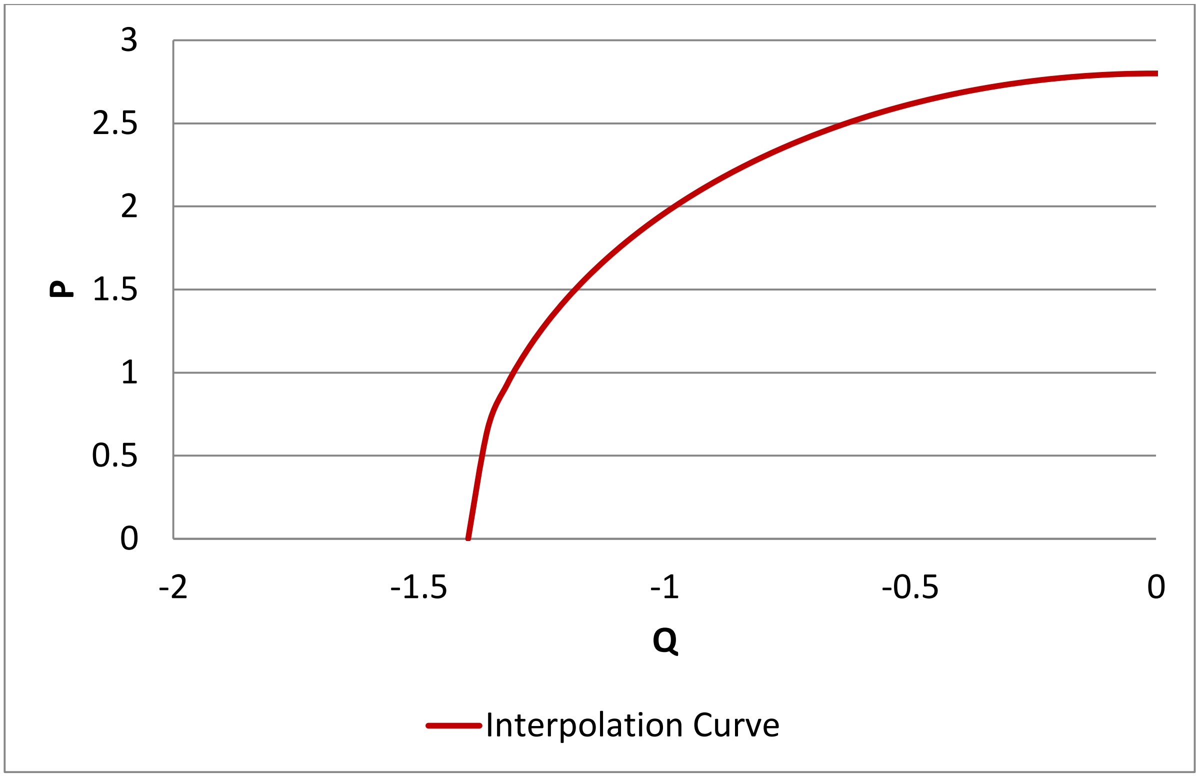

Figure 3: Interpolation curve shape when coefficients 'a' and 'c' are negative and positive, respectively.

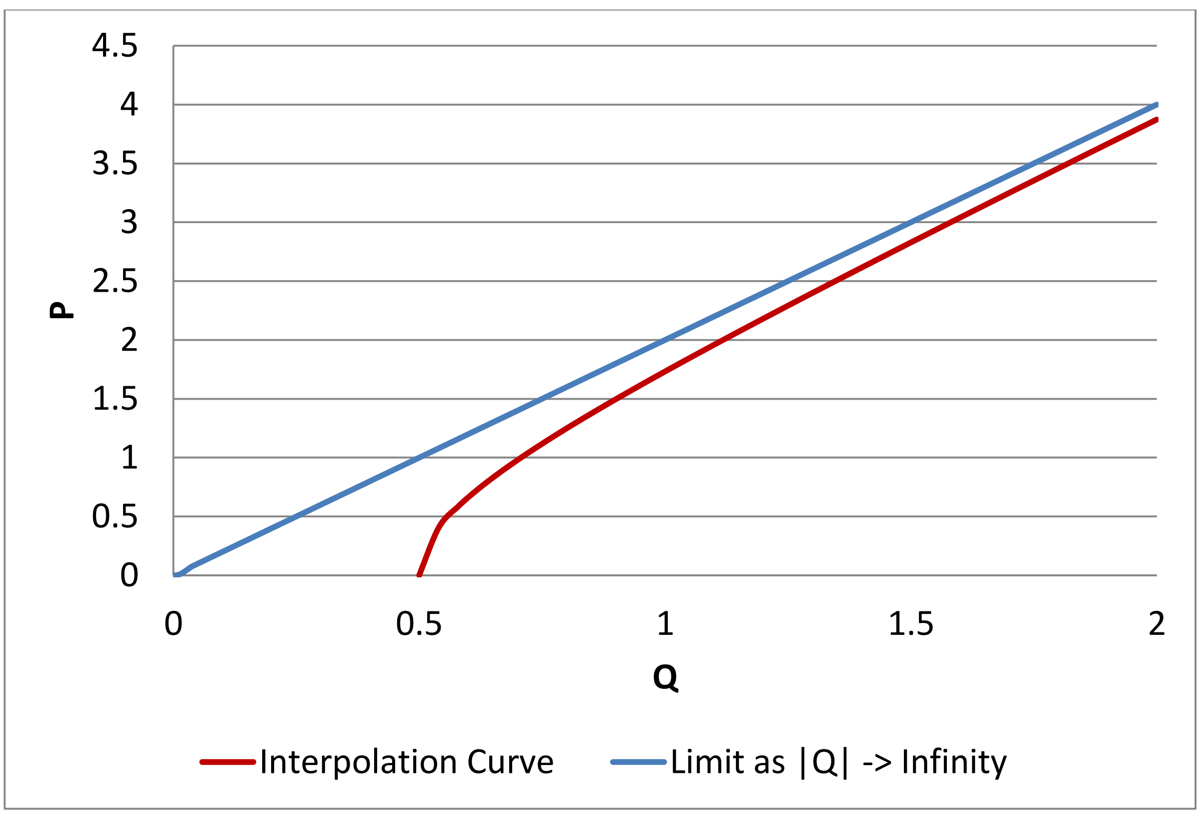

Figure 4: Interpolation curve shape when coefficients 'a' and 'c' are positive and negative, respectively.

Utilization of this interpolation function between each pair of adjacent values of flow rate for which power is known allows the power to be calculated as a function of any flow rate in each turbine's domain. At this point, determining the flow rates that solve the sub-problem detailed in Eq. 2.6 – Eq. 2.9 requires only that the correct value of the slope be determined. This determination is accomplished via an iterative approach which yields the value of λ (and thus the flow through each turbine) that corresponds to the desired net flow rate.

Once all of the preliminary analysis is complete, the problem definition is reduced to the following:

Maximize:

$$Revenue = R = \sum_{t}^{T}\sum_{j}^{J}\left(\rho_{j,t}P_{reg,\psi_j,j,t} + \pi_{j,t}P_{\psi_j,j,t} - SU_{j,t}\right) \qquad \mathbf{2.25}$$

Subject To:

$$\psi \in \Psi_j \qquad \mathbf{2.26}$$

$$SU_{j,t} = SU(\psi_{j,t}, \psi_{j,t-1}) \qquad \mathbf{2.27}$$

$$Q_{j,t} - Q_{max,j} \le 0 \qquad \mathbf{2.28}$$

$$Q_{min,j} - Q_{j,t} \le 0 \qquad \mathbf{2.29}$$

$$Q_{j,t} - \left(Q_{j,t-1} + \Delta Q_{max}\right) \le 0 \qquad \mathbf{2.30}$$

$$\left(Q_{j,t-1} - \Delta Q_{max}\right) - Q_{j,t} \le 0 \qquad \mathbf{2.31}$$

$$P_{\psi,j,t} = P_{\psi,j}(Q_{j,t}, V_{j,t}) \qquad \mathbf{2.32}$$

$$P_{reg,j,t} = \min\left(P_{max,\psi,j} - P_{\psi,j,t}, P_{\psi,j,t} - P_{min,\psi,j}\right) \qquad \mathbf{2.33}$$

$$V_{j,t} = V_{j,t-1} - Q_{j,t-1} + \sum_{k \in K_j} Q_{k,t-1-\tau_{k,j}} + Q_{in,j,t} \qquad \mathbf{2.34}$$

$$V_{j,t} - V_{max,j} \le 0 \qquad \mathbf{2.35}$$

$$V_{min,j} - V_{j,t} \le 0 \qquad \mathbf{2.36}$$

$$V_{j,1} - \mathrm{V}_{\mathrm{j}}^{1} = 0 \qquad \mathbf{2.37}$$

$$V_{j,T+1} - \mathrm{V}_{\mathrm{j}}^{\mathrm{T+1}} = 0 \qquad \mathbf{2.38}$$

$$Q_{j,0} - Q_j^0 = 0 \qquad \mathbf{2.39}$$

$$\psi_{j,0} - \psi_j^0 = 0 \qquad \mathbf{2.40}$$

Thus, the preliminary analysis has reduced what was previously an optimization problem involving three nested summations to a problem with only two nested summations.

2.2 Proposed New Optimization Approach

The final goal of this thesis is to produce a method whereby optimal system commitment can be achieved in a set of hydropower facilities owned by various individuals or corporations, which inherently seek out only their own optimal strategy. Along the way to achieving this goal, a considerable simplification of the system optimization objective function has been accomplished by developing the analysis described in the previous section, which led to the mathematical problem statement given in Eq. 2.25. However, the optimization objective of this equation does not yet include any means of ensuring that the individual stations in the system are optimally operated. In fact, even without this insightful simplification, the original problem as stated in Eq. 1.1 can be (and has been) solved for hydropower systems via mixed-integer optimization strategies. Yet, since the goal of this optimization includes not only optimizing the system as a whole, but also optimizing each individual plant's operations, use of mixed-integer methods become less desirable due to the increase in complexity necessary to accommodate multiple objective functions, to determine optimal operating points, as well as the corresponding financial interactions that allow for individual as well as system optimization.

In order to optimize both the system and each of the individual stations, a model is needed whereby system optimization and individual optimization are intrinsically linked. This necessarily introduces a large new set of independent variables to the optimization (which will be detailed in Section 2.4). The method developed herein incorporates a financial compensation approach whereby plant operators can regard the delivery of water at particular times as a valued commodity with specific monetary value, and pay upstream facilities for such delivery. Furthermore, the monetary value of the water is tied to the value of electrical energy so that the strategy yields results of consequence. Finally, hydropower systems are inherently non-linear, and the introduction of the requirement that both the system and the individual station operations be optimal introduces another layer of nonlinear interactions which makes previously used mixed-integer methods undesirable. Thus, the following method was developed, as outlined in the next three paragraphs, and further described in Sections 2.3 and 2.4.

In order to overcome the shortcomings of mixed integers methods, and determine the degree of cooperation between neighboring facilities necessary to reach a globally optimal commitment, the problem is approached by iterating between two stages. We label this method a dual iterative optimization. Note that the word "dual' refers only to the two stages; in each many different stations and interactions are analyzed. Iterations continue between the two stages, until a regional optimum solution is obtained.

The first stage of the method consists of optimizing the operation of each individual facility, without regard to its effects on other facilities in the system. Since the flow into each station's reservoir is determined by the flow released from upstream facilities, the order in which the facilities are optimized is based on the physical structure of the system. The farthest upstream facilities are the first to be analyzed, and the furthest downstream are the last. This ensures that all of the relevant inputs are available when each optimization is performed. The details of this stage of the optimization are given in Section 2.3.

The second stage of the optimization consists of first analyzing revenue gradients of neighboring pairs of stations with respect to changes in the upstream facility's flow release during each individual time period. Then, for each pair of facilities, the time period during which adjusting the upstream facility's flow release results in the greatest total increase in the two station's revenue is determined. Finally, the value that must be added to the flow release (in addition to the value of the electricity being produced) of the upstream facility is determined and recorded for use in the next iteration of the first stage. Further details of this stage of the optimization are given below in Section 2.4.

2.3 Iteration Stage 1: Individual Hydro Station Optimization

Optimizing a single power station with known values for price and water inflow, while certainly much simpler than optimizing a hydropower system, is still far from a trivial problem. First, the problem definition is as follows:

Maximize:

$$R^1_{opt,j} = \max\left(R_j = \sum_t^T \left(\rho_{j,t} P_{reg,\psi,j,t} + \pi_{j,t} P_{\psi,j,t} - SU_{j,t}\right)\right) \qquad \textbf{2.41}$$

The constraints of this optimization are the same as those given in Eq. 2.26 – Eq. 2.36.

The only difference between this problem statement and the one given in Eq. 2.25 – Eq. 2.36, is that now the summation over the different stations has been removed, and the optimization is only concerned with the results for a single station.

The use of mixed-integer programming for solving this simpler problem was considered, but ultimately dynamic programming was selected due to its flexibility and robustness (two characteristics that are more difficult to achieve using mixed-integer methods). An overview of dynamic programming can be found in most optimization text books, or alternatively, Wikipedia given an excellent overview of the topic and links to many other helpful references.

The following shows the definitions necessary to utilize dynamic programming to solve the optimization problem.

Value Function:

$$VtG_{j,t}(\boldsymbol{x_{j,t}}) = \max_{\boldsymbol{u_{j,t}}} \left(R_{j,t}(\boldsymbol{x_{j,t}}, \boldsymbol{u_{j,t}}) + VtG_{j,t+1}(\boldsymbol{x_{j,t+1}}) \right) \quad \textbf{2.42}$$

State Variables ($\boldsymbol{x_t}$):

$$\mathbf{x_{j,t}} = \{\mathrm{V_{j,t}}, \mathrm{Q_{j,t-1}}, \psi_{\mathrm{j,t-1}}\} \quad \textbf{2.43}$$

Input Variables ($\boldsymbol{u_t}$):

$$\mathbf{u_{j,t}} = \{\mathrm{Q_{j,t}}, \psi_{\mathrm{j,t}}\} \quad \textbf{2.44}$$

State Update Function:

$$\mathbf{x_{j,t+1}} = \left\{ \left(V_{j,t} - Q_{j,t} + \sum_{k \in K_j} Q_{k,t-1-\tau_{k,j}} + Q_{in,j,t} \right), \mathrm{Q_{j,t}}, \psi_{\mathrm{j,t}} \right\} \quad \textbf{2.45}$$

In the above equations, the only new variables introduced are VtG, which it the "value-to-go," and $R_{j,t}$, which is the revenue of station j during time period t. The value-to-go function is iteratively calculated from the state of the system at the end of the final time period to the state before the initial time period for every possible combination of states at each intermediate time. When calculating the value-to-go at each iteration stage, the values of all future state combinations are known, and thus, for each possible current state, the optimal set of inputs is determined based on the revenue the inputs would produce for the current time period and the future revenue (value-to-go) that is possible.

Additionally, the selection of state variable bears an explanation. The selection of the volume of water in the reservoir $(V_{j,t})$ as a state variable should be an obvious choice, but the reasons for the inclusion of the previous flow rate $(Q_{j,t-1})$ and operating set $(\psi_{j,t-1})$ likely are not quite so obvious. They are, however, easily accounted for. The flow during the previous time period must be known in order to enforce the constraint on the maximum change in flow rate during a single time period (shown in Eq. 2.30 and Eq. 2.31). The previously active operating set must be known in order to determine the start-up costs that would be incurred by each possible selection of the current operating set (shown in Eq. 2.27).

Additionally, this formulation of the problem introduces an extra requirement regarding the method of discretization employed for the flow rates and reservoir volume. When using dynamic programming it is advantageous for each set of possible inputs to lead directly from one state to another. While it is possible to apply an interpolation function to determine the value of the future value-to-go function in Eq. 2.42, doing so would result in an avoidable loss of accuracy, and is thus best circumvented if possible. Thus, since flow rate and volume are linked in the state update function shown in Eq. 2.45, the discretization of the volume and flow rates must be such that when the state update function is evaluated, the value found for the updated volume will lie on an existing discretized value rather than between two values in the volume discretization.

The results generated from this optimization include the optimal operating set, net flow rate, and reservoir volume for each facility during each time period of the planning horizon, as well as the resulting optimal total revenue of each facility. These results (shown below in Eq. 2.46) are then used in the next stage of the system optimization routine. The superscript of each term indicates that these are the results from the first optimization of each individual facility.

$$\{\Psi^1_{\mathrm{opt,j,t}}, Q^1_{opt,j,t}, V^1_{opt,j,t}, R^1_{opt,j}\} \quad \mathbf{2.46}$$

2.4 Iteration Stage 2: River System Optimization through Gradient Analysis

This section uses the results of the individual station optimizations and links the optimization of neighboring stations. The procedure developed herein is the mechanism that ties together the individual station optimizations with the total system optimization. Using this procedure, it becomes possible for the optimization of the individual components of the system to drive the overall system optimization. This is achieved by allowing communication of the future, downstream value of water between neighboring stations. This value takes the form of payments offered for altered release levels from downstream facilities to upstream facilities.

The basic concept behind this approach can be easily demonstrated using a simple example. Figure 5 shows a sample electricity price profile, the resulting revenues with no communication, and the resulting revenue if the value of water is communicated between neighboring facilities. In this example there are two facilities that are connected by a river. The upstream facility has a reservoir and is able to hold water until the most beneficial time, but the downstream facility is run-of-river, and can only use water when it is received from an upstream facility. In this example, the flow time between the two facilities is one hour, the upstream facility has enough water to flow for a two hour period, and it is assumed that both facilities can convert water to electricity at the same rate. As shown in Figure 5(a), without communication of the future downstream value, the upstream facility would choose to operate during the two hours with the highest price (the $50 and $40 time periods). This would then constrain the downstream facility to operate at times during which the electricity price is much lower (the $40 and $35 time periods). Without communication, the total

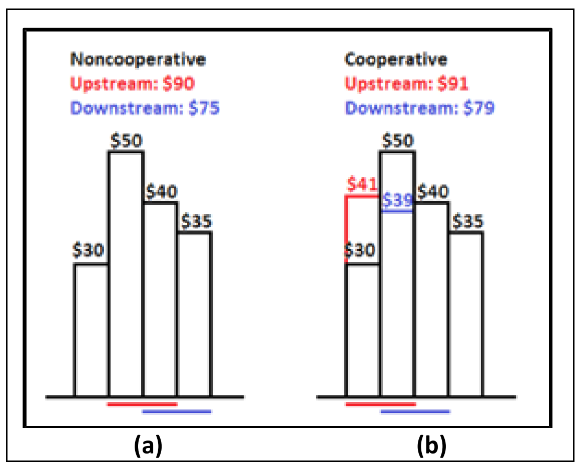

Figure 5: Water value concept example.

revenue of the two plants amounts to $165. If, however, the changes in total revenue were evaluated for various different possibilities then the total revenue of the system would be maximized if the upstream facility operated during the time periods with prices of $30 and $50 (a decrease of $10), which would allow the downstream facility to operate during the time periods with prices of $50 and $40 (an increase of $15), generating a total revenue of $170. However, the upstream facility would not be willing to operate during these time periods because of the decrease in revenue it would experience. Thus, in order to incentivize the upstream facility to release water during the earlier time, it is necessary for the downstream facility to use some of its increased profits to pay for water to be released during the $30 time period. A possible result is shown in Figure 5(b). In this scenario both plants generate greater revenue than previously, and the system as a whole is optimized. It is further seen here that there are many possible payments that would result in the desired change in release. The payment that is necessary to induce the desired change in a real market is a matter of negotiation between neighboring stations and is beyond the scope of this thesis. However, this example successfully demonstrates that enabling communication between neighboring reservoirs using water payments can drive the system toward optimal operation.

Of course, applying this concept to a real system is much more complex, but the basic idea transfers over quite nicely. In order to detail the methods used in this section of the optimization, it is first necessary to first describe the discretization scheme that is used for the variables in the optimization. Specifically, it is necessary to describe the discretization of the operating sets, the net flow, and the volume. Since the operating sets are already divided into discrete groups, it is unnecessary to do anything further with them in this regard. Hence, only the net flow and the reservoir volume must be considered. As was stated at the end of Section 2.3, the discretization of these two variables is linked. First, a discretization size is selected for the net flow, which will be referred to as ΔQ_{grid}. The same discretization is used for all operating sets in a station. This then determines the discretization that must be used for the volume according to Eq. 2.47.

$$\Delta V_{grid} = \Delta Q_{grid} * 3600\ \mathrm{s} \qquad \mathbf{2.47}$$

Thus, the possible flow rates for any given combination of turbines will consist of all of the values of flow rate that are divisible by ΔQ_{grid} and between the minimum and maximum flow limits for that particular operating set.

In order to apply the gradient analysis to a hydropower system consisting of multiple hydropower plants, the facilities are divided into pairs, each consisting of two neighboring facilities (the water released from one flows into the reservoir of the other). Since the same analysis is performed for every pair of stations, the equations developed below apply to calculations within a pairing. After all of the pairs of plants are determined, the first step of this second stage of the optimization is to determine, for each upstream facility in each pair, the effect on the optimal commitment of increasing or decreasing the net flow during each time period in the planning horizon. In particular, the effect of requiring an increase or decrease in the net flow during time period "s" on the optimal station revenue and the optimal net flow profile are sought (Eq. 2.48). Since the procedure required to determine the effect of an increase is basically the same as that required for a decrease, everything developed below will be for the analysis of a flow increment, and can easily be extended to the analysis of a flow decrement. Further, the subscript "j," which refers to which facility is being considered, is replaced with either "up" or "$down$," indicating whether the calculation is for the upstream or downstream facility within a pair.

$$\Delta R_{opt,up}^{1,s+} = R_{opt,up}^{1,s+} - R_{opt,up}^{1}\ ,\ Q_{opt,up,t}^{1,s+} \qquad \textbf{2.48}$$

Of course, requiring an increase in the net flow during one time period means that the net flow must decrease at some other time in order for the constraint on the final reservoir volume to be met. Furthermore, it has been found that requiring the flow to increase at a particular time commonly results in changes in the optimal flow profile during several time periods. Thus, to determine the exact effect of increasing the flow at any particular time, the optimization problem given in Eq. 2.41 (repeated below for reference) must be solved again with an additional constraint specifying the required change in the final flow profile.

