PERFORMANCE OF A SAMPLING STOCHASTIC DYNAMIC PROGRAMMING ALGORITHM WITH VARIOUS INFLOW SCENARIO GENERATION METHODS

by

JENNIFER LYNN SCHAFFER

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF APPLIED SCIENCE

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES

(Civil Engineering)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

January 2015

Abstract

We present the implementation of a Sampling Stochastic Dynamic Programming (SSDP) algorithm to maximize water value, while meeting consumer demand for the BC Hydro hydroelectric system in British Columbia, Canada. The implementation includes power generation facilities on the Columbia and Peace River systems.

Variability of natural streamflow into a reservoir is a major source of uncertainty when developing reservoir operation policies and determining the value of water within a system. This study investigates SSDP model performance with various hydrologic inputs. Sixty years of historical data are used to generate hydrologic scenarios comprised of inflow and forecast sequences as input to the SSDP model. Scenario types studied include historical record data, inflows and forecasts generated from an autoregressive lag-1 model, and BC Hydro ensemble streamflow prediction forecasts.

We present results of our implementation of the SSDP algorithm including a discussion on improved reservoir operation policy and the future value of water with various hydrologic inputs. We also present our investigation of the marginal value of water with the evolution of forecasts. Results indicate that forecasts are most valuable in determining the value of water during the early freshet, and the value added from updating future forecasts diminishes as the time in which the forecast is made progresses through the melting period.

Preface

This thesis is based on the development and testing of a model applied to the BC Hydro Hydroelectric System. The author, Jennifer Schaffer, was the lead investigator responsible for model development, results analysis, and manuscript composition for research described in Chapters 2, 3, and 4. A version of Chapter 2 was published in the conference proceedings of HydroVision 2014 (Schaffer and Shawwash 2014). Results described in Chapter 3 were presented at the Canadian Water Resources Association Workshop, "From operational hydrological forecast to reservoir management optimization" (Schaffer 2014).

Prof. Ziad Shawwash (University of British Columbia) provided guidance in algorithm formulation, determining testing procedures, and analyzing results. Prof. Shawwash also contributed to manuscript editing of all chapters. Amr Ayad (University of British Columbia/BC Hydro) and Ziming Guan (University of British Columbia) also provided assistance with model formulation and testing procedures. Calculations described in Chapters 2.4.2, and 4.2.2 were developed with assistance from Prof. Jery Stedinger and Jonathan Montagne of Cornell University. Additionally, Pascal Cote of Rio Tinto Alcan provided material regarding value function approximation methodology described in Chapter 4.2.3. Historic hydrologic data used in all chapters of this manuscript was provided by BC Hydro, and assistance with its interpretation was provided by Adam Gobena (BC Hydro). Synthetic hydrologic data for use in Chapters 2 and 3 was generated by and modified for the author's use by Ziming Guan. Much of the computational effort was via hardware systems from BC Hydro's Generation Resource Management Group supervised by Alaa Abdalla.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

AMPL	A Mathematical Programming Language
BCH	British Columbia Hydro and Power Authority (BC Hydro)
CRT	Columbia River Treaty
DP	Dynamic Programming
ESP	Ensemble Streamflow Prediction
FC	Forecast
MVW	Marginal Value of Water
SDP	Stochastic Dynamic Programming
SDDP	Stochastic Dual Dynamic Programming
SSDP	Sampling Stochastic Dynamic Programming

Acknowledgements

I would like to express my sincerest gratitude to my research supervisor, Prof. Ziad Shawwash, for his guidance in my research pursuits and imparting his vast breadth of knowledge on me throughout my studies.

I also am grateful to Prof. Jery Stedinger from Cornell University for sharing his expertise and providing valuable insight on my research.

I am appreciative of the funding of my research from grants provided to Dr. Shawwash by the Natural Sciences and Engineering Research Council of Canada (NSERC) and from BC Hydro.

I also would like to thank BC Hydro for their cooperation in support of my research by making a vast amount of resources available to me, including access to data, computing systems, and their employees. I would like to thank specific individuals in allowing me to take advantage of their wealth of experience—Amr Ayad, Ziming Guan, Adam Gobena, and Alaa Abdalla.