Maximize:

$$R_{opt,up}^{1,s+} = \max\left(R_j = \sum_{t}^{T} \left(\rho_{j,t} P_{reg,\psi_j,j,t} + \pi_{j,t} P_{\psi_j,j,t} - SU_{j,t} \right) \right) \qquad \textbf{2.49}$$

Subject To:

$$\text{Eq. 2.26 – Eq. 2.36} \qquad \textbf{2.50}$$

$$Q_{j,t=s} = Q_{opt,j,t=s}^{1} + \Delta Q_{grid} \qquad \textbf{2.51}$$

Of course, repeating the full optimization analysis for each upstream facility twice for each time period in the planning horizon is computationally quite expensive. Thus, additional constraints are added which restrict the values of reservoir volume and flow rate to be considered in the optimization. These additional constraints are given below in Eq. 2.52 – Eq. 2.54. They specify that no flow rate (or volume) that is greater than some integer multiple of the grid size away from the previously determined optimal solution is to be considered in the analysis. Further, the number of allowed grid-sized steps the volume can take away from the previously found optimal values is itself a positive integer multiple of the number of allowed grid-sized steps the flow rate can take away from its previously found optimal values.

$$Q_{j,t\neq s} \in \left\{ Q_{opt,j,t}^{1} - n \times \Delta Q_{grid}, Q_{opt,j,t}^{1} + n \times \Delta Q_{grid} \right\} \qquad \textbf{2.52}$$

$$V_{j,t} \in \left\{V^1_{opt,j,t} - m \times \Delta V_{grid}, V^1_{opt,j,t} + m \times \Delta V_{grid}\right\} \quad \textbf{2.53}$$

$$n \in \mathbb{Z}_{>0}, \frac{m}{n} \in \mathbb{Z}_{>0} \quad \textbf{2.54}$$

While the addition of these constraints greatly reduces the required computational time, it is possible that the result of this more constrained optimization will not be the true optimum of the problem given in Eq. 2.49 – Eq. 2.51. However, since only a small deviation from the previously found optimal solution is required in the constraint given in Eq. 2.51, it is expected that the flow profile obtained when adhering to this constraint will be close to the previous results. Furthermore, since the ultimate purpose of this stage in the process is to identify when it is beneficial for downstream facilities to subsidize the revenue of the upstream facilities for increases in flow, small inaccuracies likely would not lead to significant error in the resulting payment schedule. Analysis of the effect of the value of "n" and "m" is given in Section 3.4.

Once this step of the second stage is complete, several new data sets are available for further analysis. The data that is saved for use in the next step includes, for incremented flow during each time period, the revenue lost due to the forced deviation from the optimal flow, and the resulting net flow profile:

$$\Delta R^{1,s+}_{opt,up}, Q^{1,s+}_{opt,up,t} \ \forall \{j, s, t\} \in J_{up} \times T \times T \quad \textbf{2.55}$$

Note that here (and in Eq. 2.56) the subscript "up" refers to the indices of the facilities that are the upstream member in any pair of facilities. The second step of stage two consists of analyzing the effects of the changes in flow rate on the downstream facility in each pair. In this step the only effect that is important is the effect on the revenue of the downstream station receiving the modified flow profile. This also requires an additional optimization analysis for each increase (and decrease) that was analyzed for the upstream facility, though with the constrain in Eq. 2.34 (the volume updating constraint) replaced by the constraint shown below in Eq. 2.56 in order that the modified flow profile will be considered.

$$V_{j,t} = V_{j,t-1} - Q_{j,t-1} + \sum_{\{k \neq up\} \in K_j} \left\{Q_{k,t-1-\tau_{k,j}}\right\} + Q_{in,j,t} + Q^{1,s+}_{opt,up,t-1-\tau_{up,j}} \quad \textbf{2.56}$$

Again, running this optimization twice for each time period in the planning horizon would be computationally very expensive, and thus the optimizations again include the constraints in Eq. 2.52 – Eq. 2.54. The same argument is made for the reasonableness of these constraints as was made above. The resulting changes in revenue of the downstream facilities (Eq. 2.57) are then used in conjunction with the changes in revenue of the upstream facilities (Eq. 2.55) in the final step of this stage of the optimization.

$$\Delta R_{opt,down}^{1,s+} = R_{opt,down}^{1,s+} - R_{opt,down}^{1} \quad \textbf{2.57}$$

The final step is the determination of the water flow payments that, if implemented, would motivate the upstream facilities to increase their release at the appropriate times, thus driving the total system toward more optimal operations. Selecting the proper method of implementing a payment is crucial to ensure that every concerned party will benefit: primarily each individual owner/operator (since their decisions determine the operation of the system), but also the whole system by producing more power when demand is high, and less power when demand is low (as reflected by the power prices). Beyond simply benefiting the revenue of power companies, this is also of benefit to society, likely resulting in decreased CO_2 emissions and lower energy prices because of the increase in the efficiency of producing power when it is most needed using clean hydropower.

When neighboring stations both seek optimal revenue, some tension may inevitably arise. Referring back to the example in Figure 5, it is clear that the operation reflected in Figure 5(b) is optimal for the system as a whole, but whether it is optimal for each station remains questionable. Whereas both stations generate additional revenue as compared to that generated while operating independently (Figure 5(a)), the downstream facility might have paid a smaller amount for the water. On the other hand, the upstream facility might have refused to cooperate unless the downstream facility paid them a majority of the increased revenue. These possibilities lead to a re-examination of what it means to optimize the operation of each individual facility. For the purposes of this work, we define optimal station operation as any operation capable of generating higher revenue than what is accrued under independent operation. Obviously, this leaves a great deal of variability in the decision regarding how the payments will be determined. In a real-world application payment would be negotiated by the lawyers and financial officers of the companies owning neighboring facilities, but since modeling human behavior is far beyond the scope of this work, the decision has been automated via what we hope is an insightful and acceptable way of splitting additional revenue.

The algorithm developed below applies to a pair of neighboring hydropower facilities for the revenue resulting from a required increment in the flow rate. The methodology can, however, easily be extended to include the changes in revenue resulting from a required decrement and any number of pairs of neighboring stations in a system.

First, the total changes in revenue due to each flow adjustment for the pair of neighboring facilities are calculated:

$$\Delta R_{opt}^{1,s+} = \Delta R_{opt,up}^{1,s+} + \Delta R_{opt,down}^{1,s+} \quad \textbf{2.58}$$

Next, the time of the change in flow released from the upstream facility that caused the greatest increase in total revenue is determined.

$$\text{Determine}: \quad s^1 \mid \Delta R_{opt}^{1,s+} = \max_t\left(\Delta R_{opt}^{1,t+}\right) > 0 \qquad \textbf{2.59}$$

This time (s) is the time at which the upstream facility will receive payment for increasing the net flow released (applied during the next iteration of individual optimizations). Note that if there is not time period for which the total revenue of the two stations increases, then this stage of the optimization stops without determining a flow payment to be made. However, it is possible that in some future iteration changes in other parts of the system could result in possible increases in revenue, so the pair cannot yet be removed from further consideration.

The question of appropriate compensation for water released has not yet been addresses. Several options for determining this payment were considered. The first option considered was to pay the upstream facility enough to compensate it for the lost revenue $\left(-\Delta R_{opt,up}^{1,s+}\right)$ plus half of the total revenue increase $\left(\frac{1}{2}\Delta R_{opt}^{1,s+}\right)$ for any release above the previously determined optimum during time period "s". This option was quickly abandoned because of the many iterations that would have likely been required to achieve any meaningful shift in operations. Next, it was considered to pay the upstream facility an amount equal to its lost revenue $\left(-\Delta R_{opt,up}^{1,s+}\right)$ plus half of the total revenue increase for each increment beyond the previously determined optimum $\left(\frac{1}{2}\Delta R_{opt}^{1,s+} \times \frac{Q_{up,s}-Q_{opt,j,s}^{1}}{Q_{grid}}\right)$. However, this option could easily result in a downstream station continuing to pay for an upstream station to increase its flow release far beyond the value that would be optimal for the downstream station, also potentially causing the total revenue of the two stations to decrease, and was therefore discarded. The option finally implemented generates a payment schedule in which each increment of the upstream facility's release beyond the previous optimum results in an additional payment following a geometrically decreasing trend. The total resulting payment is shown below in Eq. 2.60 and Eq. 2.61. Equation 2.60 represents the payment schedule using a summation, and thus the payment for each increment in flow rate can be easily determined, while Eq. 2.61 yields the same result, but with the result of the summation presented in a simpler form.

$$\text{if: } Q_{up,s} > Q_{opt,up,s}^{1}, \text{then: } R_{flow} = -\Delta R_{opt,up}^{1,s+} + \left(\sum_{n=1}^{\frac{Q_{up,s}-Q_{opt,up,s}^{1}}{Q_{grid}}} \left(\frac{1}{2}\right)^{n}\right)\Delta R_{opt}^{1,s+} \qquad \textbf{2.60}$$

$$\text{if: } Q_{up,s} > Q_{opt,up,s}^{1}, \text{then: } R_{flow} = -\Delta R_{opt,up}^{1,s+} + \left(1 - \left(\frac{1}{2}\right)^{\frac{Q_{up,s}-Q_{opt,up,s}^{1}}{Q_{grid}}}\right)\Delta R_{opt}^{1,s+} \qquad \textbf{2.61}$$

The utility of this approach may not be immediately apparent, so the various terms bear explanation. First, the payment only occurs if the upstream facility releases more water than the amount that it previously found to be optimal. Second, since the upstream station was previously optimized, $\Delta R_{opt,up}^{1,s+}$ must be negative. In this equation this term appears to ensure that the upstream station is, at the very least, compensated for its lost revenue. Third, the payment is not linearly related to increased release (in the way that power revenue depends linearly on the amount of power produced). The amount of the payment does depend on how much the flow has increased, but not via a simple multiplicative value function. Finally, the amount of the total increase in revenue paid to the upstream facility depends on how many increments away from its previously determined optimal flow rate it operates. According to the coefficient of the total increase in revenue $\left(\Delta R_{opt}^{1,s+}\right)$, if the upstream station increases its flow release by one grid step, the upstream station receives 50% of the total increase in revenue; if it increases by two grid steps, then the upstream station receives 75% of the total increase that the two stations received for a single grid step increase. The payment schedule is further demonstrated in Figure 6, which shows the payments that would occur for various increases in flow rate based on a loss in revenue of upstream facility equal to $3 and an increase in revenue by the downstream facility of $7, resulting in a total increase in revenue of $4.

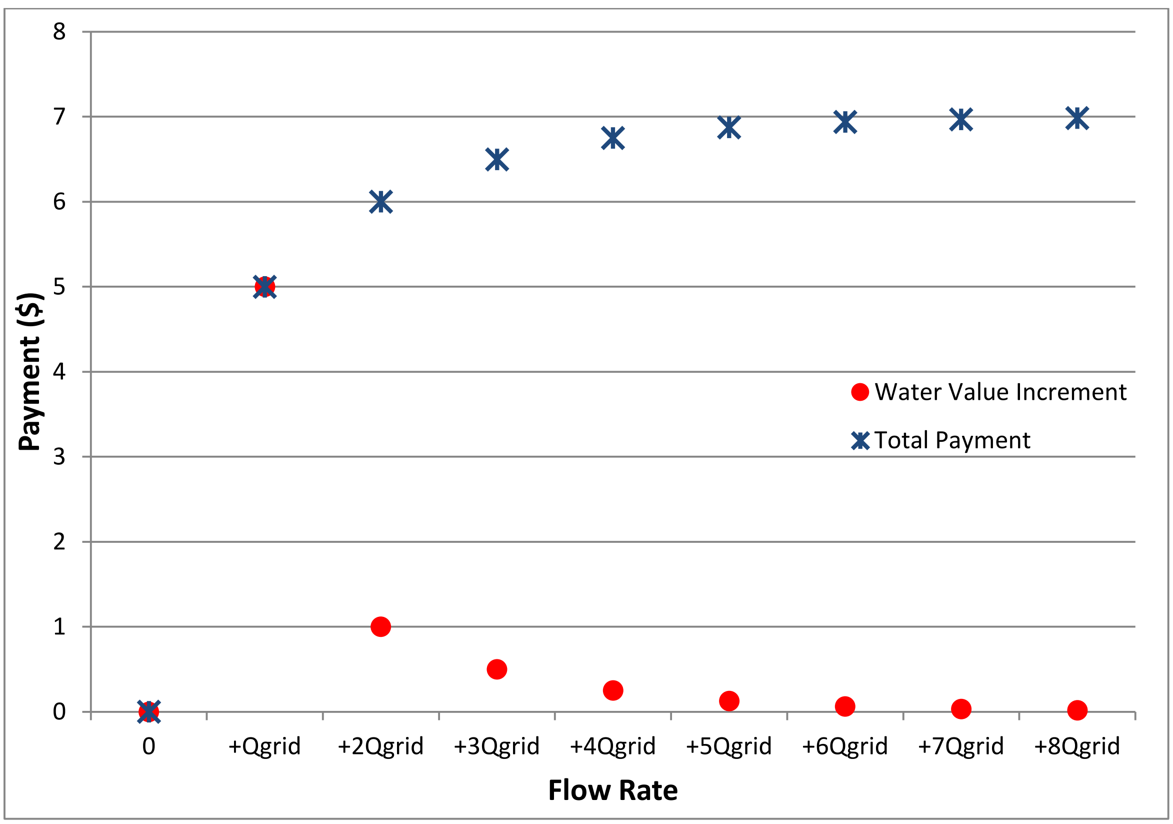

Figure 6: Payments for flow changes of an upstream facility. Payments correspond to upstream facility lost revenue of $3 and total revenue increase of $4 (for a single increment of flow).

The intent of this geometric progression is twofold. Its first purpose is to prevent the downstream station from making a payment beyond the amount its own revenue increased due to the analyzed single increment in the upstream station's flow rate. This is important since there is no guarantee that further increases in the flow released by the upstream facility will be beneficial for the downstream facility. Second, it allows for the optimization to identify directions of steepest ascent and make more than a single step in the identified direction. As shown in Figure 6, the upstream facility can continue to receive higher payments for release beyond the single increment that was analyzed. This is important to reduce the number of iterations (and thus computational time) required to converge on a final solution.

The final result of this stage of the optimization is the payment schedule shown in Eq. 2.61. In order to communicate the proper payment to the next iteration, several pieces of information must be saved. These include the time period for which a new payment has been calculated, whether the payment corresponds to an increment or a decrement in flow (indicated by a sgn taking a value of 1 or -1), the previous optimal flow rate at that time, the revenue loss of the upstream station, and the total increase in revenue of the pair of stations (all listed in Eq. 2.62). This information is then added to a list of payments to be made to the upstream facility and utilized by the updated value function shown in Section 2.5, which details the iterative procedure.

$$\mathrm{PS}_{j=up}^{1} = \left\{s^1, sgn^1, \mathrm{Q}_{opt,up}^{1}, \Delta R_{opt,up}^{1}, \Delta R_{opt,tot}^{1}\right\} \qquad \mathbf{2.62}$$

2.5 Iterative Procedure

The procedure and equations developed above in Sections 2.3 and 2.4 refer specifically to the first iteration of the optimization (superscripts indicate the iteration number). A few slight modifications are necessary for the second and following iterations. Hence, the equations developed above in Section 2.4 all have a superscript of “1” indicating that they result from analysis of the first iteration. In order for the equations to apply to all subsequent iterations of the optimization, this “1” only needs to be replaced by the index “n.” Next, the objective function of the single station optimization procedure of Section 2.3 requires a modification to include the effects of the water payments. This modification is determined using two steps. First, the water payments for each allowable flow rate are determined according to Eq. 2.63:

$$WP_{j,t}^{n}(Q_{j,t}) = \sum_{m=1}^{n-1} \begin{array}{l} \text{if}\Big(\left(s^m = t\right) \wedge \left(sgn^m \times \left(Q_{j,t} - Q_{opt,j,t}^{m}\right) > 0\right)\Big), \\ \text{then}\left(-\Delta R_{opt,up}^{m,s} + \left(1 - \left(\dfrac{1}{2}\right)^{\frac{Q_{j,t} - Q_{opt,j,s^m}^{m}}{Q_{grid}}}\right) \Delta R_{opt}^{m,s}\right), \\ \text{else}\, 0 \end{array} \qquad \mathbf{2.63}$$

This equation results in a matrix of values of flow rate for each facility. The resulting matrices are then included in the optimization objective function previously given in Eq. 2.41, and shown below in Eq. 2.64 with the new water payment term included.

$$R_{opt,j}^{n} = \max\left(R_j = \sum_{t}^{T}\left(\rho_{j,t}P_{reg,\psi,j,t} + \pi_{j,t}P_{\psi,j,t} - SU_{j,t} + WP_{j,t}^{n}\right)\right) \qquad \textbf{2.64}$$

The same modification must also be made to Eq. 2.49 in all further iterations.

Finally, there are a couple of problems that could arise as the iterations progress. The first and most obvious potential problem is that because this is a gradient shooting method, applying a payment schedule may result in the upstream installation releasing more water than is beneficial for the downstream installation, causing the revenue of the downstream installation to decrease after water payments are made. This would indicate that the value applied to increased flow rates beyond the already considered single increment is too large. If the method developed in Section 2.4 were followed, a payment would be made for decreased flow at the same time. However, such a payment would result in the downstream facility first paying for increased flow, and subsequently paying for the flow to be decreased. Obviously, this would not make sense. Fortunately, this problem is easy to rectify. Since it is known that a single increment in flow rate by the upstream facility is beneficial, the base payment that compensates the upstream facility for lost revenue must stay the same. However, the magnitude of payments for further increases in flow rate can be easily reduced by decreasing the value of the recorded total revenue change between the two stations $\left(\Delta R_{opt}^{m}\right)$ by half until the problem is rectified. Reducing this recorded value would cause the geometric series portion of the payment to decrease, in turn causing the incremental benefit of increased flow by the upstream facility to decrease, and thus the optimum release of the upstream facility will be lower at the relevant time.

The second potential problem is that the payment schedule determined in the current iteration n may indicate a payment at the same time as a payment determined by a previous iteration m, but in the opposite direction (a decrease in flow instead of an increase, for example). As was stated above, simultaneously paying for a decrease and an increase is not conducive to a valid solution. The method for dealing with this problem is only slightly more complex than the method for dealing with the previous problem. The first step is the same as for the previous problem, namely, to decrease ΔR_{opt}^{m} by half until the analysis indicates a payment at a different time. However, in some instances that may not be enough, and the flow of the upstream facility will need to be decreased (or increased) further than is possible by simply decreasing ΔR_{opt}^{m}. This possibility bears further explanation. Suppose that a previous iteration m indicated that the flow of a facility should be increased at time s^{m}, and water payments occur when flow at that time is greater than $\mathrm{Q}_{opt,up}^{m}$ (the flow that was determined to be optimal at time s^{m} during stage 1 of iteration m). Suppose further that in the current iteration n it has been determined that at that same time (s^{m}) the flow ought to be decreased, and the first step of decreasing ΔR_{opt}^{m} has been performed until $\mathrm{Q}_{opt,up,s^{m}}^{n} = \mathrm{Q}_{opt,up}^{m} +$

ΔQ_{grid} (i.e. the flow released is only large enough for the upstream facility to receive the smallest possible payment indicated by Eq. 2.60). At this point decreasing ΔR_{opt}^{m} further will have no effect since the facility is receiving a base payment for this flow rate of $\Delta R_{opt,up}^{m}$, which cannot be decreased. In this instance the previous payment must be completely removed, and the analysis continues as before.

The contour plot shown below in Figure 7 gives a two dimensional graphical demonstration of the concepts discussed above. The contours represent constant values of the total revenue of a system consisting of two neighboring hydropower facilities, while the parameters represent flow rates. Individual optimization of the two facilities could result in initial operation on the lower left side of the outermost curve, while system optimal operation occurs near the upper left side. Analysis of the system via a gradient technique would then indicate that the greatest increase in revenue would result from increasing parameter 1, and so a payment would be made to the upstream facility to increase it. Once parameter 1 is increased sufficiently, the greatest rise would instead result from increasing parameter 2, and again, a payment would be made for its increase. The gradient optimization would switch back and forth between increasing parameter 1 and parameter 2 until parameter 2 increased beyond the zero value. At that point parameter 1 would have to start decreasing to approach the global optimum. As discussed above, paying for an increase in a parameter, then subsequently paying for it to be decreased does not make sense. Thus, the payments previously determined to increase parameter 1 would be decreased or reversed until convergence on the final optimal value is achieved.

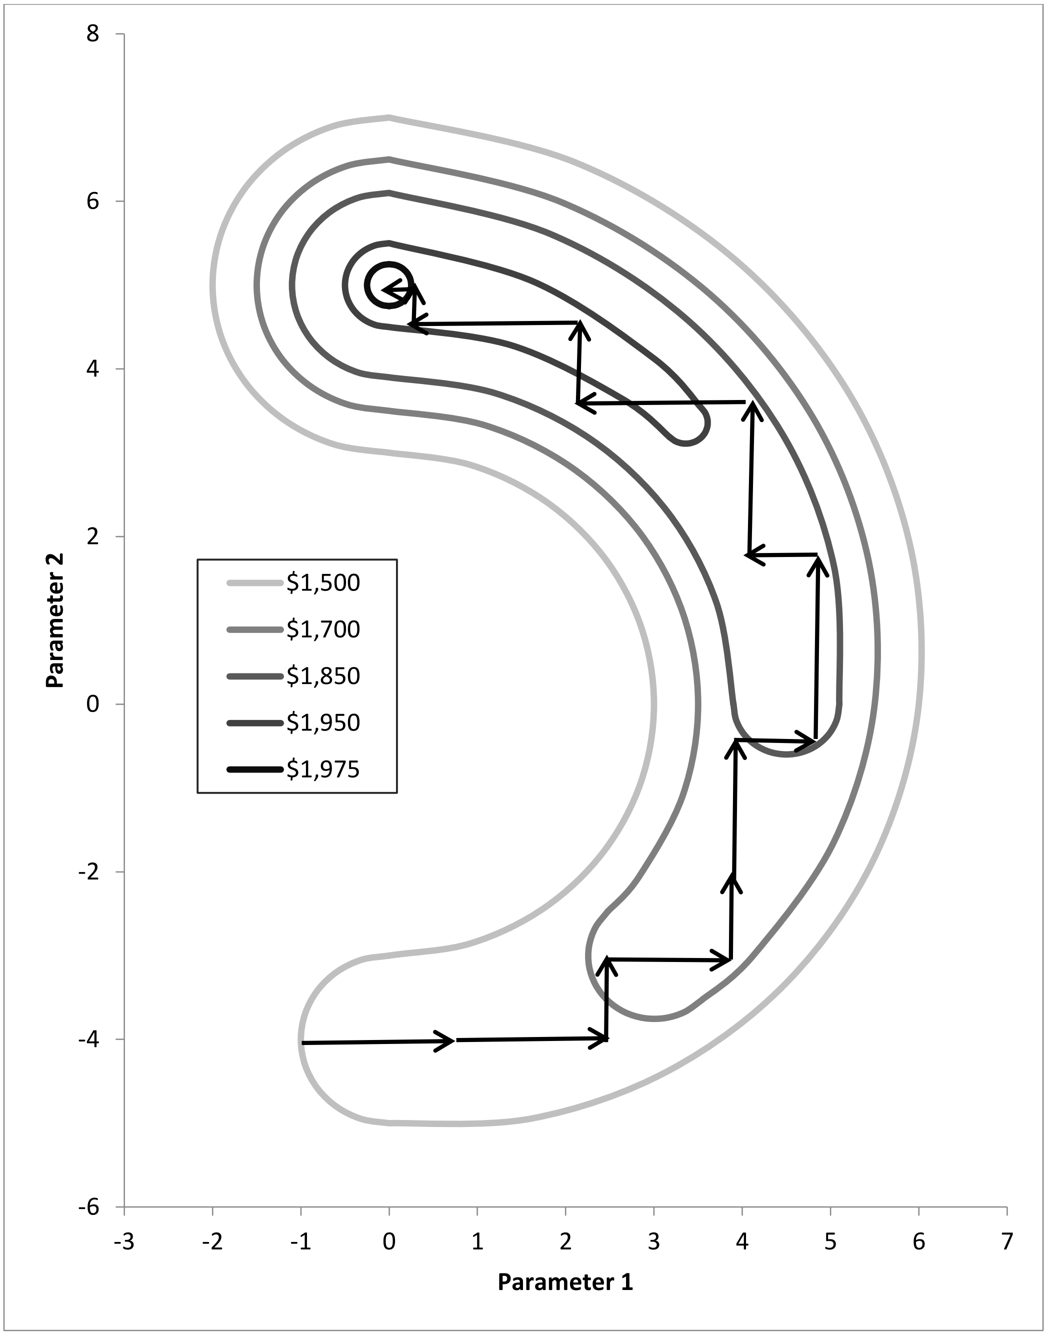

Figure 7: Two dimensional graphical demonstration of the necessity of the ability to reverse or eliminate previously determined payments to change a parameter in a particular direction. Also shown is a possible gradient based optimization path.