For my parents, Karen and John, and my husband, Nik

1 Introduction

British Columbia, Canada generates over ninety five percent of its energy from hydroelectric facilities located throughout the province. The system includes 31 generation stations and over 75,000 kilometers of transmission lines that are linked with the province of Alberta and the western United States allowing the exchange of energy over a large market as seen in Figure 1 ("BC Bulk Transmission System" 2007). BC Hydro is the crown corporation responsible for generating, purchasing, distributing, and selling electricity to its customers both within and outside the province. Hydroelectric generation gives the system flexibility to store energy and purchase electricity from the market when prices are low and then generate energy to sell to the market when prices are high.

Figure 1. BC Hydro transmission system (BC Hydro)

The operations policy determining when to store and release water throughout a hydroelectric power system is typically based on the results of an optimization problem where the objective is to maximize benefits to the system throughout the planning horizon. A technique called dynamic programming (DP) has been extensively used in solving the reservoir operations problem. In dynamic programming, a multistage planning problem is broken into a series of smaller single stage problems that are solved successively. The problem is described in each stage by the state of the system, which is often reservoir storage. The algorithm optimizes the decision to release water to maximize the sum of current benefits to the system and future benefits achieved when making that decision (Bellman 1957). The recursive equation is solved at every stage, starting at the last stage and moving backwards in time, and for every state of the problem. While dynamic programming is a powerful optimization tool, the algorithm assumes that inputs are known.

In reservoir operations planning, many uncertainties exist that affect operations decisions including pricing, load demand, and reservoir inflows. The optimization problem becomes more complicated when model inputs are uncertain and is one of the main challenges in modelling. Yet when uncertainty is considered, models are better able to provide estimates of the value of system resources and provide better projections of expected values of revenues, energy generation, and market transactions given potential future conditions (Abdalla et al. 2013). Knowing the value of system resources and more specifically the marginal value of system resources (water in storage) is extremely useful in reservoir operations planning as it is the driving force in policy decisions.

Reservoir operations decisions consider the tradeoff between releasing water to gain immediate benefits and storing water to realize benefits in the future. The marginal value of water is the incremental benefit associated with the change of the amount of water in reservoir storage (Tilmant et al. 2008). An optimal decision is made when the value of releasing an additional volume of water is equal to the value of storing water for future use, i.e. the marginal values are equal. This is the point where the total value of system resources, the sum of the immediate and future benefits, is maximized.

BC Hydro's Water Value Project aims to incorporate uncertainties in reservoir inflows into a long term planning model to estimate the value of system resources. Many stochastic optimization algorithms that are being investigated including Stochastic Dynamic Programming (SDP), Stochastic Dual Dynamic Programming (SDDP), Sampling Stochastic Dynamic Programming (SSDP), and Reinforcement Learning.

The SDP algorithm is solved using the same methodology as dynamic programming but is able to consider uncertain reservoir inflows. The algorithm, first described by Little (1955), maximizes the expected benefits to the system using probabilities of possible inflow realizations. It may be assumed that inflow probabilities are independent from one time period to the next; however, sequential flows are often related, and probabilities may be more accurately described when they are conditioned on some event, which is incorporated into the model by a hydrologic state variable. The hydrologic state variable provides additional information about the current condition of the system and improves the representation of future inflows by helping to maintain realistic spatial and temporal relationships from hydrologic processes. The use of hydrologic state variables in reservoir operations modelling has been shown to improve reservoir operations policies (Faber and Stedinger 2001; Stedinger et al. 1984; Tejada-Guibert et al. 1995).

The choice of a hydrological state variable depends on the characteristics of the system and the information available. Previous month's inflow and current period's inflow have been commonly used. For example, Little (1955) represented inflows by a first-order Markov chain, where the conditional probabilities of inflows in the current stage were dependent on flows realized in the previous stage. Stedinger et al. (1984) used forecasts of future flows. In areas where hydrology is dominated by the seasonal events of snowmelt accumulation and subsequent melting, the snowmelt runoff forecast has shown to be a useful indicator of future flows (Tejada-Guibert et al. 1995; Kelman et al. 1990; Faber and Stedinger 2001).

In another effort to enhance the representation of streamflow in stochastic modelling, the SSDP optimization algorithm was developed. SSDP is an extension of SDP that captures the uncertainty of inflows by considering a set of intact hydrologic scenarios simultaneously to calculate expected future benefits and to determine an optimal operating policy, where hydrologic scenarios are comprised of sets of inflow sequences and time based forecasts. In considering scenarios, the temporal and spatial correlation of flows due to seasonal hydrology is better represented than by independent inflow outcomes because the model implicitly captures the relationships that exist during an annual cycle and between watersheds without requiring a complex description of inflows (Kelman et al. 1990). Hydrologic scenarios used with SSDP may be from the historical record or synthetically generated. Kelman et al. implemented SSDP in their model of the Feather River in California using historical data (1990). Ensemble streamflow prediction (ESP) sequences have been used with the algorithm in Faber et al. (2001) and Kim et al. (2007).

The SSDP algorithm is especially appealing for use with the BC Hydro system because of the structure of requirements mandated by the Columbia River Treaty (CRT). Since the SSDP algorithm works with sequences, CRT complexities based on hydrologic conditions including treaty accounts and operational requirements are not difficult to model.

SSDP models have been historically used to make immediate policy decisions, yet it is often required that future decisions are made before the event takes place. Forecasts used in making current decisions contain information that extends into the future planning horizon, and they may be used to make decisions for future time periods. A future event's decision making process changes as the forecast evolves. The degree of improvement in expected future operations resulting from the additional information gained from an updated forecast is investigated.

This thesis describes the application of the SSDP algorithm to the BC Hydro system and explores the relationships between forecasts, inflows, and the value of water. It has been prepared to present the complete scope of the author's research by incorporating several chapters that were originally composed as stand-alone works. Chapter 2 develops the optimization model using an SSDP algorithm and applies it to the BC Hydro system. The algorithm is adapted for use with various types of inflows and forecasts. Differences in each models' calculation of value functions and policies are evaluated for several months. In Chapter 3, the analysis from the previous chapter is extended to examine the marginal value of water. Marginal values are calculated at every month for each model by simulating SSDP policies from three hydrological years of historical forecast and inflow data. Chapter 4 investigates how valuation of BC Hydro system resources evolves as forecasts are updated. A methodology is developed to calculate future policies and the expected future value of water given forecasts made in previous months. Model variants are extended to include two reservoirs. Finally, a summary is provided and conclusions relating to the complete scope of the author's research are drawn in Chapter 5.

2 Performance of an SSDP Algorithm Using Various Inflow Generation Methods

2.1 Introduction

BC Hydro is a provincial Crown corporation serving British Columbia, Canada that is mandated to generate, purchase, distribute, and sell electricity. Over ninety five percent of energy generation in British Columbia is from renewable sources including hydropower with minor thermal generation. Efficient reservoir operations management and planning is vital in ensuring BC Hydro is able to meet its goals.

A hydroelectric power system's operation policy is typically based on the results of solving an optimization problem to find optimal policies that maximize the benefit to the system. The dynamic programming (DP) technique breaks the problem into states and stages. It then finds the optimal solution in each stage and state that maximizes current benefits plus expected future benefits achievable with the optimal policies. This problem becomes more complicated when model inputs are uncertain.

Handling the uncertainty of reservoir inflows is one of the main challenges in modelling. The current decision making process for planning does not capture uncertainty in reservoir inflows, and BC Hydro is investigating the development of an improved long and medium term operations planning model that is better able to account for uncertainties in reservoir operation. Several models have been considered including a Sampling Stochastic Dynamic Programming (SSDP) method.

The SSDP optimization method is an extension of dynamic programming that captures the stochastic nature of reservoir inflows. It employs a number of sample streamflow sequences, or scenarios, and considers them all simultaneously to calculate expected future benefits and to determine an optimal operating policy (Kelman et al. 1990).

The use of streamflow sequences in SSDP is advantageous as it avoids the discretization of inflows needed in many other modelling methods, and it captures spatial and temporal correlation of annual flows from the historical record (Kelman et al. 1990). However, the selection of the type of streamflow scenario data for use with SSDP is not required to be from the historical record, and ESP traces have been used with the algorithm in Faber et al., (2001) and Kim et al., (2007). These studies show increased model performance when employing ensemble streamflow prediction (ESP) forecast scenarios for short term planning.