3 Results

3.1 Preliminary Analysis

3.1.1 Preprocessing of Data

As the results of the work rest on the foundation formed in the preliminary analysis, the results of those initial steps are shown here. The available data describing the efficiency of a hydro turbine consisted of iso-efficiency curves for varying levels of efficiency (as shown in Figure 1 in Section 2.1.1). In order to apply numerical methods to this problem, power production was needed as a function of evenly discretized values of flow rate and net head. As previously stated in Section 2.1.1, the first step to producing this is creating a triangulation of the points on the head – flow rate plane for which the efficiency is known for use in an interpolation function. Four stages of this process and the resulting interpolation of the efficiency surface are shown in Figure 8 – Figure 15. Note that in each of these figures, the axes are nondimensional head (EwD) and nondimensional flow rate (QwD).

Figure 8 and Figure 9 show the initial triangulation and the resulting interpolation of the efficiency surface. It is immediately apparent that that this is not a good interpolation because of the flat sections half way down the slope to the left of the peak and immediately to the right of the peak.

Examination of the triangulation quickly showed that some of the edges connecting points together spanned across efficiency levels. This observation resulted in the addition of the first constraint to the triangulation, namely, a requirement that neighboring points of the same efficiency be connected. This resulted in the triangulation and interpolation shown in Figure 10 and Figure 11, respectively.

While the addition of the iso-efficiency curve connection constraint greatly improved the triangulation and the smoothness of the resulting efficiency interpolation, the approach is still not yielding a sufficiently orderly triangulation or a sufficiently smooth surface. The next step taken was to normalize the axes with respect to each other, so that the triangulation was performed on points for which the x-axis and y-axis both ranged from 0 to 1. This resulted in the far superior triangulation and interpolation shown in Figure 12 and Figure 13, respectively.

Finally, there were still a few locations in which the interpolation resulted in flat areas on the efficiency surface due to the connection of non-neighboring points of the same efficiency. The locations in which this occurred were identified, and additional constraints were strategically added to cross the offending edges, thus eliminating them. The resulting triangulation, efficiency surface, and nondimensional power surface are shown in Figure 14, Figure 15, and Figure 16, respectively.

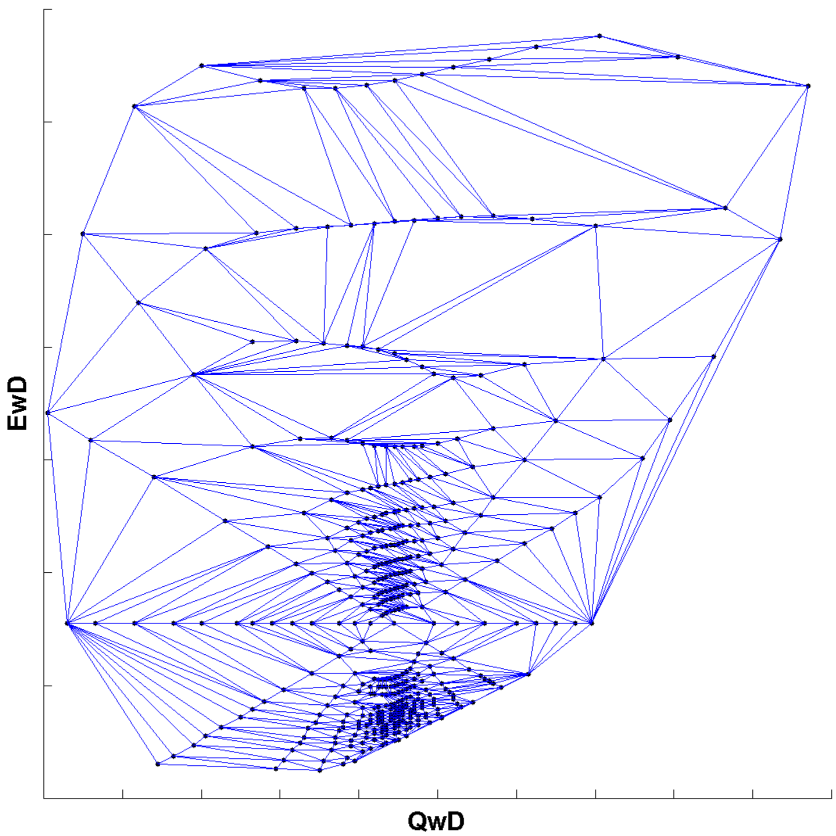

Figure 8: Initial unconstrained triangulation.

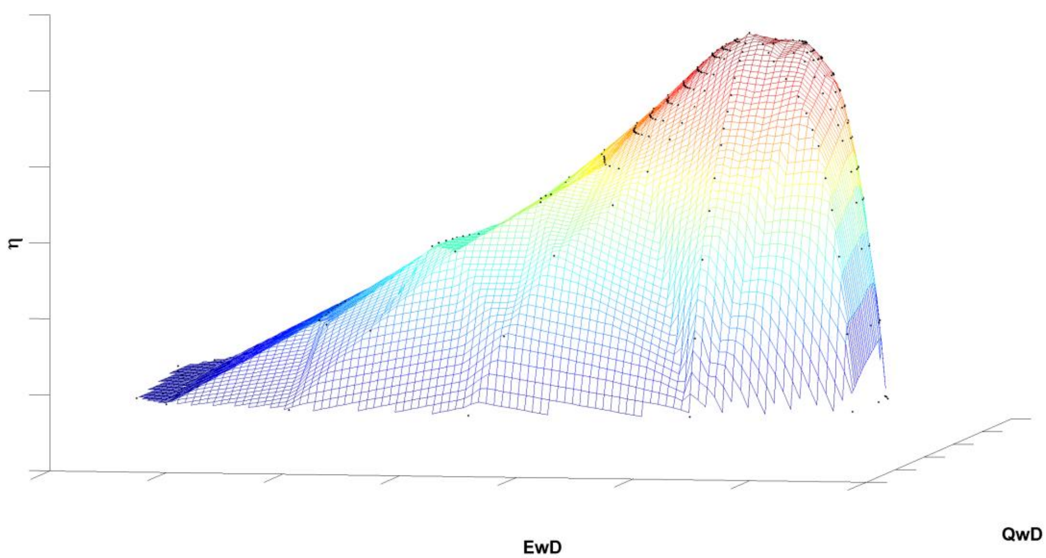

Figure 9: Interpolated surface from Figure 8 triangulation.

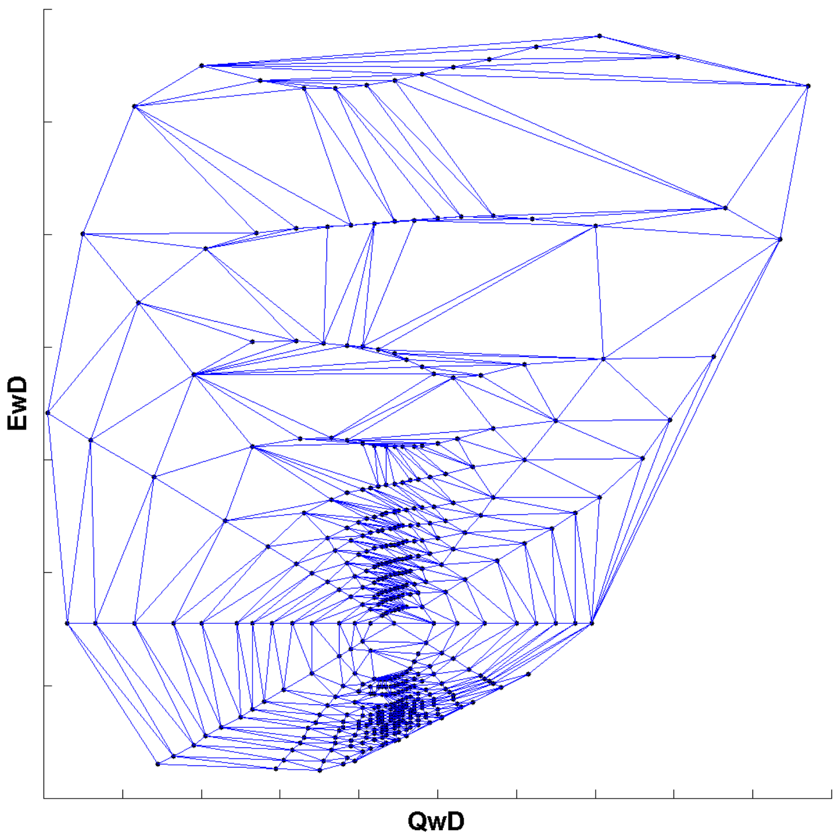

Figure 10: Triangulation with constraints along iso-efficiency curves.

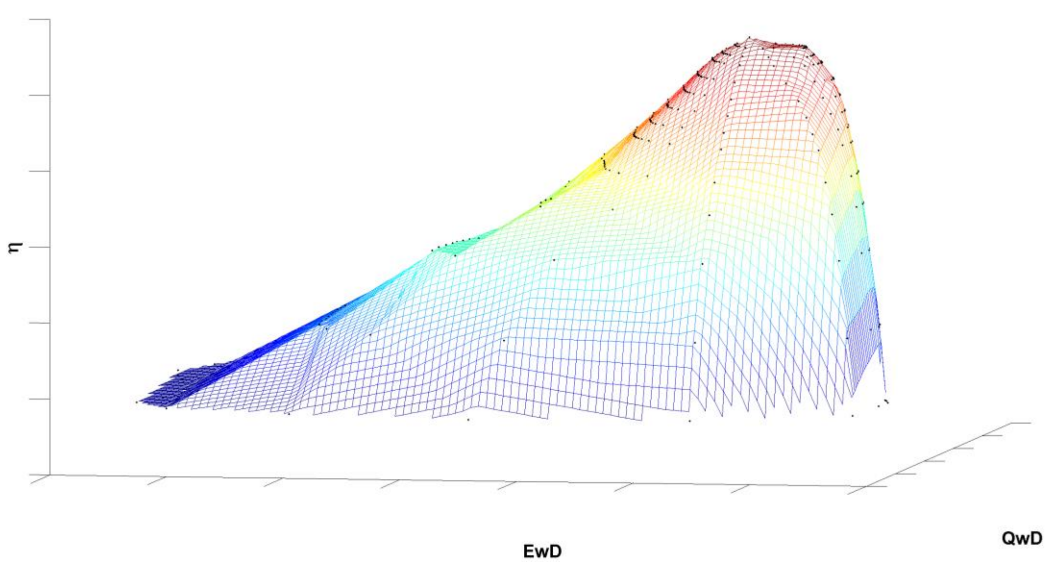

Figure 11: Interpolated surface from Figure 10 triangulation.

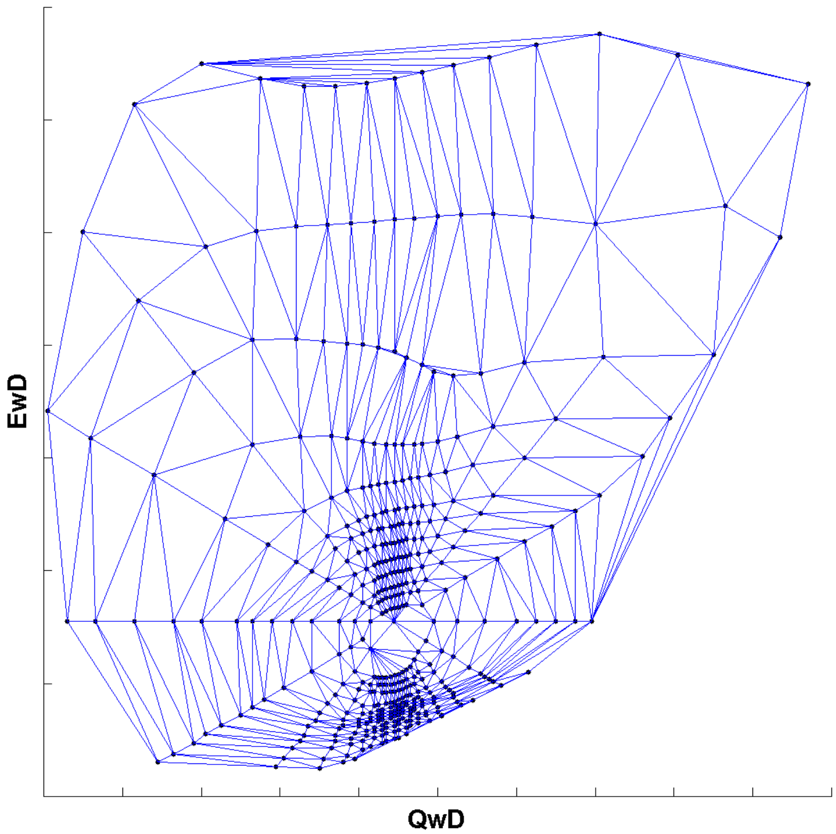

Figure 12: Triangulation with iso-efficiency curve constraints and axes normalized to range from 0 to 1.

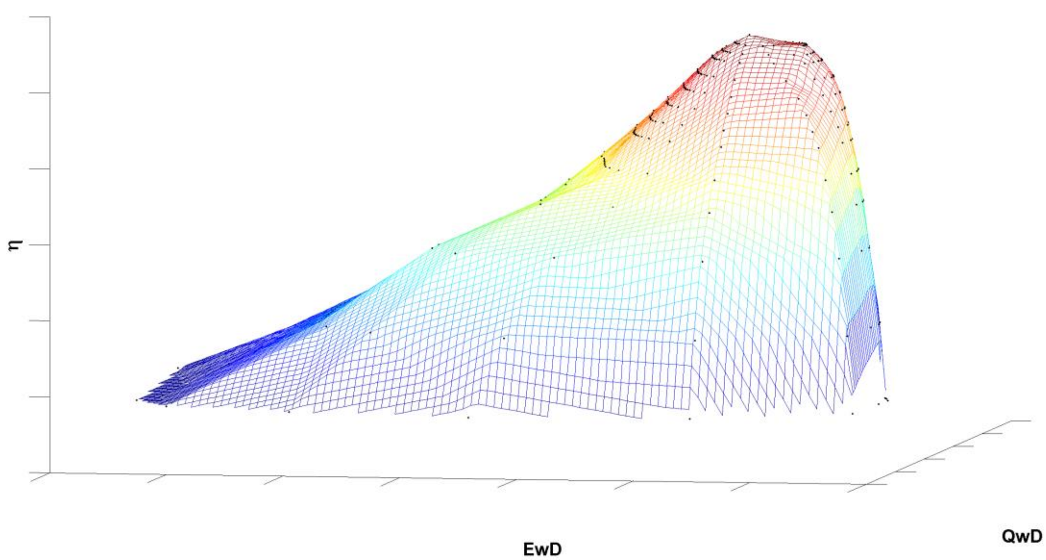

Figure 13: Interpolated surface from Figure 12 triangulation.

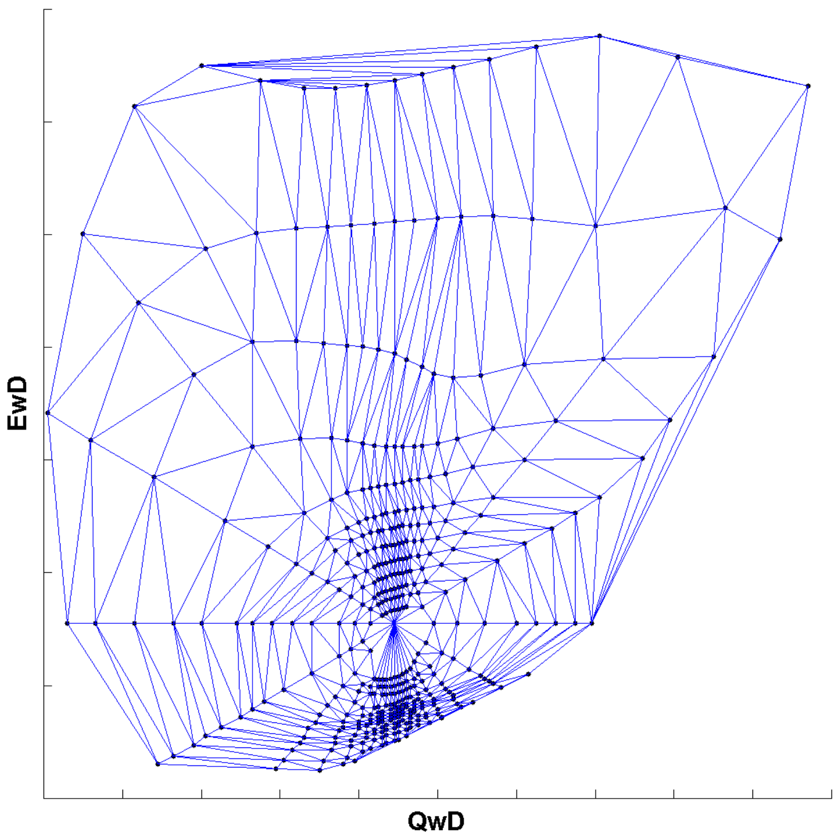

Figure 14: Final triangulation with iso-efficiency curve constraints, normalization, and constraints added as needed to prevent flat areas on the interpolated efficiency surface.

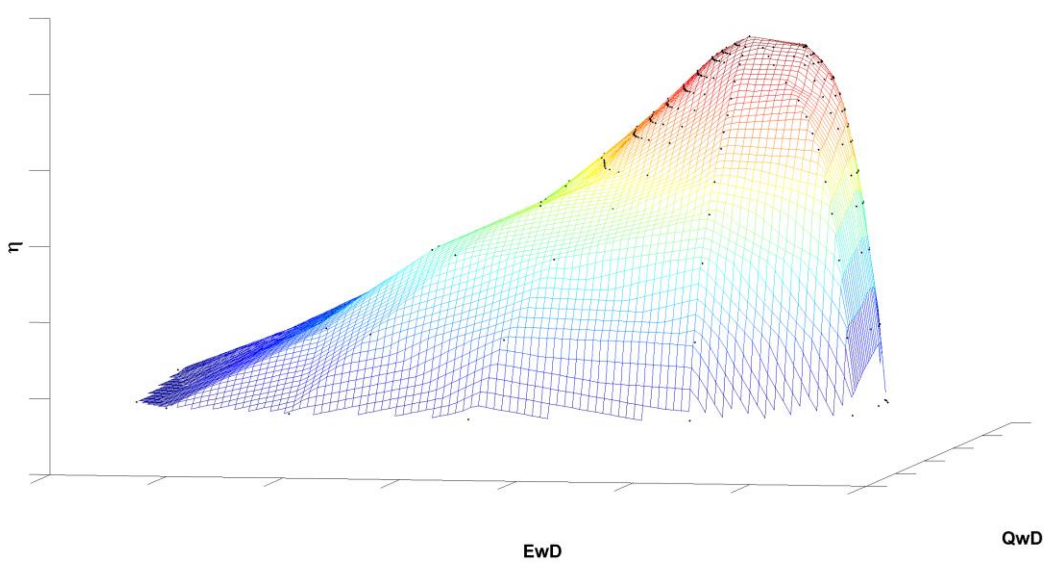

Figure 15: Interpolated surface from Figure 14 triangulation.

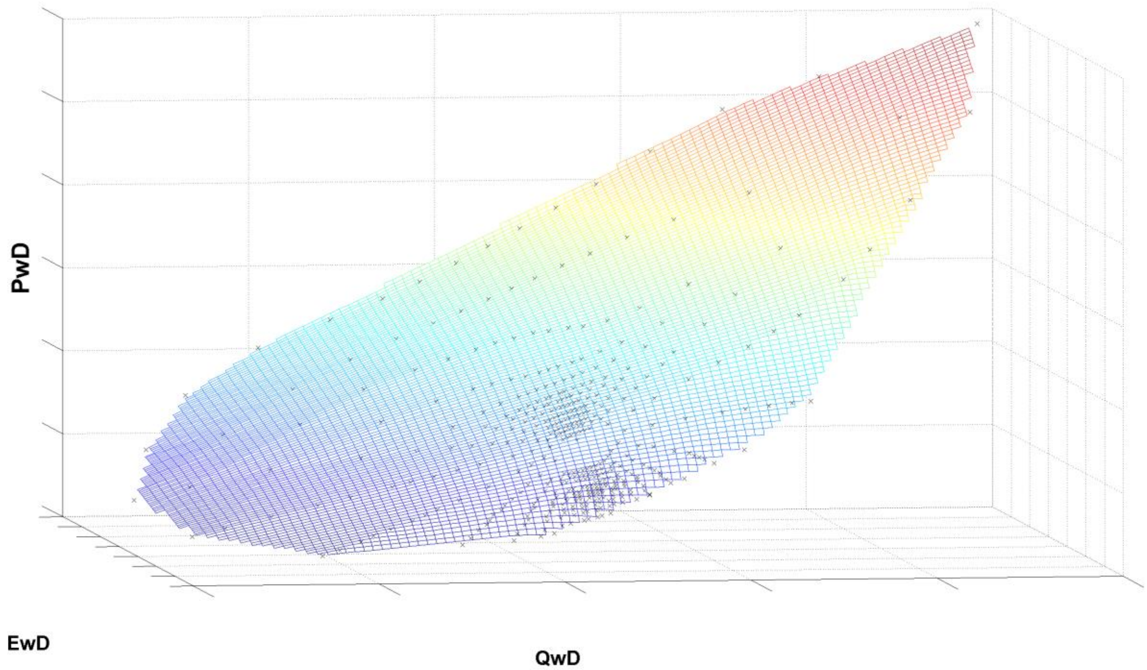

Figure 16: Power surface resulting from the final triangulation shown in Figure 14.

3.1.2 Preliminary Analysis of Unit Combinations

As outlined in Section 2.1.2, the goal of the preliminary analysis of unit combinations was to determine the maximum power surface of each possible combination of operating units as a function of net flow through the hydropower plant and the volume of water held in the reservoir. This combination of known parameters ensures that the operating head can be known, and thus, finding the maximum power only requires determination of the flow rate through each active turbine. The required steps to apply the KKT conditions to solve this portion of the problem include, in order: (1) using the previously determined power surface interpolation functions, extract the power curves as a functions of flow rate for each active turbine at the calculated value of net head, (2) ensuring that those curves are convex so that the KKT conditions are necessary and sufficient, (3) identifying the coefficients of the interpolation functions between the points of the extracted curves, and finally, (4) identifying the flow through each active turbine resulting in the maximum power output for the given net flow rate.

The first step is a straightforward application of the power surface interpolation function that was previously generated to the calculated head and the full range of possible flow rates. The second step immediately analyzes the extracted curve. All data points which cause the curve to be non-convex are first removed, then the center points of any group of collinear points are removed, and finally, the remaining points are shifted vertically to minimize the error between the points interpolated using power surface interpolation function and a linear interpolation of the power using the remaining points. The various stages of this process are shown in Figure 17. The red line shows the power as determined by interpolation using MALAB's built in TriScatteredInterp function and the previously determined triangulation of the available data. This curve is very nearly convex, as expected, but it is not perfectly so, as it must be to directly apply the KKT conditions. Since the curve is reasonably expected to be convex as previously established, all points that cause the curve to be non-convex are removed first, resulting in the set of points identified by the squares. Next, any sets of collinear points further complicate application of the KKT conditions. Thus, any point that is collinear with the points immediately before and after was removed. Again, it was reasonable to do so since the curve is expected to be convex. The points remaining after this step are identified by circles in the figure. Finally, the remaining points are shifted vertically to minimize the difference between the original interpolation and a linear interpolation between the remaining points. The percent difference between the final piecewise linear curve and the original curve is shown by the blue line. As expected, it is, in all cases, quite small, never differing by more than 1.5% from the original interpolation of the curve in all cases studied.

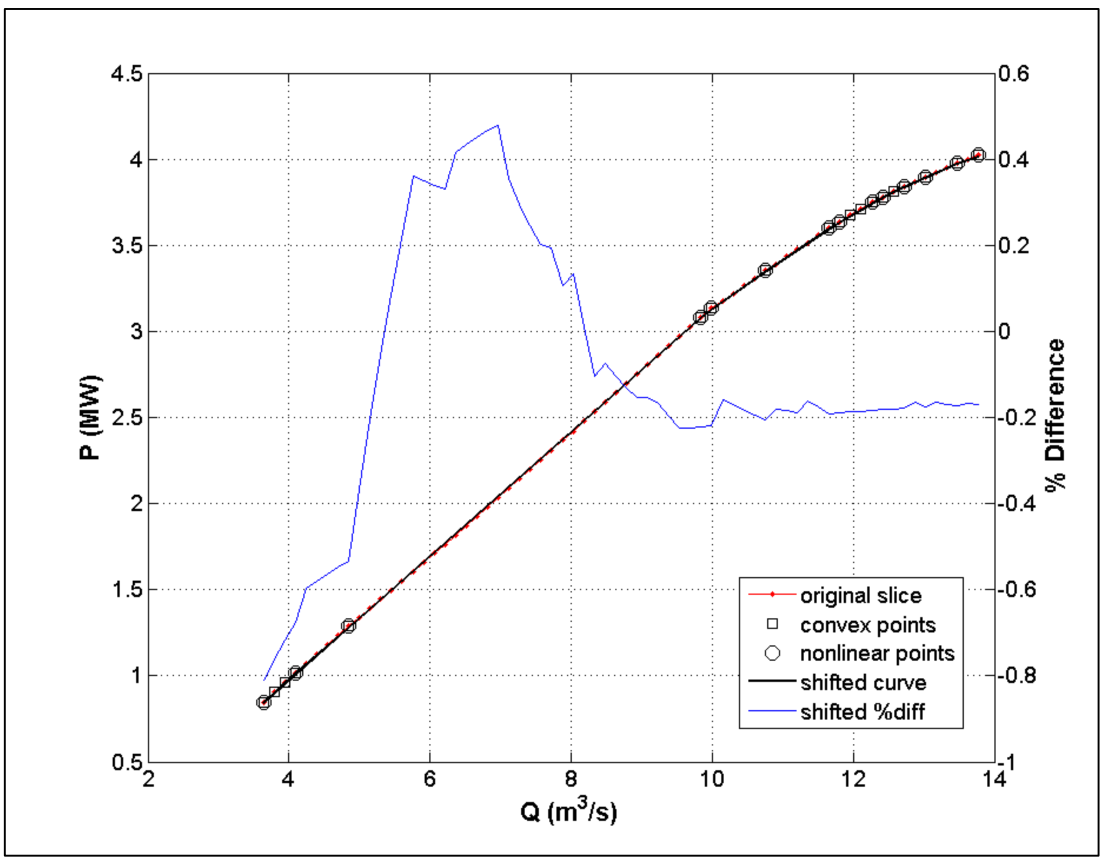

Figure 17: Various stages of the process of extracting the convex power curve for a single turbine at a known value of net head.

These curves were generated for values of head spanning the range of possible values for each turbine. After all of the relevant curves were generated for any particular turbine, the vertical shifts were averaged, and the amount by which each curve was shifted was adjusted to the average. These curves were then used to solve for the optimal division of flow among each set of active turbines for all allowed flow rates. The results for one of each of three differently sized turbines are shown below in Figure 18 – Figure 21. Figure 18 shows three surfaces representing the optimal division of net flow rate (shown on the x-axis) among three active turbines of differing design capacity based upon the net flow through the three turbines and the volume of water currently in the reservoir that feeds them. Figure 19 is very similar to Figure 18, except it shows the power generated by each of the three turbines. Figure 20 shows the total power produced by the combination of three turbines. Note that though the power surfaces of the individual turbines appear rough, the net effect when the power of the individual turbines is added together is that the total power produced by the combination is a very smooth function of the net flow rate and reservoir volume. Finally, Figure 21 shows the net efficiency of the combination when the net flow rate is distributed optimally among the active turbines. Again, note that the efficiency surface is quite smooth.

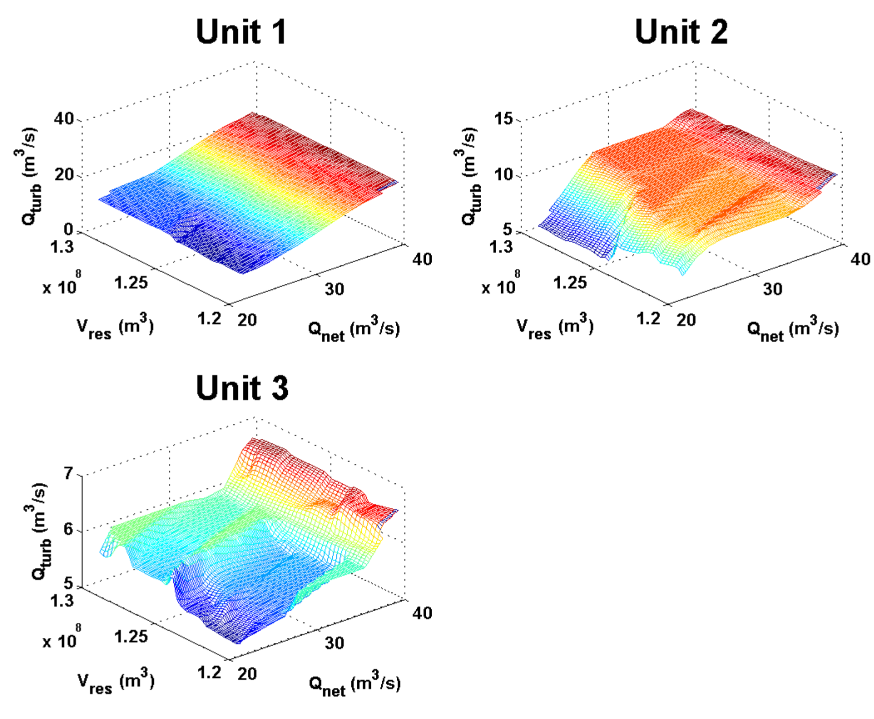

Figure 18: Optimal division of net flow rate among three turbines of differing design capacity.

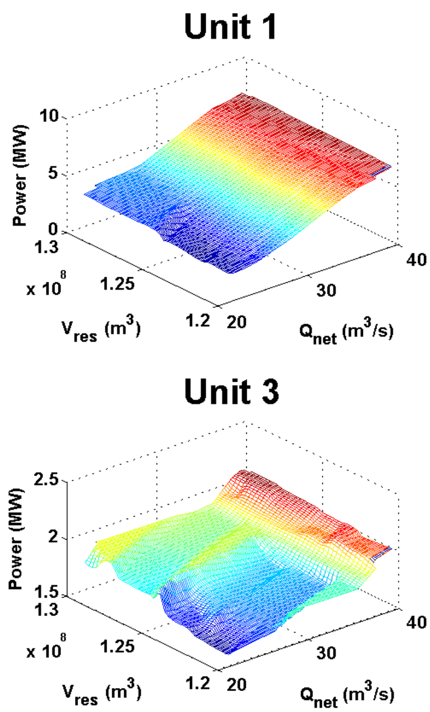

Figure 19: Power production of three turbines when operating at the flow rates shown in Figure 18

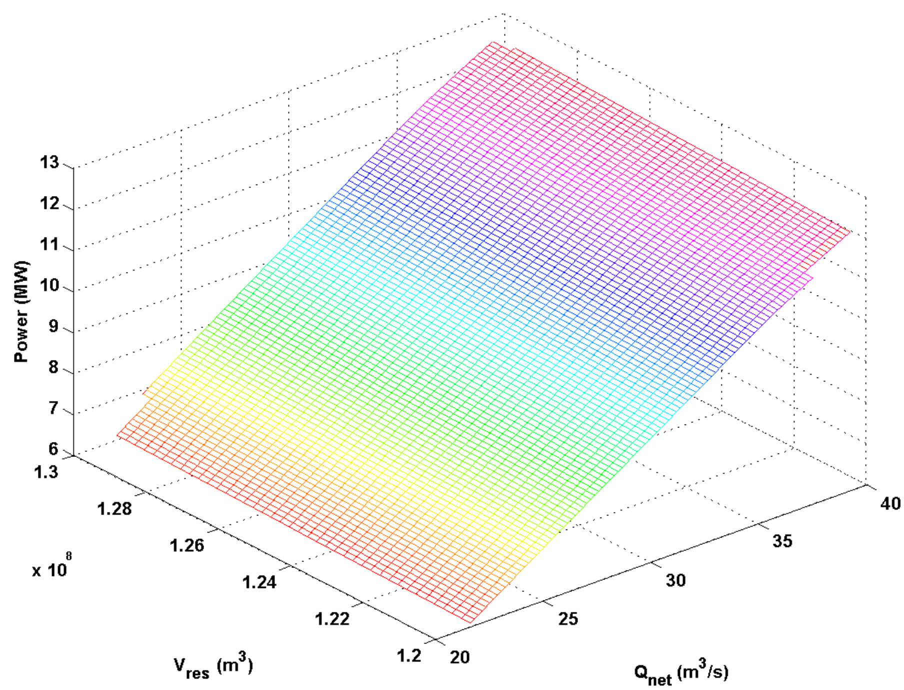

Figure 20: Optimal power of the combination when each turbine is operated as shown in Figure 18.

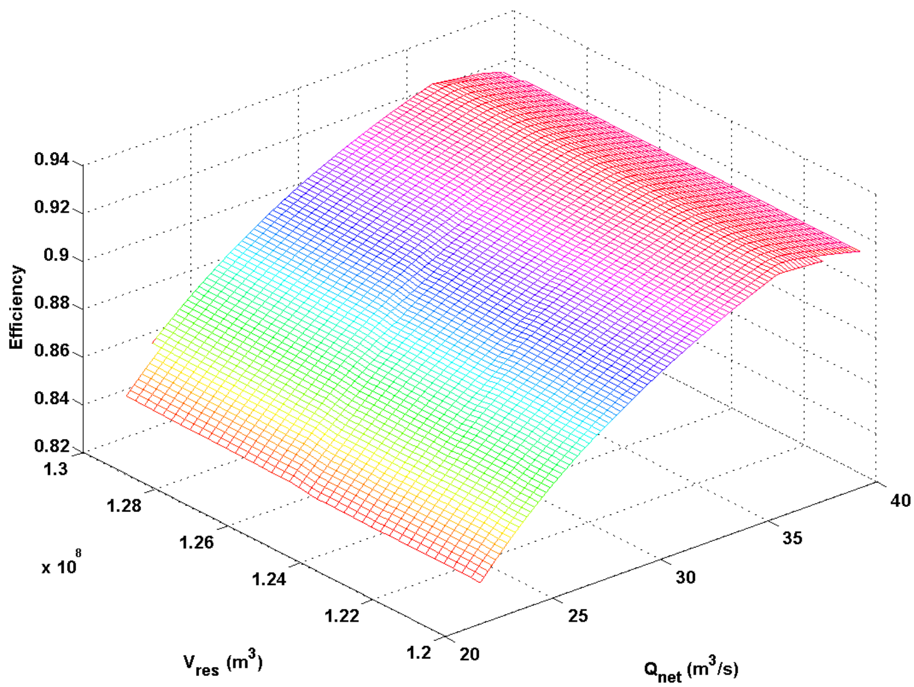

Figure 21: Optimal net efficiency of the combination when each turbine is operated as shown in Figure 18.

Finally, the actual savings in the number of operating states that must be examined during each stage of the DDP analysis bears discussion. The savings for various levels of each of the parameters (number of turbines and discretization size) are shown in Table 1 and Figure 22. Table 1 shows the savings for three levels of each parameter, while the other parameter is held constant. The Discretization Size row indicates the number of discrete flow rates possible for each turbine (not including the value of zero for when the turbine is off). The Before Pre-Analyzing and After Pre-Analyzing rows indicate the number of operating states before and after the preliminary optimization analysis. It becomes immediately apparent that though the savings for a facility with a small number of turbines are not very significant (though still worthwhile), as the number of turbines increases and the discretization is refined, the savings become very significant, with only a tiny fraction of the previous possible number of states remaining to be considered. The data from the table (with the data for n = 3, 4, 5, 7, 8, and 9 added) is plotted in Figure 22 to further demonstrate the savings. For each curve the independent parameter is varied from 2 to 10, while the constant parameter is held at a value of 6. Additionally, curve fits were performed for each of the resulting curves, which indicate that the savings roughly follow a negative exponential with respect to the number of turbines, and a negative power with respect to the discretization size.

Table 1: Computational savings of the preliminary unit combination analysis.

Discretization Size	6	6	6	2	6	10
Number of Turbines	2	6	10	6	6	6
Before Pre-Analyzing	49	117,649	282,475,249	729	117,649	1,771,561
After Pre-Analyzing	24	1,024	26,624	256	1,024	1,792
Percentage Remaining	49%	0.87%	0.0094%	35%	0.87%	0.10%

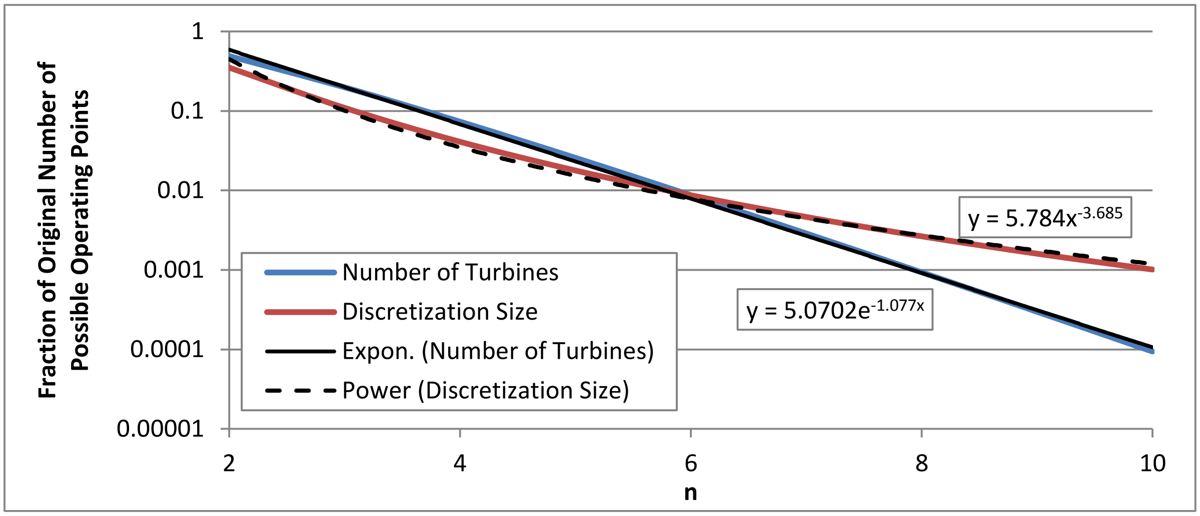

Figure 22: The fraction of operating states that must be analyzed during each single station analysis, as compared to the number of operating states before the preliminary analysis of unit combinations.

3.2 Individual Hydro Station Optimization through DDP

As with the preliminary analyses, the final results depend heavily upon the foundation established by the single station optimization. Thus once the single station DDP optimization was completed, several analyses were performed to ensure that all of the constraints were being met, and that modifying various constraints and inputs in certain ways yielded reasonable results in the final outcome of the optimization.

The constraints analyzed included hourly flow rate change limits, initial and final volume (assumed herein that these are the same), and the grid size used. Input parameters considered included the start-up cost, regulation price, and flow into the reservoir. The variations in the determined optimal inputs and the resulting revenues were then examined for variations of each input parameter, while the other parameters took a median value. The results of each study are shown below.

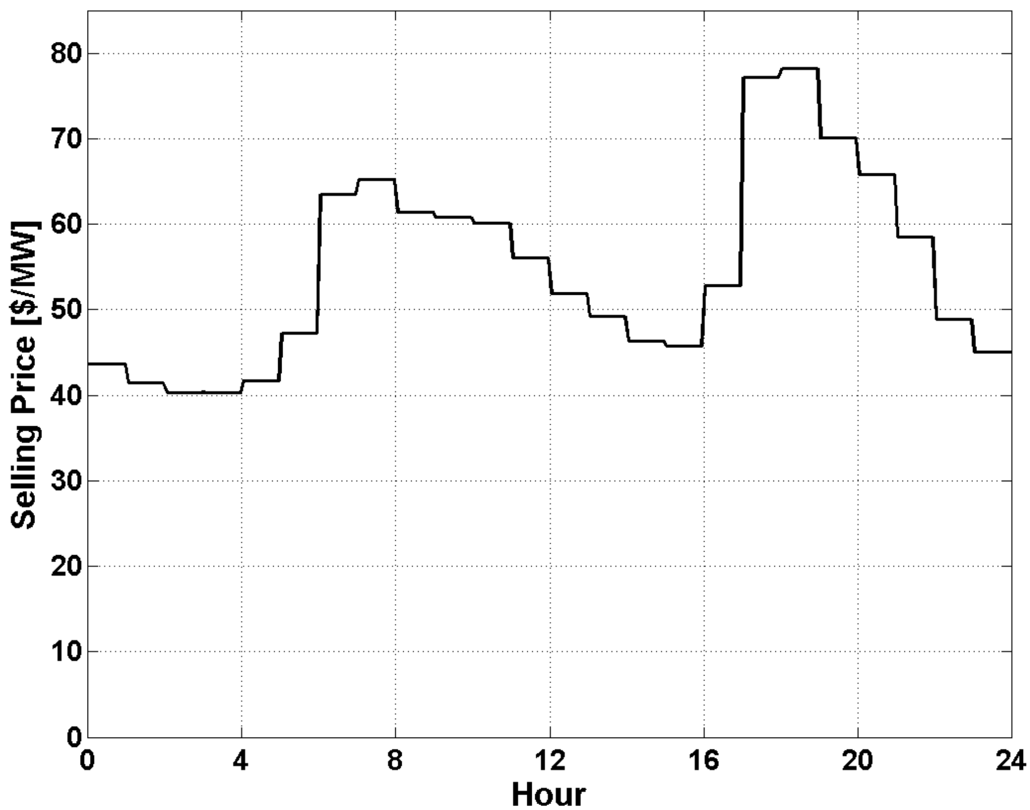

Figure 23: Price profile for power sales on a typical winter day.

Each of the studies was performed using a standard winter price profile as shown above in Figure 23. This winter price profile was chosen because the double peak provides the clearest demonstration of the effects of many of the parameters being studied. Each figure demonstrates the effect of varying a single parameter. The lines indicate the optimal trajectories of flow rate, the numbers indicate the optimal operating sets, and the resulting revenues from each analysis are shown in the legends.

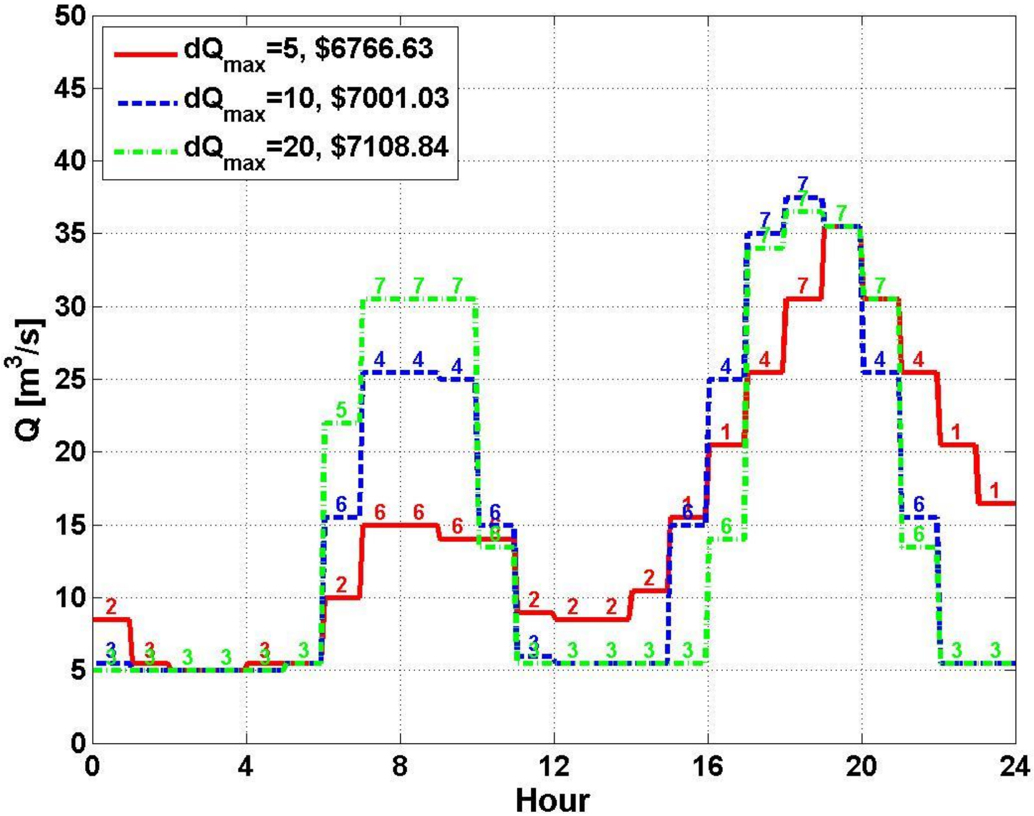

Figure 24: Optimal flow rates and operating sets for various levels of maximum allowed hourly change in flow rate (in m^3⁄s).

The first constraint studied was the maximum allowed hourly change in flow rate. Three levels were studied, and the results are shown above in Figure 24. It can be seen that when allowed, the flow rate released by the station will quickly peak during the times at which the price is the highest. As expected, the revenue is higher when the constraint on the hourly change in flow rate is looser.

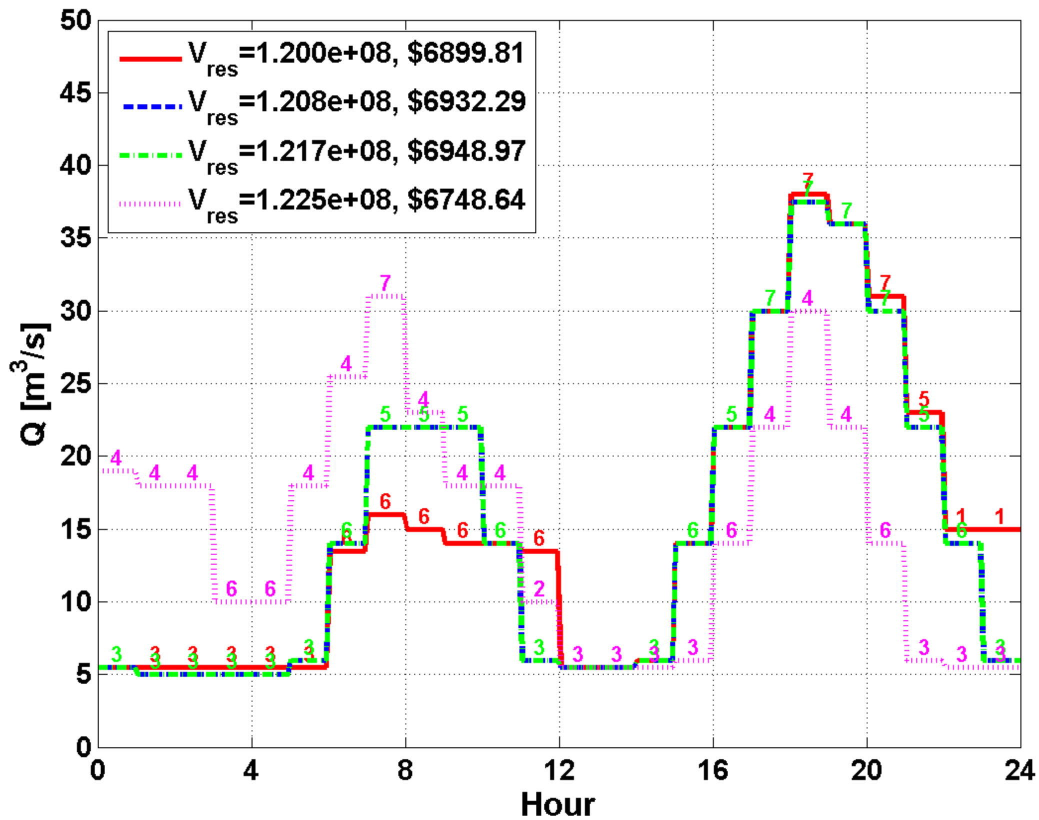

Figure 25: Optimal trajectories for initial reservoir volume levels (in m^3) ranging from minimum to maximum. Final volume is always equal to initial volume.

Figure 25 shows the optimal flow rates for various levels of the initial and final volume of the reservoir. This study confirmed several expectations about the optimization results. First, it confirmed that it is better for the reservoir to be near capacity than for the level to be low (as long as the water level doesn't encounter an upper or lower bound). This is indicated by the higher revenue when the reservoir volume is at the third level ($V = 1.217 \times 10^8\ m^3$) as compared to the revenue when the volume is at the second level ($V = 1.208 \times 10^8\ m^3$). Second, it confirmed the expectation that encountering the upper or lower volume limits causes the resulting optimal revenue to be less than that which would result from operating away from the volume limits. The effect of operating near the volume limits can also be seen in the deviation of optimal flow rate trajectories in the first and fourth cases from the optimal trajectories found in the second and third cases. Finally, the fact the second and third trajectories are identical indicates that precisely tracking the volume of the reservoir may not be necessary for optimization unless an upper or lower limit is encountered. This, however, is beyond the scope of the current work, and would be a topic for future studies.

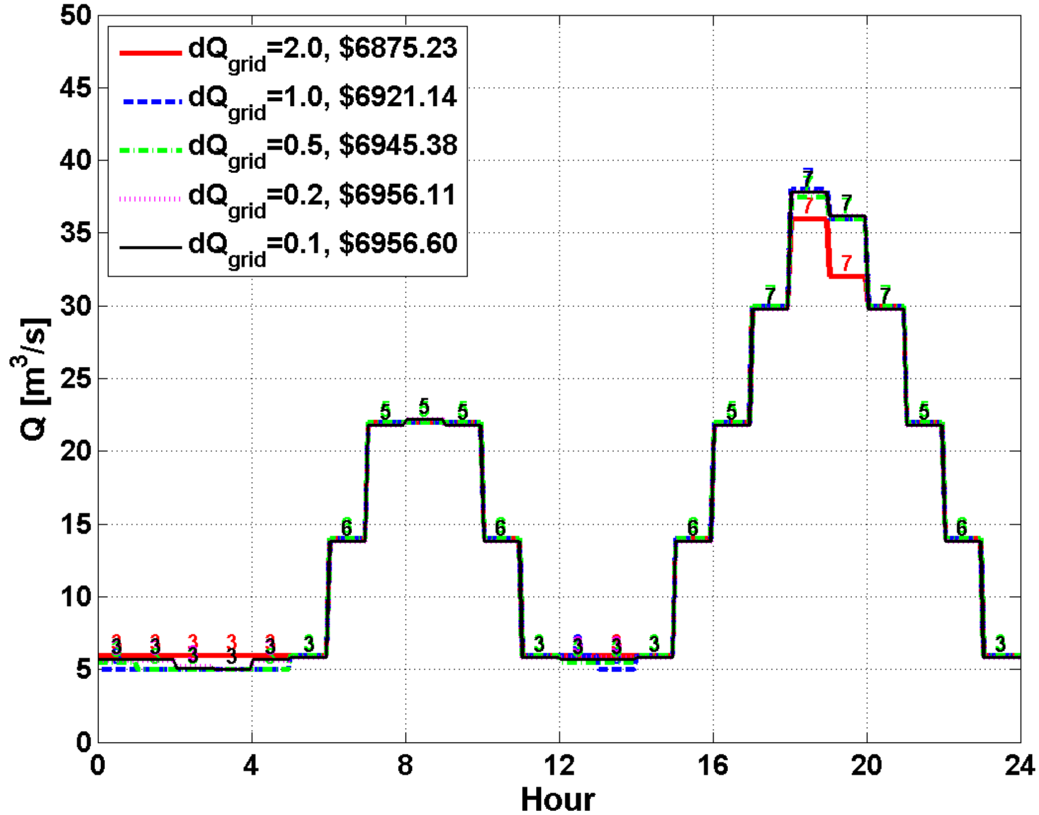

Figure 26: Optimal trajectories resulting from various discretization resolutions (in m^3⁄s).

The results from the variation of the resolution of the discretization shown in Figure 26 and Figure 27 were not entirely as expected. Finer discretizations, of course, always resulted in higher revenue than coarser discretizations. However, what was a bit unexpected was that the flow rate trajectories determined using finer discretizations were not always close to the trajectories obtained from coarser discretizations. Upon further consideration it was realized that this happens for two reasons. First, such cases can result from the presence of two (or more) trajectories that yielded revenues close to the optimal one, as seen in the difference between the red line and all of the other lines in Figure 26. When such a case occurs, the exact values of flow rate considered by each discretization become very important, as a coarse discretization may include points very close to one of the trajectories, but not include values as close to the other trajectory.

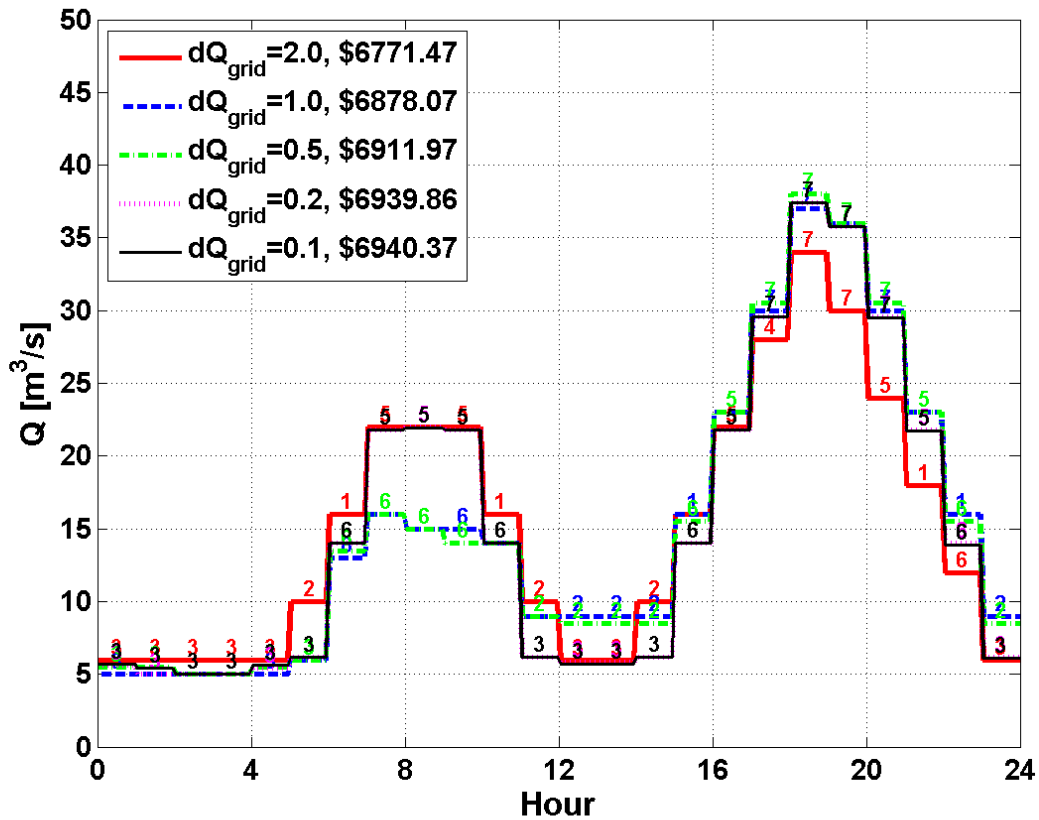

Figure 27: Alternate results for optimal trajectories resulting from various discretization resolutions (in m^3⁄s) obtained using slightly different input parameters.

Another source of the discrepancy is strict adherence to some of the input parameters. In particular, maximum hourly change in flow rate was manipulated to obtain the results shown above. The results shown in Figure 26 were obtained when using a maximum hourly change of $8\frac{m^3}{s}$, whereas the results shown in Figure 27 were obtained when using a maximum hourly change of $7.8\frac{m^3}{s}$. Thus the effective maximum hourly change was always $8\frac{m^3}{s}$ in Figure 26, but the effective maximum hourly changes for the results in Figure 27 were $6\frac{m^3}{s}, 7\frac{m^3}{s}, 7.5\frac{m^3}{s}, 7.8\frac{m^3}{s}$, and $7.8\frac{m^3}{s}$ when the grid size was $2\frac{m^3}{s}, 1\frac{m^3}{s}, 0.5\frac{m^3}{s}, 0.2\frac{m^3}{s}$, and $0.1\frac{m^3}{s}$, respectively.

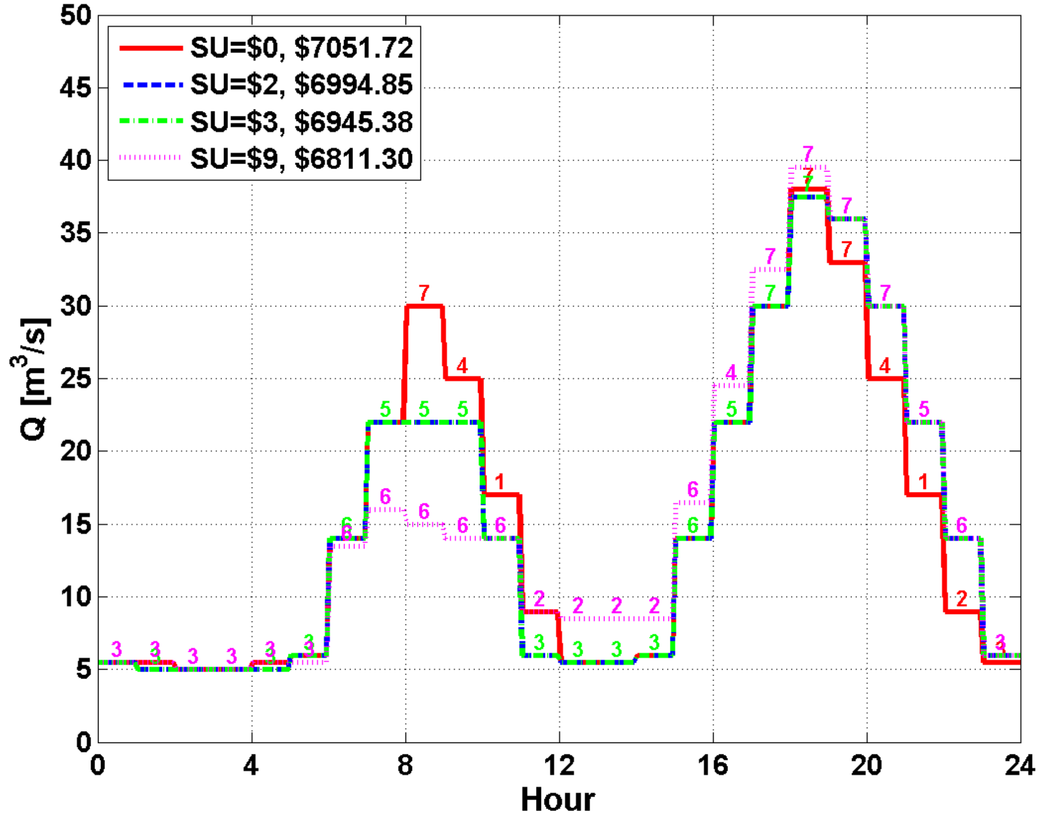

Figure 28: Variations in optimal operations resulting from varying levels of the turbine startup costs.

The start-up cost was a very interesting parameter to study. As expected, higher start-up costs resulted in lower revenues, as seen in the legend of Figure 28. However, the higher revenues obtained when the start-up costs were lower are caused by two factors. First and most obviously, lower costs incurred when starting the turbines correspond to higher incomes. However, that is not the most significant effect of the start-up cost. Incorporating the start-up cost causes the benefits of starting a turbine to be weighed against the cost of turning it on. If starting up a new turbine to either add it to the set of operating turbines or to replace a currently operating turbine cannot result in an increase in revenue at least equal to the start-up cost, then that turbine will be left inactive. This effect is most clearly seen by observing the flow rate and active operating set during the time of the earlier peak in the daily price profile when the start-up cost went from $0 to $2 and from $3 to $9. It is also quite interesting to note that the differences between the resulting optimal flow rates for start-up costs of $2 and $3 (per rated MW of capacity) are very small, indicating that exact determination of the start-up costs may be unnecessary for the determination of the optimal flow rates, so long as the start-up costs are known to be within a certain range.

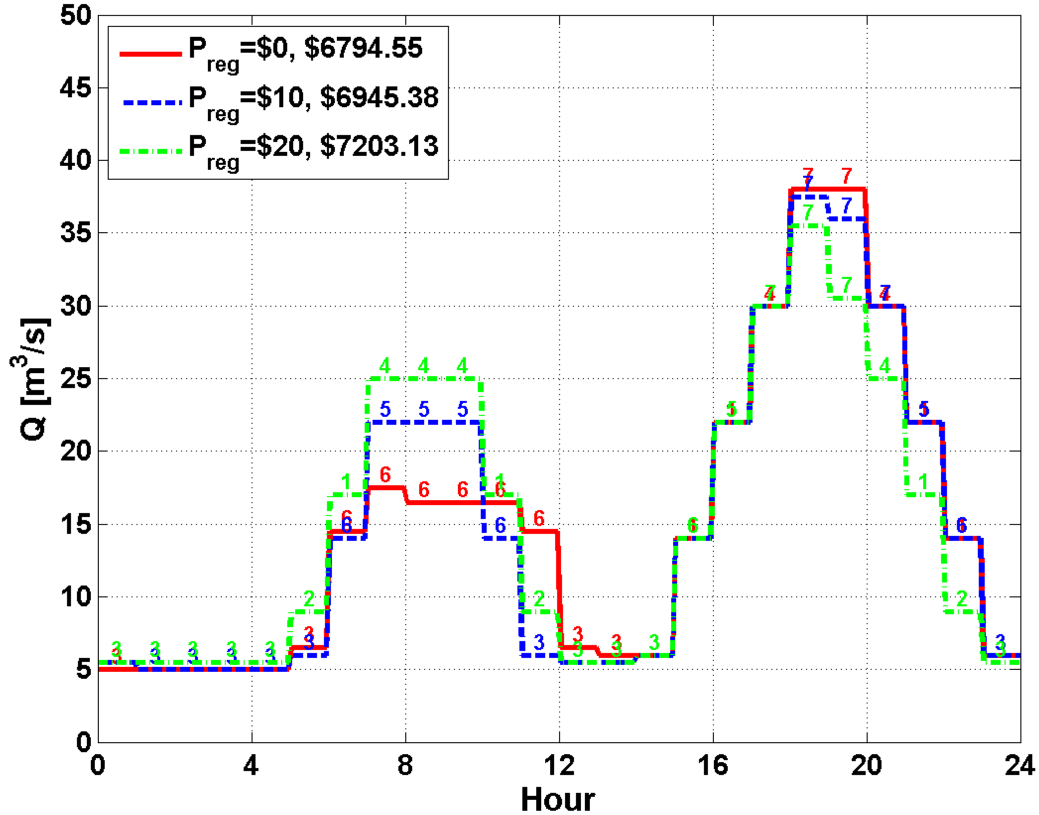

Figure 29: Variations in the optimal operations due to changes in the price of regulation capacity.

The regulation price contains considerable variability, and so this analysis has simply considered scenarios in which regulation has a constant value during each hour of the day. The results shown in Figure 29 indicate two effects. The most obvious result of varying the value of regulation is that when regulation is valued more highly, a higher net revenue results. The second observation relates to the resulting optimal flow rate trajectories. It is seen that it is possible for a higher regulation price to cause operation during times of lower power price to be more beneficial; resulting in the net flow rate being more evenly distributed between the two price peaks.

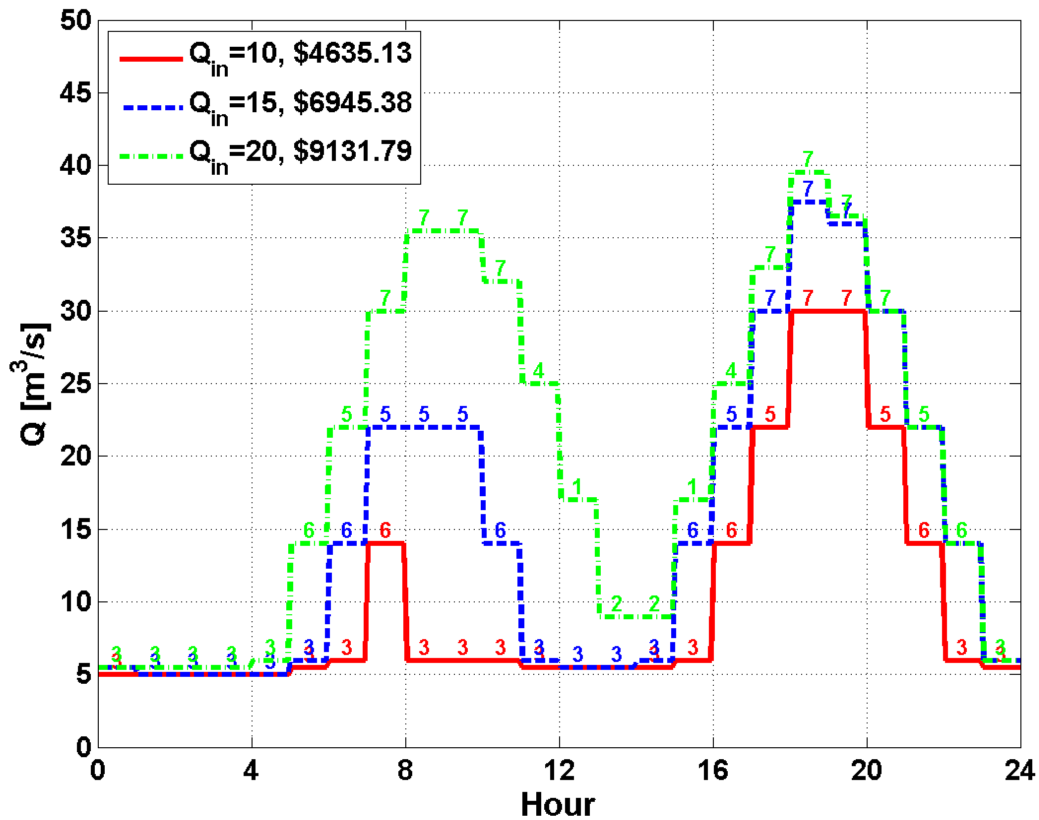

Figure 30: Effect of the hourly flow (in m^3⁄s) into the reservoir on the optimal trajectories.

The final parameter examined was the hourly flow into the reservoir. Again, the results from this study were entirely as expected. Increasing the flow rate into the reservoir required that the net flow through the turbines over the course of the day be higher so that the target final volume could be reached. This, of course, also resulted in higher revenue.

Once the functionality of the DDP optimization of a single station was confirmed, consideration had to be made of how it would be incorporated into the overall optimization. Of primary interest in this regard is the trade-off between the discretization resolution and the optimization time required. Relating to this trade-off, it is interesting to examine the difference between the results when different levels of discretization are used.

The first set of interesting results relates to the optimal revenue resulting from various levels of discretization. Several single station optimizations were performed for a hydropower facility with a large reservoir (so that the reservoir capacity limits would not affect the results) and three turbines of differing sizes. The nominal output power of the three turbines occurs at flow rates of 6, 10, and 20 $\frac{m^3}{s}$. It is assumed that each turbine can operate from 60% to 110% of its nominal flow rate. (Note that all turbines have unique ranges in which they can operate without incurring cavitation damage.

This range is only an approximation of the operating range of a real turbine.) Finally, the optimization is performed for winter operations, during which time the price typically has two peaks each day. The results are shown below in Table 2.

Table 2: Results of grid refinement for DDP optimization.

Optimization #	Flow Rate Grid Size $\left(\frac{m^3}{s}\right)$	Max Difference		Optimal Revenue	% Increase of Revenue	Optimization Time (s)
		Flow	Volume			
1	2.0	N/A	N/A	$6771.47	N/A	0.44
2	1.0	3.5	17.5	$6878.07	1.574%	1.38
3	0.5	1.0	1.5	$6911.97	0.493%	5.85
4	0.2	15.6	41.8	$6939.86	0.404%	58.02
5	0.1	2.0	3.5	$6940.37	0.007%	493.20

Note: The Max Flow Difference and Max Volume Difference columns are comparisons with the results of the previous optimization and have units based on the grid size of the previous optimization. For instance, optimization #4 shows a maximum flow difference of 15.6. The previous grid size of flow rate was $\mathbf{0.5\,\frac{m^3}{s}}$. Thus the maximum difference between the flow rates determined by optimization #3 and #4 was $\mathbf{15.6 \times 0.5\,\frac{m^3}{s} = 7.8\,\frac{m^3}{s}}$.

This table contains two interesting results. First, note that the revenue increase from using a grid size of $0.2\,\frac{m^3}{s}$ to using a grid size of $0.1\,\frac{m^3}{s}$ is very small (0.007%). It is obvious that the point of diminishing returns has been reached, and that further refinement past $0.2\,\frac{m^3}{s}$ is unnecessary for this facility.

Second, observe the large change (15.6 grid steps) in the optimal flow rate between optimizations 3 and 4. This large difference is indicative of a change in the optimal operating set at the time at which the difference arises (and likely at other nearby times as well). This type of change is difficult to predict, and thus renders any attempt to utilize successive refinement to speed up the run time also difficult. However, the run time of nearly one minute required for the grid size of $0.2\,\frac{m^3}{s}$ is far too long for it to be utilized in the iterative procedure in the next stage of the optimization.

In order to attempt to decrease the run time of the single station optimization, a successive refinement strategy was implemented. The refinement strategy utilizes the results of each successive optimization to eliminate from future consideration any flow rate or volume trajectory that deviates more than a set amount from the previously determined optimal trajectory. This can significantly decrease the number of unique states and inputs that must be considered, thus decreasing the run time significantly. The results shown below in Table 3 are for a flow rate limitation of 10x the grid size and a volume limitation of 30x the grid size.

Table 3: Results of grid refinement for DDP optimization with refinement restrictions.

Optimization #		Max Difference				

	Flow Rate Grid Size $\left(\frac{m^3}{s}\right)$	Flow	Volume	Optimal Revenue	% Difference of Revenue	Optimization Time (s)
1	2.0	N/A	N/A	\$6771.47	0%	0.47
2	1.0	3.5	17.5	\$6878.07	0%	1.45
3	0.5	1.0	1.5	\$6911.97	0%	3.04
4	0.2	4.6	12	\$6934.35	-0.079%	7.84
5	0.1	0.5	1.5	\$6935.67	-0.068%	16.64

Note: The % Difference of Revenue column is the difference with respect to the values shown in Table 2.

These results are similar to those in Table 2. Even though the large difference between the optimal flow rates that previously occurred between optimizations 3 and 4 is no longer present, the optimal revenue resulting from the analysis with successive refinement restrictions is within 0.1% of the optimal revenue determined in the full analysis.

It is also of note that the optimization time when successive refinement is used is only slightly difference for the first two analyses. This is because for the first two cases the savings in the optimization time were offset by the increased calculation time required to run multiple optimizations to zero in on the final result. However, the savings in the running time of the program quickly become quite significant as the grid size is refined past $0.5\,\frac{m^3}{s}$.

In order to perform the iterative analysis in a reasonable amount of time, a grid size of 0.5 was used for all applications of the DDP optimization of a single station. Although this grid size does not provide results quite as good as might be desired, it provides to best balance of required run time and quality of results of the grid sizes studied.

3.3 River System Optimization through Iterative DDP and Gradient Analysis

3.3.1 Optimization Results

The results of the analysis of a system consisting of three hydro plants along the same river are presented and discussed here. The river flow time is three hours between each pair of neighboring plants. The first one (farthest upstream) had a large reservoir, while the reservoirs of the second and third facilities were small enough that they could be considered to be run-of-river installations. Each plant has three turbines of varying sizes, and each one operates at approximately the same head.

The predicted price profile used in the analysis was a typical summer price profile in order to most clearly show the shifting of the expected single peak in flow rate. This price profile is shown below in Figure 31.

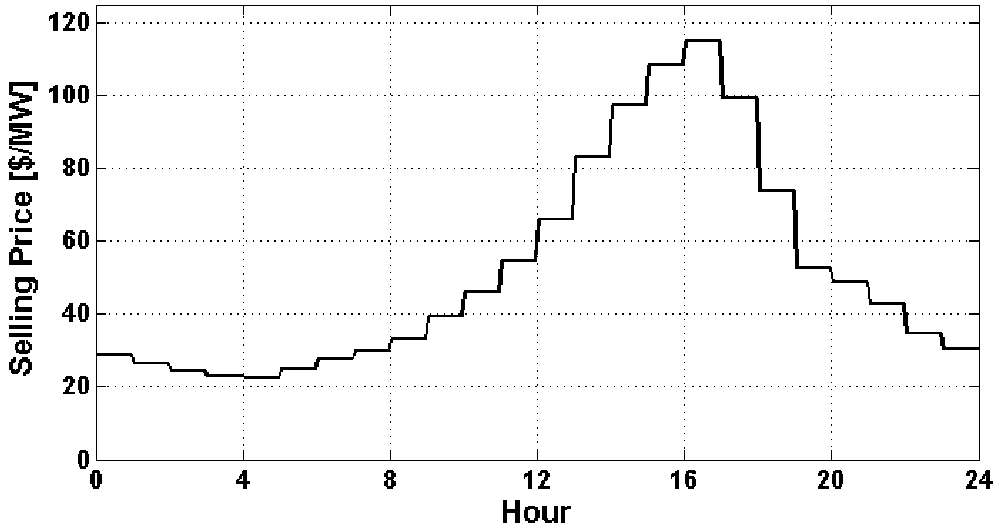

Figure 31: Typical summer price profile used in system optimization analysis.

The initial and final flow rate profiles for each facility are shown below in a set of three figures, namely Figure 32, Figure 34, and Figure 36, while the group composed of Figure 33, Figure 35, and Figure 37 shows the hourly distribution of power revenue, water revenue, and water payments. Finally, the progression of the iterative process is shown for each facility in yet another set of similar graphs, Figure 38, Figure 39, and Figure 40. Several observations and comments can be made regarding the general results as well as those depicted in each of these three groups of figures.

First, the analysis has succeeded in increasing the expected revenue for each facility in the system, and thus also for the system as a whole, using only communication of water values between stations. From upstream to downstream, the facility revenues were increased by 2.8%, 10.0%, and 21.1%, for a net increase in revenue for the system of 10.1%. Further details regarding the final revenues of the facilities are given below in Table 4.

Second, the water payments indicated are for both increases and decreases in flow. Whether payment at any particular time is for an increase in flow or a decrease in flow can only be determined by examining the payment schedule, though in general payments made before the peak in the price are generally for increased release, and payments made after the price peak are generally for decreased flow. This is as expected because before the plants begin communicating water values, the first installation releases the most water during the peak price time, delivering that water to the run-of river plants after the peak price time. Thus the run-of river facilities desire more water at earlier times and less water at later times. It is also certain that at any particular time payment is only made for either increases or decreases in flow due to the methodology discussed in the last paragraph of Section 2.5.

Another observation is that in the final results the temporal flow profile of the middle and last facilities resemble the initial profiles of the first and middle facilities, respectively. This meets expectations. Since each facility operates at approximately the same head, highest total system revenue would be expected to result from a distribution of flow such that as much water is released around the peak price time as possible. Thus, for an optimally operated system, the middle reservoir should be releasing the most water around the time of the peak price, while the upstream and downstream facilities should be releasing the most a few hours earlier and a few hours later, respectively. This is achieved by a chain reaction effect in which water value at the last plant are communicated to the middle plant, which in turn communicates its water value (which is influenced by the last facility) to the first installation, which controls the propagation of water throughout the whole system.

Finally, the water transactions make up a relatively small, though significant portion of the final revenue of each facility. In this example the total value of the water transactions (summed over the course of the day) tended to be approximately 5% – 8% of the final total revenues. Strategically, located, these payments were sufficient to increase the system revenue by more than 10% of the value obtained by optimization of the three plants without coordination.

Table 4: Final breakdown of revenue sources for each facility in the system.

	Facility 1	Facility 2	Facility 3
Water Payment	$ -	$ 645.64	$ 423.36
Water Revenue	$ 645.64	$ 423.36	$ -
Power Revenue	$ 8,088.22	$ 8,365.42	$ 7,268.07
Total Revenue	$ 8,733.87	$ 8,143.14	$ 6,844.70

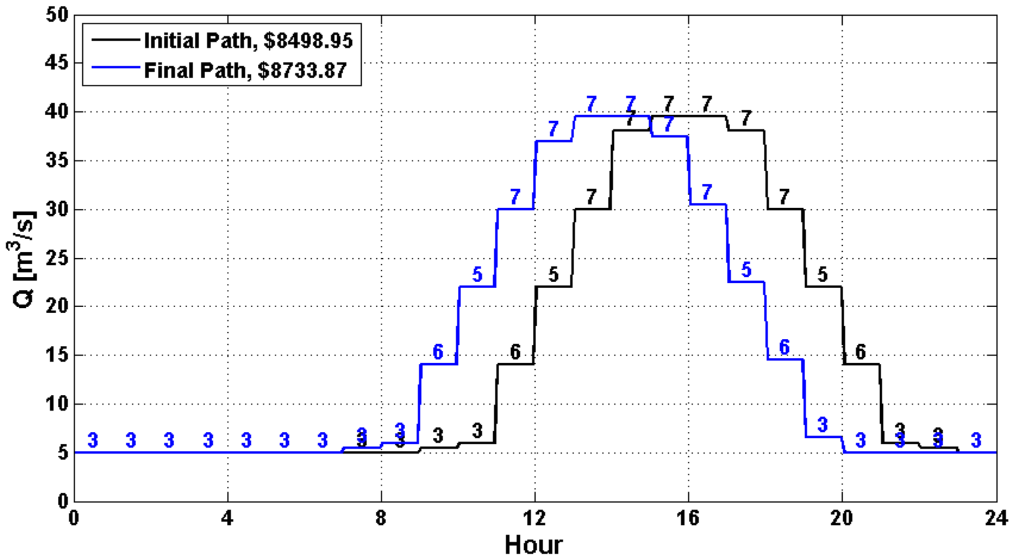

Figure 32: Initial and final flow rate trajectory of farthest upstream facility with a large reservoir.

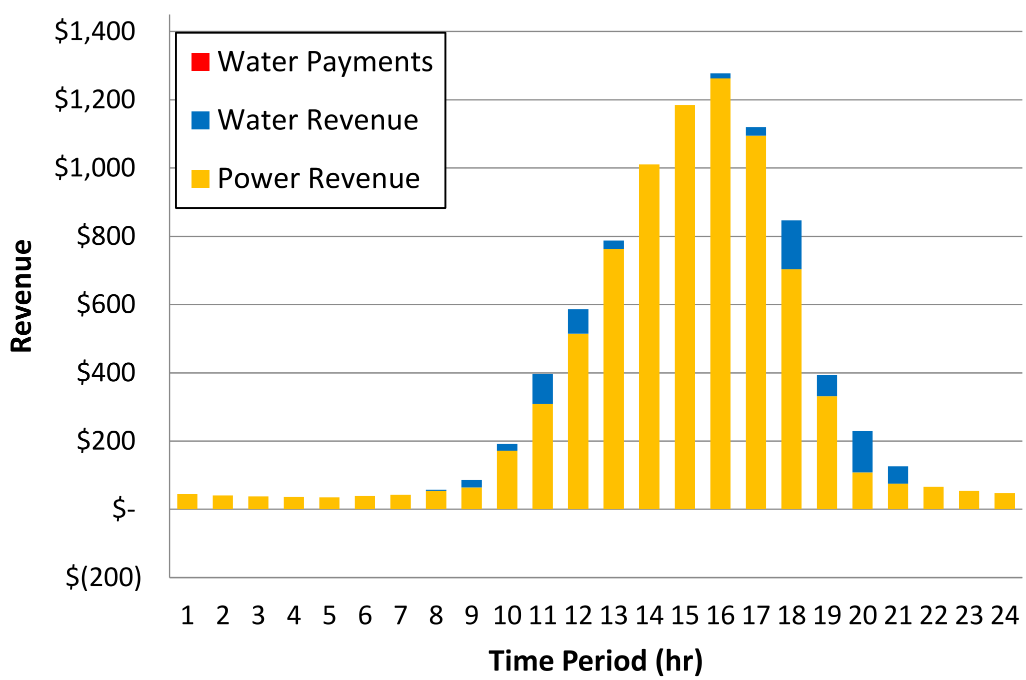

Figure 33: Hourly revenue distribution of farthest upstream facility.

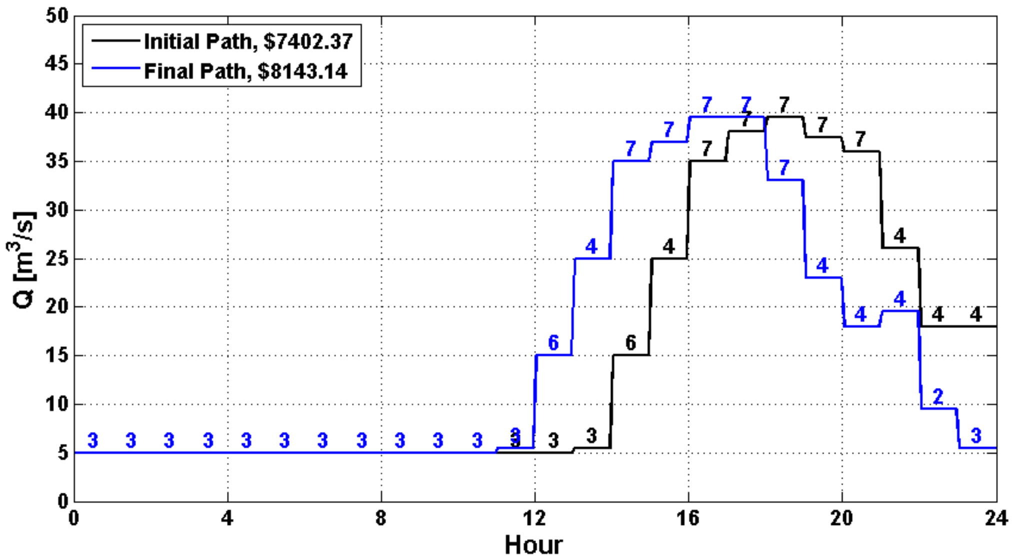

Figure 34: Initial and final flow rate trajectory of middle run-of-river facility.

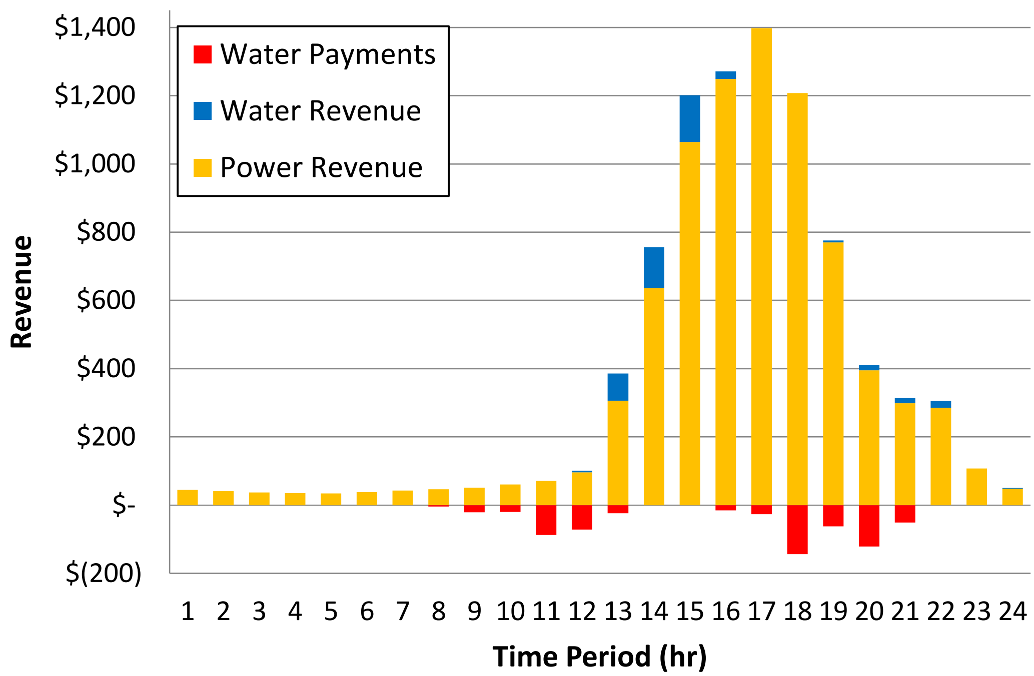

Figure 35: Hourly revenue distribution of middle facility.

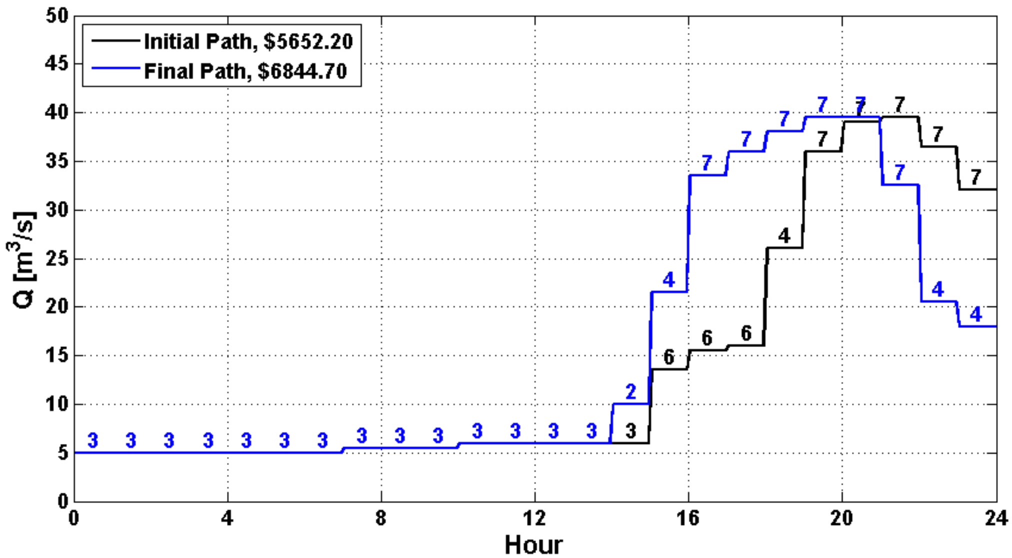

Figure 36: Initial and final flow rate trajectory of farthest downstream run-of-river facility.

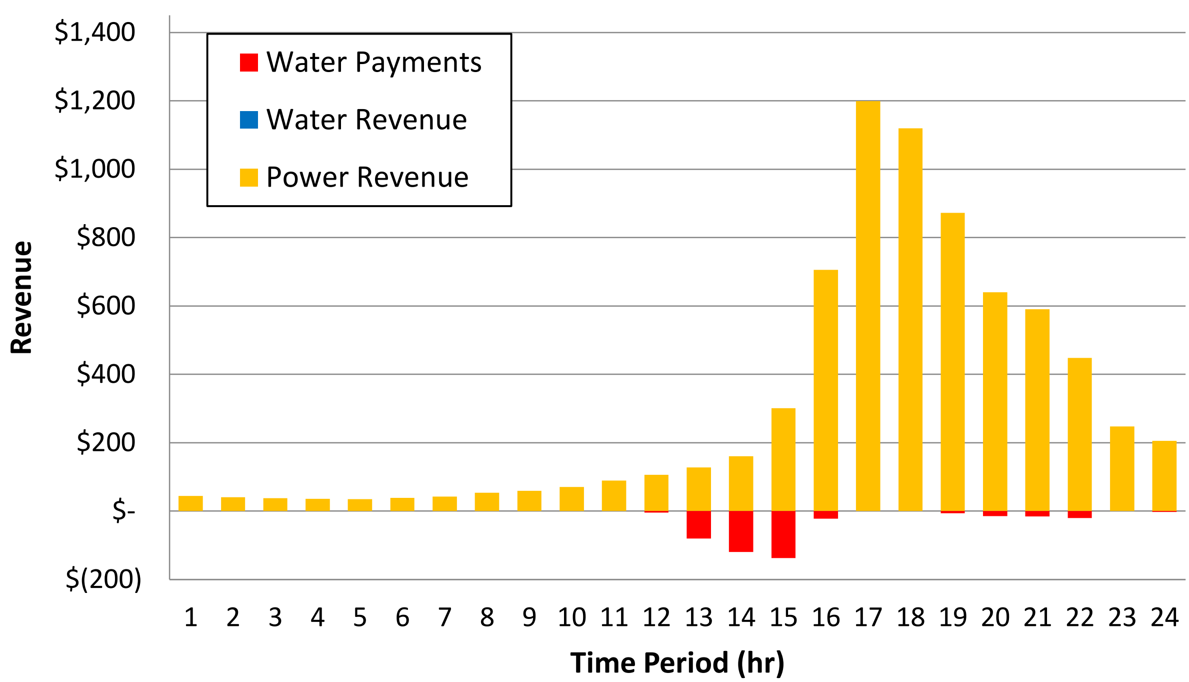

Figure 37: Hourly revenue distribution of farthest downstream facility.

It is also interesting to consider the progression of the iterative method from the first iteration to the final one. The first item of note is that not every step of the iteration is guaranteed to result in increasing the revenue of each facility. At times the revenue can fail to increase because of the complexity of the interactions between neighboring plants. For instance, a change in the flow released by the first one due to a payment by the second one could potentially result in unexpected changes to the flow into the third facility, causing a temporary decline in its revenue. However, such events are a large part of the reason the process needs to be iterative, allowing for corrections to arise in later iterations. This type of interaction cannot, however, account for the large decline in the revenue of the first facility that begins after the 33rd iteration. This decline results due to a completely different reason. In this particular instance the determination had been made in a previous iteration that increased release by the first facility during the 9th hour of the day had a very high value, prompting this installation to increase its release during that time period. However, the 33rd iteration determined that the flow should be decreased (from its now higher value) during the 9th hour. As was detailed in the last paragraph of Section 2.5, paying for a change in flow in one direction, then subsequently paying for the change to be reversed is not reasonable, so the prior payment for the flow to increase during the 9th hour was decreased. The previously determined value of increased release was quite high, so it took several iterations for it to decrease sufficiently.

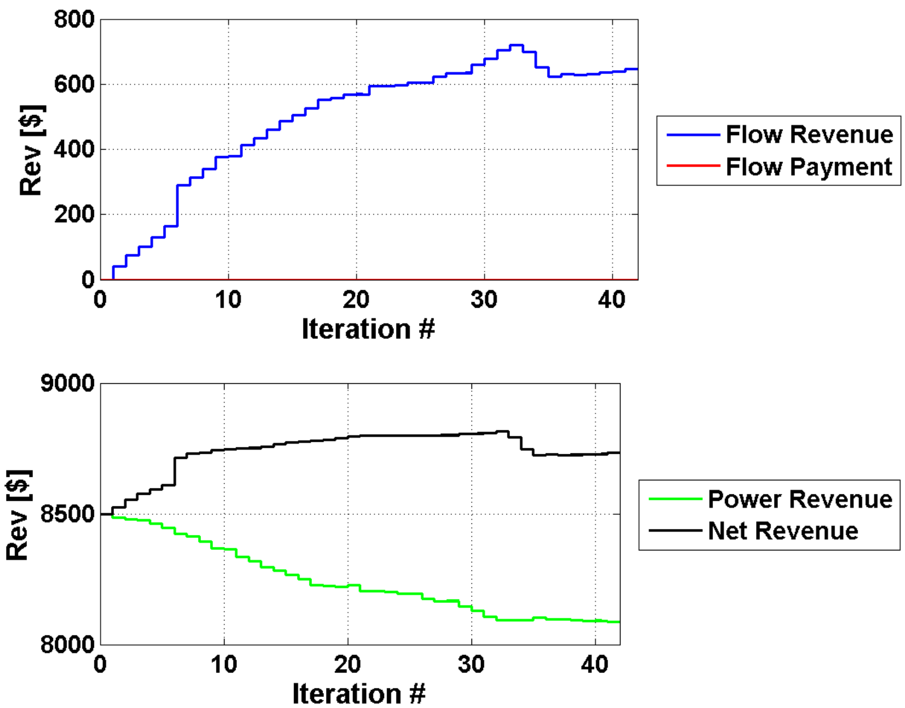

Figure 38: Revenue progression of first facility. No upstream facility, so flow payment is zero.

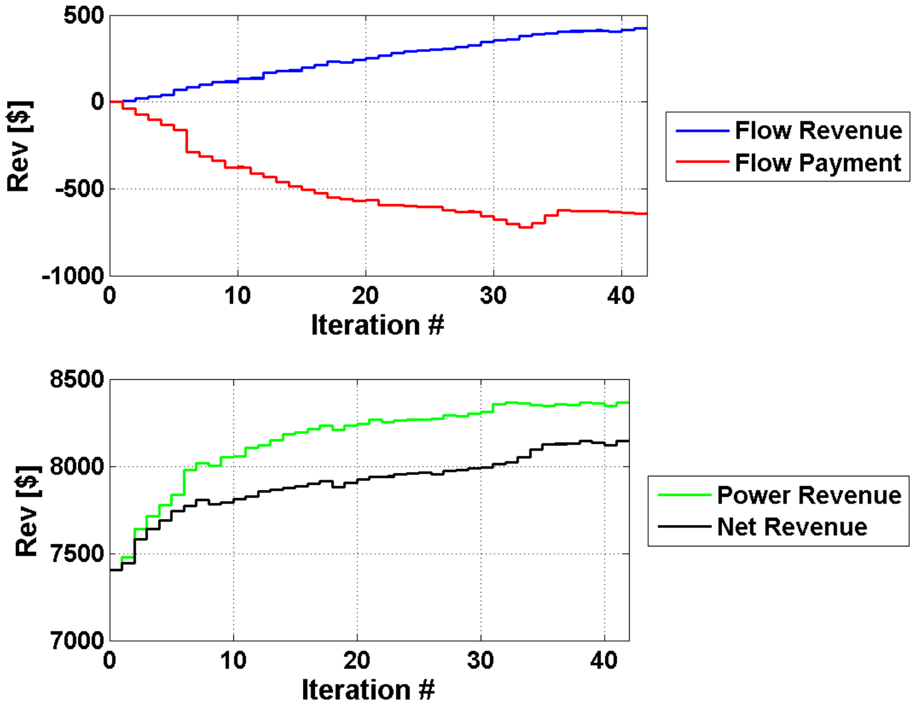

Figure 39: Revenue progression of middle facility. Flow payments are being made to the first plant, and flow revenue is from payments made by the last plant.

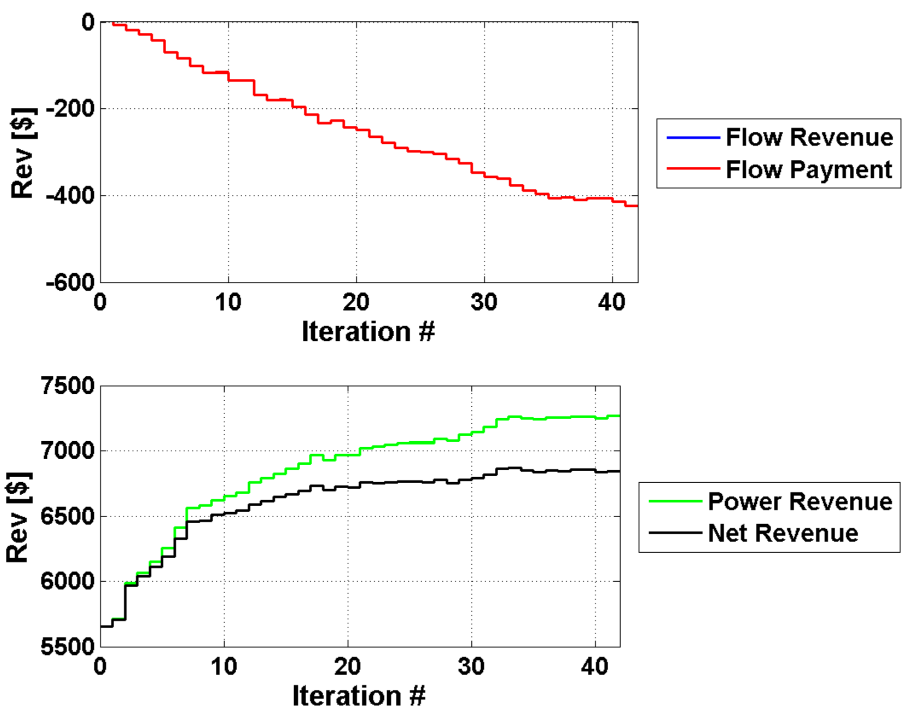

Figure 40: Revenue progression of last facility. No flow revenue because there is no downstream facility.

3.3.2 Comparison to another Optimization Technique

Since the goal of the optimization is to maximize both the revenue of the individual stations and of the system as a whole, the results shown above need to be compared to some measure of optimality. Since the problem is highly complex, determining a global optimum can be a difficult task. However, using some knowledge of the constraints of the system, a reasonable prediction of the inputs resulting in the global optimum can be made.

In this case, since the system is composed of three facilities with similar heads and sets of operating turbines, and because the second and third installations are run-of-river, approximating the global optimum was a relatively straightforward matter of adjusting the price profile of the first plant to reflect the expected total value of the power production as the water flows through the system. Thus, in order to reflect the revenue of all three facilities in the price profile of only the first one, three different price profiles were added together: the base price profile, a price profile shifted 3 hours earlier in the day (representing the expected revenue of the second facility), and a price profile shifted 6 hours earlier in the day (representing the expected revenue of the third facility). This price profile, shown below in Figure 41, was then used to analyze the first facility only, and the resulting optimal flow rates for the first facility were sent downstream to the second installation. The operation of the two downstream power plants was then determined using the iterative method, which had very little effect on their operations as determined by DDP alone since they are run-of-river installations and their available water is predetermined by the release of the first plant.

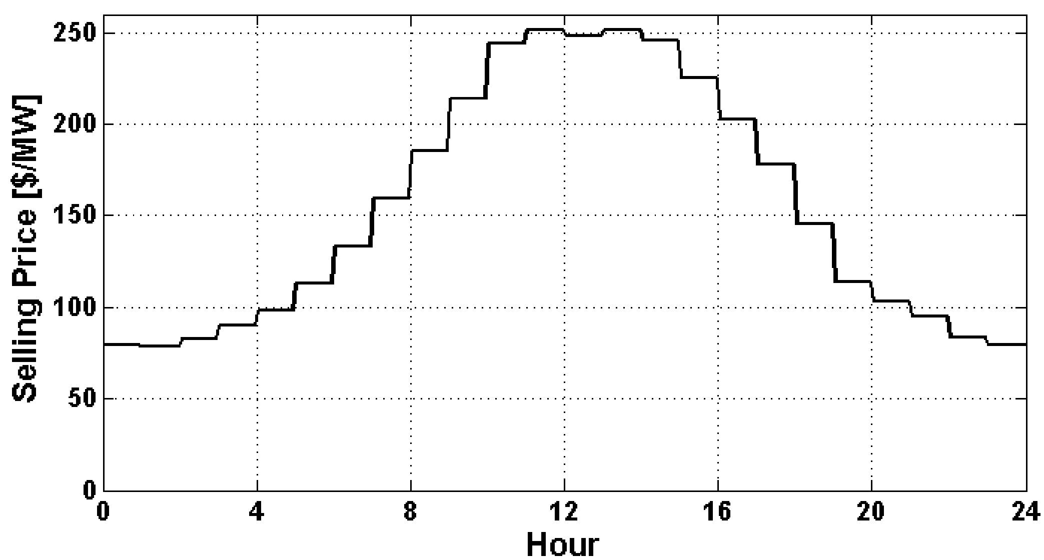

Figure 41: Combined price profile used in the comparison analysis.

The results of this optimization are compared with the results from the iterative optimization results below in Figure 42 – Figure 44. It is seen that the general result is that the flow is essentially shifted one hour earlier than it was mapped by the iterative system optimization. This result is indeed closer to what might be expected for the system optimization since the total shift away from the initial analysis in which each facility was optimized independently is now three hours, matching the flow time between each pair of power plants. Thus the flow rate peaks before the peak in price for the first facility, during the peak in price for the middle one, and after the peak in price for the last one.

Quantitatively, the revenue results of the two different methods are shown below in Table 5. Note that while both methods increase the expected revenue above what was determined when each facility was optimized independently of the others, the more direct method described above resulted in a slightly higher final expected revenue for the system as a whole (a 12.1% increase above the initial, independently determined total revenue as opposed to a 10.1% increase), but it did nothing to distribute the resulting higher net revenue among the different facilities. Thus, while the second and third facilities would enjoy higher revenue, the first facility would generate a significantly lower one than the case in which they operate independently. Furthermore, though the direct method results in higher total revenue for the system, there is no clear indication on how a group of independent facility owners should cooperate to reach this operation method.

It is prudent at this point to highlight the value of the system iterative method subject of this thesis over the more direct method described above. The iterative method developed herein is finely balanced via power and water valuation, and seeks to unearth not only the optimal solution for the system, but a solution that is also optimal for each operator. The operational strategy thus unveiled is bound to interest all significant parties. Although the direct method achieves greater total revenue, it does not show why facilities generating lower revenue than they would if they operated independently should act cooperatively with the other facilities.

Table 5: Revenue comparisons between iterative method and direct method.

	Facility 1	Facility 2	Facility 3	Total
Initial Revenue	$ 8,498.95	$ 7,402.37	$ 5,652.20	$ 21,553.52
Iterative Final Revenue	$ 8,733.87	$ 8,143.14	$ 6,844.70	$ 23,721.71
Direct Final Revenue	$ 7,461.87	$ 8,630.80	$ 8,061.61	$ 24,154.28

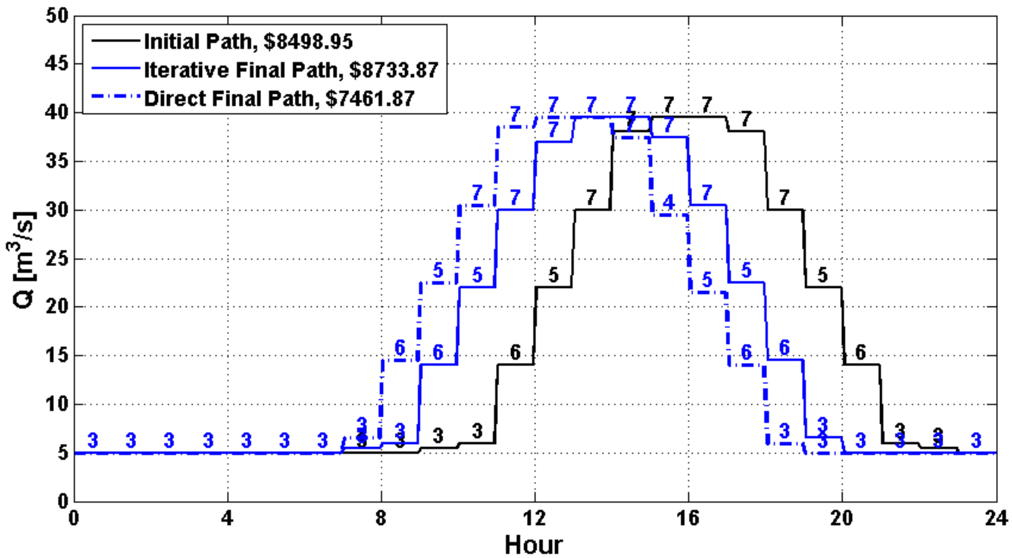

Figure 42: Comparison of the flow rates for the first facility determined by the iterative method for system optimization, the direct method, and the initial, independently optimized flow rates.

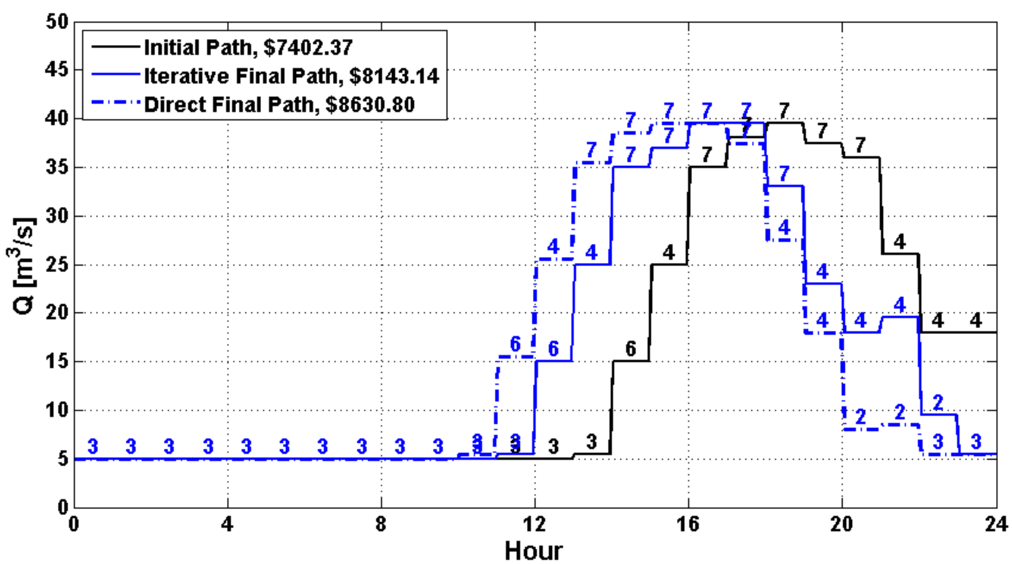

Figure 43: Comparison of the flow rates for the second facility determined by the iterative method for system optimization, the direct method, and the initial, independently optimized flow rates.

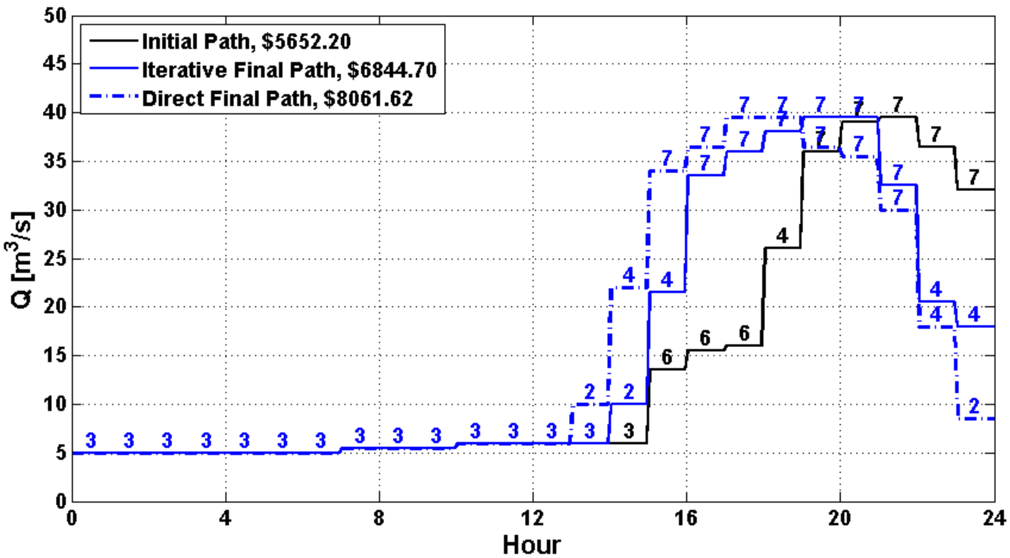

Figure 44: Comparison of the flow rates for the third facility determined by the iterative method for system optimization, the direct method, and the initial, independently optimized flow rates.

3.4 Sensitivity Analysis

3.4.1 Sensitivity to Price Profile

Since the price profiles used are only predictions of the day's prices, it is important to look at the effect of variation in the price profile on the resulting optimization. Good predictions of the day-ahead price profile can have mean absolute percentage errors as high as 10% [22]. Thus, in order to test the sensitivity to the accuracy of the predicted price profile, 20 variations of the summer price profile were randomly generated in such a way that the general shape and smoothness of the price profile was preserved, while positive and negative variations were introduced. The resulting set of price profiles is shown below in Figure 45.

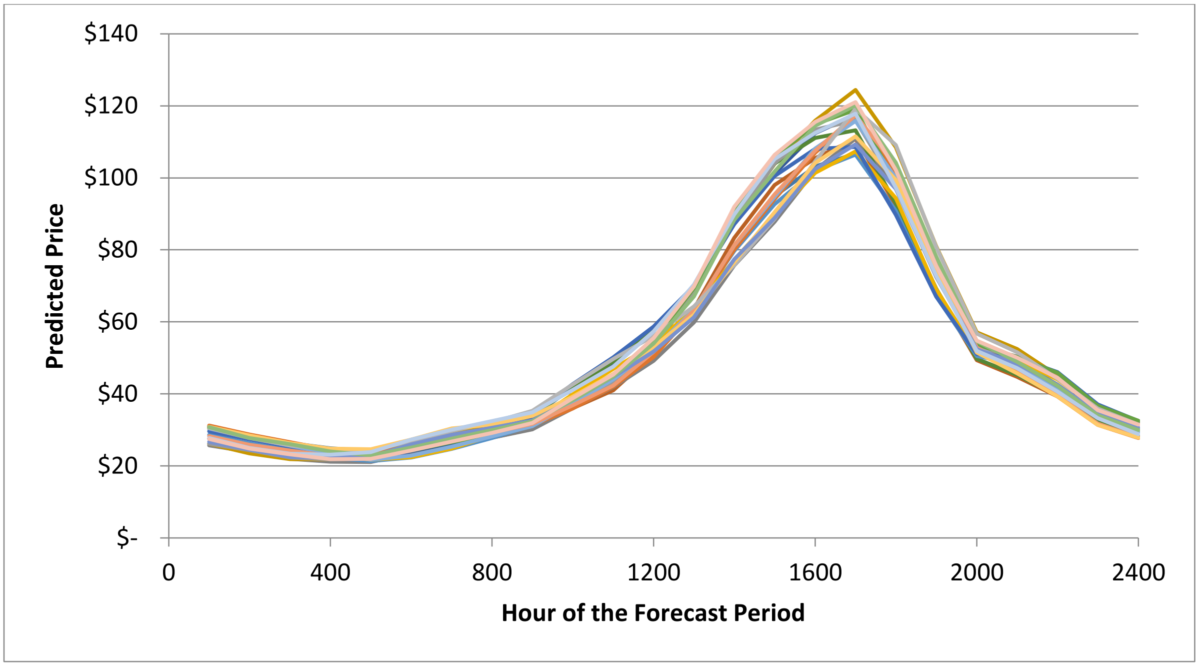

Figure 45: Randomized price profiles used in sensitivity analysis.

The iterative analysis was performed for each of these price profiles to determine the expected optimal revenue for each one. Then, again for each price profile, the total revenue was calculated using the optimal flow rates for the original price profile, and the differences between these values and the optimal expected previous values were calculated. Since the iterative water valuing method is not guaranteed to return globally optimal inputs, the differences found were both positive and negative. The average value of the differences was 0.36% of the predicted revenue of the original price profile, the absolute value of the differences was 0.61%, and the standard deviation of the differences was 0.99%. When compared to the expected 10.1% increase in revenue beyond that which would be produced without cooperation and water valuation (from Section 3.3.1), this is an acceptably small variation.

3.4.2 Sensitivity to Gradient Analysis Constraints

During the determination of the revenue gradients to be used in the second stage of the iteration (ΔR in Equation 2.55 and Equation 2.57) limits were imposed upon the flow rates and reservoir volumes that would be included in the analyses in order to reduce the run time. These limits took the form of a narrow band of flow rates and volumes to be considered, centered about the optimal path determined in the first stage of each step of the process. These bands restricted the flow rates being considered to those within a certain number of grid steps away from the previously determined optimum. The band of volumes to be considered consisted of those within a number of volume grid steps, equal to a multiple of the size of the flow rate band, of the previous optimal path.

In order to determine the effect of the grid size on the final result, the analysis was performed for many different combinations of band sizes for flow rate and volume, and for both a summer price profile and a winter price profile. The results are shown below in Table 6 and Table 7.

Table 6: Sensitivity analysis results for a summer price profile.

		Flow rate band size in number of grid sized steps					
Volume band size as a multiple of the flow rate band size	**Total Rev.**	1	2	4	6	10	50
	x1	$21,882	$21,933	$23,481	$23,332	$23,481	$23,481
	x2	$21,933	$23,271	$23,486	$23,486	$23,481	$23,481
	x4	$23,028	$23,335	$23,486	$23,486	$23,481	$23,481
	x6	$23,475	$23,335	$23,486	$23,486	$23,481	$23,481
	x10	$23,475	$23,335	$23,486	$23,486	$23,481	$23,481
	% of Max	1	2	4	6	10	50
	x1	93.17%	93.39%	99.98%	99.35%	99.98%	99.98%
	x2	93.39%	99.09%	100.00%	100.00%	99.98%	99.98%
	x4	98.05%	99.36%	100.00%	100.00%	99.98%	99.98%
	x6	99.95%	99.36%	100.00%	100.00%	99.98%	99.98%
	x10	99.95%	99.36%	100.00%	100.00%	99.98%	99.98%
	Run Time (minutes)	1	2	4	6	10	50
	x1	2.27	1.63	23.02	18.32	37.92	187.20
	x2	1.38	11.05	22.87	29.32	37.95	202.95
	x4	10.52	14.65	23.12	29.73	38.83	227.02
	x6	15.62	14.72	23.48	29.70	39.53	241.42
	x10	15.62	14.77	23.68	29.85	40.12	248.22

Table 7: Sensitivity analysis for a winter price profile.

		Flow rate band size in number of grid sized steps				
Volume band size as a multiple of the flow rate band size	**Total Rev.**	1	2	4	6	10
	x1	$18,172	$19,070	~~$13,713~~	~~$13,707~~	$19,608
	x2	$19,126	$19,407	$19,576	~~$13,700~~	$19,610
	x4	$19,617	$19,560	$19,619	~~$13,700~~	$19,610
	x6	$19,492	$19,560	$19,619	~~$13,700~~	$19,610
	x10	$19,492	$19,560	$19,619	~~$13,700~~	$19,610
	% of Max	1	2	4	6	10
	x1	92.62%	97.20%	~~69.89%~~	~~69.87%~~	99.94%
	x2	97.49%	98.92%	99.78%	~~69.83%~~	99.95%
	x4	99.99%	99.69%	100.00%	~~69.83%~~	99.95%
	x6	99.35%	99.69%	100.00%	~~69.83%~~	99.95%
	x10	99.35%	99.69%	100.00%	~~69.83%~~	99.95%
	Run Time (minutes)	1	2	4	6	10
	x1	6.88	28.45	~~19.14~~	~~24.87~~	57.68
	x2	19.20	12.98	26.93	~~25.30~~	54.10
	x4	11.92	15.73	23.32	~~25.89~~	56.20
	x6	9.07	15.80	23.56	~~26.32~~	58.25
	x10	9.10	15.88	23.81	~~26.81~~	59.85

Note: The crossed out cells correspond to analyses which returned a result in which the final installation was not able to find any inputs which adhered to the constraints.

For simplicity, the band sizes will be referred to using integers representing the number of grid steps away from the previously determined optimal path for both flow rate and volume.

Examination of these results indicated that using a flow rate band size of 2 or 4 with a volume band size of 8 (x4 and x2, of 2 grid steps and 4 grid steps, respectively). Using a flow rate band size of 4 and a volume band size of 8 resulted in the higher total revenues, but also in longer run times. On the other hand, the flow rate band size of 2 with a volume band size of 8 resulted in slightly lower revenue, though the run time was significantly lower. Furthermore, it was found that utilizing a flow rate band size larger than 4 or a volume band size more than 4 times the flow rate band size typically did not result in any further increases in revenue. For the purpose of this work, identifying the inputs that result in the highest total revenue is more important than low run times. Therefore the results presented herein were obtained with band sizes of 4 and 8 for flow rate and volume, respectively. These selections result in an expected run time on the order of 25 minutes.

3.4.3 Sensitivity to Payment Amount

The final sensitivity analysis regards sensitivity to the amount paid for increased flow rate released by upstream facilities. The amount paid that yielded the results shown above, as detailed in Section 2.4, consisted of a base payment equal to the loss incurred by the upstream facility, plus an additional amount determined by a finite geometric series with a base value of one half the net increase in revenue due to the evaluated increment in flow, a ratio of one half, and a number of terms equal to the number of increments the upstream facility increases its flow rate beyond the previously determined optimal value.

In order to examine the sensitivity to payment amounts, different values of the base value used in the geometric series were tested in order to examine the effects on the final solution. The resulting revenues are shown below in Table 8. Note that although the total increases in final system revenue are largely very similar (in fact, several of them are identical), the water revenues, and thus payments and total revenue of facilities receiving payments, differ a great deal. This occurs due to the fact that the final optimal flow rates for each facility are very similar in each case, as shown below in Figure 46 – Figure 48, but the payments being made are changing a great deal.

These results indicate that the method is consistent in its predictions of the optimal flow rate profiles (though there is room for improvement as was demonstrated in Section 3.3.2). However, the payments necessary to arrive at the target flow rate profiles are not unique.

Table 8: Results of the analysis of the sensitivity to the base value of the geometric portion of the payments.

		Facility 1	Facility 2	Facility 3	Total
	Initial Revenue	$8,498.95	$7,402.37	$5,652.20	$21,553.52
5% Geometric Base	Water Revenue	$460.22	$287.02	$0.00	$747.24
	Power Revenue	$8,074.46	$8,376.76	$7,285.59	$23,736.81
	Total Revenue	$8,534.68	$8,203.57	$6,998.57	$23,736.81
	Percent Increase	0.42%	10.82%	23.82%	10.13%
20% Geometric Base	Water Revenue	$508.69	$280.54	$0.00	$789.22
	Power Revenue	$8,074.46	$8,405.58	$7,408.87	$23,888.91
	Total Revenue	$8,583.15	$8,177.43	$7,128.34	$23,888.91
	Percent Increase	0.99%	10.47%	26.12%	10.84%
40% Geometric Base	Water Revenue	$646.58	$444.24	$0.00	$1,090.81
	Power Revenue	$8,049.39	$8,381.82	$7,350.07	$23,781.28
	Total Revenue	$8,695.97	$8,179.48	$6,905.84	$23,781.28
	Percent Increase	2.32%	10.50%	22.18%	10.34%

50% Geometric Base	Water Revenue	$645.64	$423.36	$0.00	$1,069.01
	Power Revenue	$8,088.22	$8,365.42	$7,268.07	$23,721.71
	Total Revenue	$8,733.87	$8,143.14	$6,844.70	$23,721.71
	Percent Increase	2.76%	10.01%	21.10%	10.06%
60% Geometric Base	Water Revenue	$690.90	$432.84	$0.00	$1,123.74
	Power Revenue	$8,074.46	$8,376.76	$7,285.59	$23,736.81
	Total Revenue	$8,765.36	$8,118.71	$6,852.75	$23,736.81
	Percent Increase	3.13%	9.68%	21.24%	10.13%
75% Geometric Base	Water Revenue	$848.86	$464.85	$0.00	$1,313.71
	Power Revenue	$8,074.46	$8,376.76	$7,285.59	$23,736.81
	Total Revenue	$8,923.33	$7,992.75	$6,820.74	$23,736.81
	Percent Increase	4.99%	7.98%	20.67%	10.13%
100% Geometric Base	Water Revenue	$949.43	$451.44	$0.00	$1,400.86
	Power Revenue	$8,194.35	$8,271.98	$7,071.72	$23,538.05
	Total Revenue	$9,143.78	$7,773.99	$6,620.28	$23,538.05
	Percent Increase	7.59%	5.02%	17.13%	9.21%

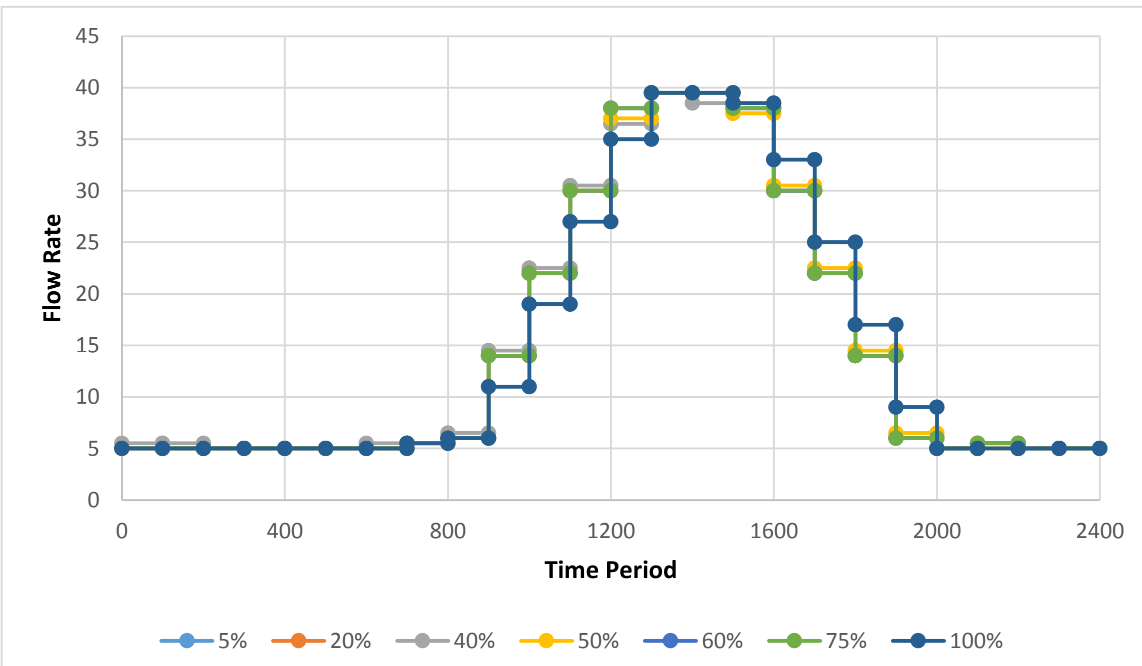

Figure 46: Flow profiles of first facility resulting from various payment amounts.

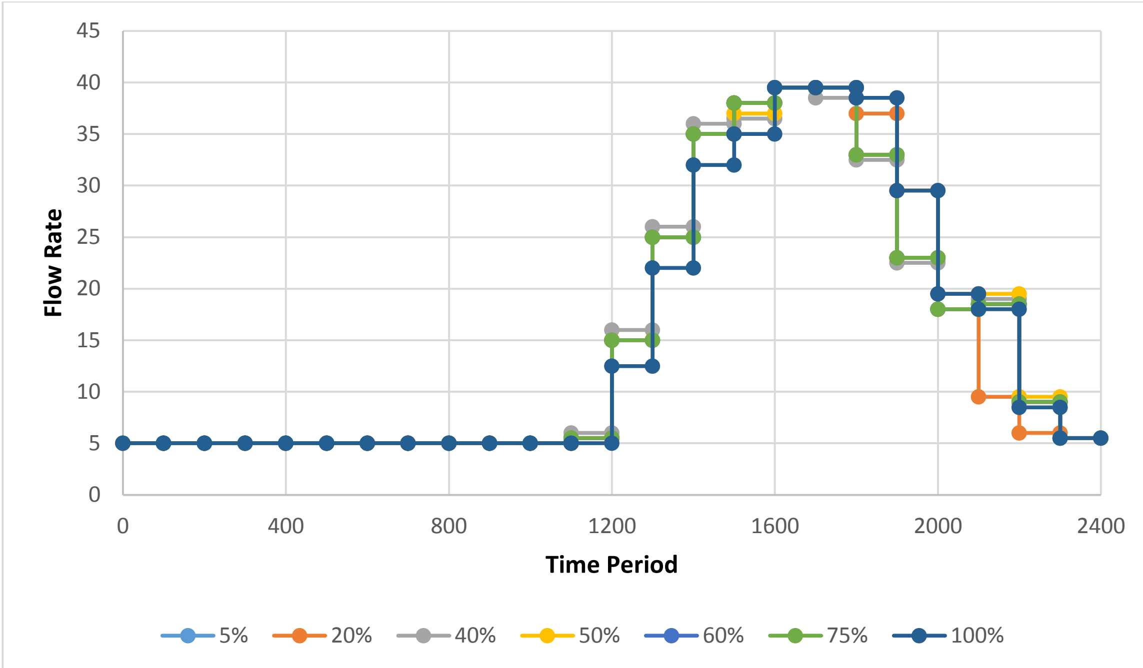

Figure 47: Flow profiles of middle facility resulting from various payment amounts.

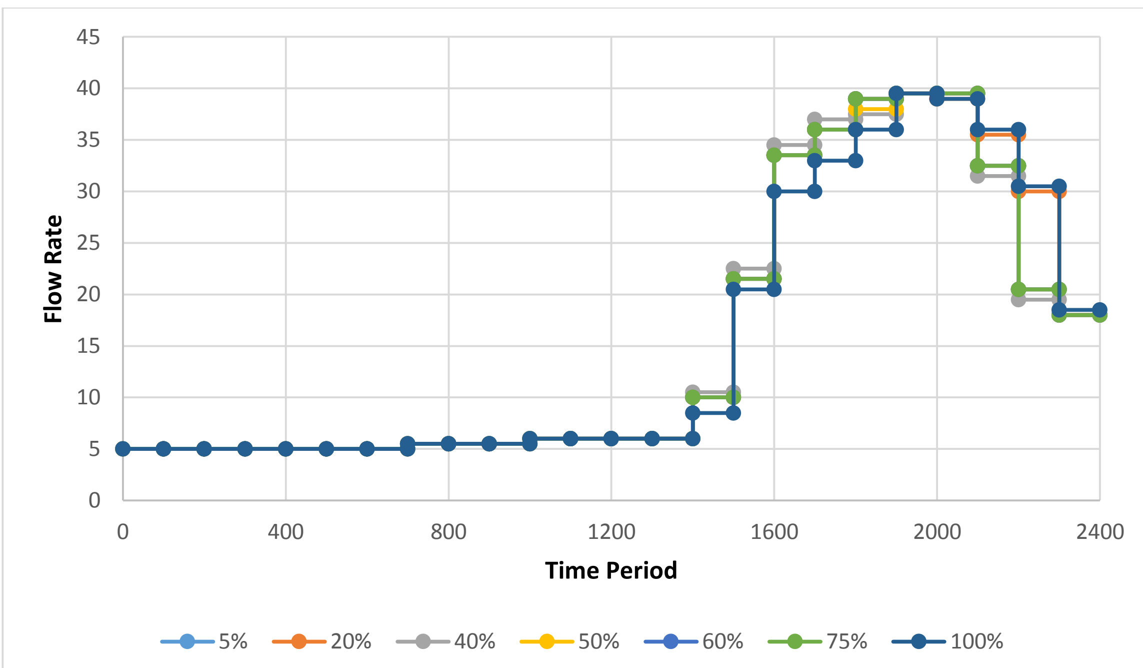

Figure 48: Flow profiles of last facility resulting from various payment amounts.

4 Conclusions

In conclusion, the methodology developed successfully increased the revenue of all parties in a multiple owner hydropower system. Although the global optimum for the system was not achieved, the final results of the iterative DDP-gradient analysis compared favorably to those obtained using a more direct optimization method. Furthermore, a number of insights into the nature of hydropower optimization and the value of water were gained.

First, regarding the modeling of hydropower facilities, many improvements were made over previously existing methods. Those methods frequently failed to account for the variation of the operating head with the net flow rate through a facility. However, more importantly than that, most previous methods either attempted to use a single power versus flow rate curve for the entire facility, or treated each turbine in the facility as an individual unit. None were found that utilized a preliminary analysis such as that described in Section 2.1.2. This preliminary analysis sacrifices none of the accuracy provided by modelling each turbine individually, but provides drastically reduced optimization run times.

Second, regarding the optimization of a single hydropower facility, application of DDP is not a new approach. However, much work was done to accelerate the run time of optimizing a single plant, primarily focused on eliminating infeasible states and inputs from the analysis and utilizing a successive refinement technique. Ultimately the run time was reduced to the order of a few seconds for a hydropower facility with a large reservoir and three turbines. Furthermore, DDP was found to be particularly useful in optimizing hydropower facilities because of its flexibility, which allows it to easily account for both discrete and continuous constraints, states, and inputs.

Finally, the chief contribution of this work is the hybrid DDP-gradient shooting iterative method. Nothing resembling this method was found in any of the reviewed literature, either in its attempt to address the problem of optimizing a system with multiple owner/operators or in its valuation of water based on future downstream revenue potential. The hybrid method was successful in achieving its goal of increasing the expected revenue of both the system and of each individual installation. It is interesting to note that the problem is essentially a Nash game in which players solve dynamic programs that capture multi-period optimization problems representing the operations of each facility. Notably, each player’s problem is parametrized by the strategies of its competitors and the game allows for the possibility for each player to offer payments to its neighbor(s) for specific changes in operations. A Nash equilibrium point is then defined by the final flow rates and operating sets of each facility, along with the payments used to motivate those changes. As was discussed in Section 3.4.3, the final flow rates determined are relatively independent of the payments being made. In fact, it is expected that if the method were able to consistently arrive at the global optimum, then the final flow rates determined would all be identical, no matter the method of dividing the increased revenue. Thus, the Nash equilibrium is not unique, though the final flow rates are expected to be. This does not diminish the utility of the approach; rather, it allows neighboring facilities to enter into pricing schedule negotiations.

In conclusion, although the proposed scheme did not produce the centralized optimal solution, the numerical studies suggest that the obtained solution does not differ significantly from its global counterpart (the estimated global optimum revenue was 12.1% higher than total individually optimized revenue, but only 1.8% higher than the solution obtained using the proposed technique). More importantly, perhaps, considerable insight was gained into the valuation of water in a hydrologically connected system of hydropower facilities. This insight, and the methods developed will hopefully provide an excellent foundation for future work in this area.

5 Future Work

There are a number of avenues that this work has revealed as having potential for future investigation. Some of the possibilities for future work relate to the optimization of single hydropower facilities and the DDP used to optimize single facilities, while others relate to the gradient approach used to link the stations together.

The optimization of a single hydropower facility could be improved in several ways. Of particular interest is any way in which the optimization could be improved so that the analysis would be sped up. Several possibilities for this were mentioned earlier. Using an interpolation function to determine the future value-to-go function based on the inputs and states currently under consideration could drastically decrease the number of volume states required, thus significantly decreasing the computational time required. Additionally, as was noted in the discussion of the results shown in Figure 25, it is possible that the optimal flow rates and operating sets depend very little on the volume of the reservoir, as long as the reservoir limits are not encountered. Thus it may be possible for future refinements of the DDP optimization of a single hydropower station to incorporate a volume discretization that not only does not match the flow rate discretization, but perhaps also volume discretization that is very coarse when not near a limit, and more refined nearer to the operating limits.

Another way in which the optimization of a single hydropower facility could be improved is the inclusion of the ability to apply a final value to the water in the reservoir so that the target volume of the reservoir does not have to be known. Such a final value would be reflective of the expected future value of water in the following days, and would be determined a mid-term optimization analysis. This would allow the daily optimization to not only optimize the operation of the facility for the day, but also extend the optimization to including determination of the optimal final volume of water left in the reservoir at the end of the day.

There are also several ways in which the hybrid system optimization might be enhanced by further work. As was noted earlier, it does not successfully achieve a global optimum, and it is believed that development and application of these ideas would greatly improve the final results obtained.

First, due to the tight constraints on the volume of water held by run-of-river facilities, the methods used to determine the revenue gradients and corresponding payments frequently cannot properly analyze the gradients when the upstream member of a pair of facilities is run-of-river. Thus, another method may better serve this type of situation. One possibility would be to simply determine the appropriate payment by analyzing the effect of increased flow into the downstream member during each time period. Such a method would have to take into account that the downstream installation would be receiving more water than the upstream one can release, but that could easily be accounted for by normalizing the resulting revenue.

A second possible modification that could improve the final results of the iterative method would be the addition of an additional phase to the analysis to take place before the analysis described above. This new phase would consist of two stages that would be identical to the two stages described above, with one key difference: the power surfaces of each operating set in each facility would be adjusted so that they would conform to a single convex power surface for the entire station. This would eliminate the possibility of the optimization stopping due to encountering the local low efficiency operating regions that occur near the transition from one operating set to another. This new first phase of the analysis would iterate to convergence, and then the water payments found would be used with the more accurate system model to zero in on the final solution. Such a method would likely provide final results much closer to the global optimum than were achieved with only the two stages developed in this thesis.

Bibliography

[1] U.S. Bureau of Reclamation, "The History of Hydropower in the United States," [Online]. Available: http://www.usbr.gov/power/edu/history.html. [Accessed 8 May 2013].

[2] U.S. Department of Energy, "Water Power Program: History of Hydropower," [Online]. Available: http://www1.eere.energy.gov/water/hydro_history.html. [Accessed 8 May 2013].

[3] National Hydropower Association, "History of Hydro," [Online]. Available: http://www.hydro.org/tech-and-policy/history-of-hydro/. [Accessed 8 May 2013].

[4] National Historic Landmarks Program, "Roosevelt Dam," [Online]. Available: http://www.nps.gov/history/nhl/DOE_dedesignations/Roosevelt.htm. [Accessed 13 5 2013].

[5] R. A. Duncan, G. E. Seymore, D. L. Streiffert and D. J. Engberg, "Optimal Hydrothermal Coordination for Multiple Reservoir River Systems," *IEEE Transactions on Power Apparatus and Systems,* Vols. PAS-104, no. 5, pp. 1154-1159, 1985.

[6] A. J. Conejo, J. M. Arroyo, J. Contreras and F. A. Villamor, "Self-Scheduling of a Hydro Producer in a Pool-Based Electricity Market," *IEEE Transactions on Power Systems,* vol. 17, no. 4, pp. 1265-1272, 2002.

[7] S. Mariano, J. Catalão, V. Mendes and L. Ferreira, "Optimising power generation efficiency for head-sensitive cascaded reservoirs in a competitive electricity market," *International Journal of Electrical Power & Energy Systems,* vol. 30, no. 2, pp. 125-133, 2008.

[8] S. Mariano, M. Calado and L. Ferreira, "Profit-Based Optimal Operation of a Head-Dependent Hydroelectric Power Station in the Bilateral Market," in *International Conference on Renewable Energies and Power Quality*, Granada, Spain, 2010.

[9] H. Pousinho, V. Mendes and J. Catalão, "Scheduling of a hydro producer considering head-dependency, price scenarios and risk-aversion," *Energy Conversion and Management,* vol. 56, pp. 96-103, 2012.

[10] D. S. O. Nilsson, "Hydro Unit Start-up Costs and Their Impact of the Short Term Scheduling Strategies of Swedish Power Producers," *IEEE Transactions on Power Systems,* vol. 12, no. 1, pp. 38-44, 1997.

[11] W. A. Hall, W. S. Butcher and A. Esogbue, "Optimization of the Operation of a Multi-Purpose Reservoir by Dynamic Programming," *Water Resources Research,* vol. 4, no. 3, pp. 471-477, 1968.

[12] R. C. Harboe, F. Mobasheri and W. Yeh, "Optimal Policy for Reservoir Operation," *Journal of the Hydraulics Division,* vol. 96, no. 11, pp. 2297-2308, 1970.

[13] M. Heidari, V. T. Chow, P. V. Kokotović and D. D. Meredith, "Discrete Differential Dynamic Programming Approach to Water Resources Systems Optimization," *Water Resources Research,* vol. 7, no. 2, pp. 273-282, 1971.

[14] J. S. Windsor and V. T. Chow, "Multireservoir Optimization Model," *Journal of the Hydraulics Division,* vol. 98, no. 10, pp. 1827-1845, 1972.

[15] W. W.-G. Yeh and L. Becker, "Multiobjective Analysis of Multireservoir Operations," *Water Resources Research,* vol. 18, no. 5, pp. 1326-1336, 1982.

[16] J. J. Shaw, R. F. Gendron and D. P. Bertsekas, "Optimal Scheduling of Large Hydrothermal Power Systems," *IEEE Transactions on Power Apparatus and Systems,* Vols. PAS-104, no. 2, pp. 286-294, 1985.

[17] M. V. F. Pereira and L. M. V. G. Pinto, "Stochastic Optimization of a Multireservoir Hydroelectric System: A Decomposition Approach," *Water Resources Research,* vol. 21, no. 6, pp. 779-792, 1985.

[18] A. I. Cohen and V. R. Sherkat, "Optimization-Based Methods for Operations Scheduling," *Proceedings of the IEEE,* vol. 75, no. 12, pp. 1574-1591, 1987.

[19] P.-H. Chen and H.-C. Chang, "Genetic Aided Scheduling fo Hydraulically Coupled Plants in Hydro-Thermal Coordination," *IEEE Transactions on Power Systems,* vol. 11, no. 2, pp. 975-981, 1996.

[20] J. García-González, E. Parrilla and A. Mateo, "Risk-Averse Profit-Based Optimal Scheduling of a Hydro-Chain in the Day-Ahead Electricity Market," *European Journal of Operational Research,* vol. 181, no. 3, pp. 1354-1369, 2007.

[21] J. W. F. D. G. R. A. Ponrajah, "Systems to Optimise Conversion Efficiencies at Ontario Hydro's Hydroelectric Plants," *IEEE Transactions on Power Systems,* vol. 13, no. 3, pp. 1044-1050, 1998.

[22] Z. D. K. M. e. a. X. Chen, "Electricity Price Forecasting with Extreme Learning Machine and Bootstrapping," *IEEE Transactions on Power Systems,* vol. 27, no. 4, pp. 2055-2062, 2012.

VITA

David Beevers

David B. Beevers was born on March 15, 1986 to Eric and Kathy Beevers in Morgantown, West Virginia. He was home-schooled for most of his early education and obtained his degree in Mechanical Engineering from Grove City College. During his time at Grove City, David discovered his passion for teaching, and decided to pursue an advanced degree with the intent of becoming a teacher. He chose to pursue his doctorate at The Pennsylvania State University, where he met his wife. David and Jessica were married on May 28, 2011, and have one son, Dillan, who was born on October 21, 2014. During his first couple of years at Penn State, David worked as a teaching assistant for a variety of undergraduate courses before joining the Hydropower Research Gfroup founded by Dr. Cimbala in 2010. During his time with the hydropower research group, David studied the development of component based computer modeling of hydropower systems, design of the turbine capacities of a pump-storage plant for the greatest continuous operating range, and finally, optimization of the operation of a hydropower system for his thesis work. David earned his Master of Science Degree in Mechanical Engineering upon his satisfactory completion of his comprehensive exam in December, 2013. In addition to his research work at Penn State, David also worked as the primary instructor for M E 201 – Introduction to Thermal Sciences. Currently, David is preparing to begin his work as a lecturer in mechanical engineering at Penn State University, The Behrend College in August, 2015.