A hydrologic state variable may be used to describe the current condition of a watershed and account for the persistence in streamflow. Remaining seasonal runoff (Kelman et al. 1990; Faber 2001) and a combination of snow water equivalent and soil moisture (Côté et al. 2011) have been used as hydrologic state variables with SSDP.

2.2 Model Description

2.2.1 SSDP Algorithm

The SSDP formulation extends the DP algorithm. Optimal decisions that maximize expected benefits over the planning horizon are found. However, SSDP captures the uncertainty in inflows by considering the probability of transitioning between inflow scenarios in one stage to the next. The SSDP model uses a two-step algorithm; the first step (Step 1) finds an optimal decision that maximizes the value function under uncertainty, then the value function is updated by carrying out that decision on a single scenario. Step 2 is used for operations planning where current hydrological information is considered in a one stage forward re-optimization model using the value functions derived in Step 1 (Faber and Stedinger 2001; Tejada-Guibert et al. 1993).

The model is as follows:

Step 1: For each scenario *i,* and all discretized S_t, at each time t in the planning horizon:

$$f(S_t, i) = \max_{R_t} \left\{ B_t(S_t, Q_t(i), R_t) + \alpha \underset{j|i}{E} [f_{t+1}(S_{t+1}, Q_t(j))] \right\} \quad (1)$$

$$f(S_t, i) = B_t(S_t, Q_t(i), R_t) + \alpha \, [f_{t+1}(S_{t+1}, Q_t(i))] \quad (2)$$

B_t	benefit function at stage t
S_t	reservoir storage at stage t
$Q_t(i)$	inflow to reservoir at stage t for scenario *i*
R_t	power release for a given scenario *i* from reservoir at stage t
i	scenario *i*, representing flows of a particular inflow sequence
j	scenario *j*, representing flows of a particular inflow sequence
α	discount factor
$\underset{j\|i}{E}$	expectation assessed using transition probabilities of the remainder of scenario *j* starting in period *t+1* given scenario *i* in period *t*

Step 2: For Hydrologic State, *H*, and all discretized S_t at each time, t:

$$f(S_t, H) = \max_{R_t} \underset{j|H}{E}\{B_t(S_t, Q_t(j), R_t) + \alpha\, [f_{t+1}(S_{t+1}, j)]\} \quad (3)$$

H	hydrologic state
$\underset{j\|H}{E}$	expectation assessed using transition probabilities of flow from scenario *j* occurring given hydrologic state *H*

A new optimization model using a Sampling Stochastic Dynamic programming algorithm was developed to find the value of water in storage and release policy to maximize water value for the BC Hydro system. A value iteration approach is used until a steady state solution is found. The value of water is a function of immediate benefits including internal and external energy trading and potential future benefits from energy production. An outline of the optimization model is presented in the following sections.

2.2.2 Objective Function

The objective of the optimization model is to maximize the present value of water at each time step and throughout the discretized state space by maximizing revenues. Revenues are generated by selling energy to meet internal demand and exporting and importing energy to and from external markets.

2.2.3 Time Horizon

Monthly time steps are used with planning horizon of one water year.

2.2.3.1 *Decision Variables*

The decision variables of the SSDP model are the following: reservoir outflows for power generation, power generation, energy sales to external markets, energy bought from external markets, energy sold to local market to meet demand, and spills from the reservoir.

2.2.3.2 *Constraints*

The model is constrained by a number of physical and procedural bounds. BC Hydro must conduct operations to meet internal customer demand, while also complying with environmental and political constraints. Demand may be met by power generation through hydro, wind, or fossil fuels, or purchasing energy from outside of the province. Once internal demand is met through generation or imports, excess energy may be exported at market prices, which vary throughout the year.

2.2.3.3 *Mass Balance*

This constraint requires that the reservoir storage at any time step is the sum of the storage at the previous time step, current period inflows, and releases through turbines and spillways.

2.2.3.4 *Storage Constraints*

Reservoir storage must be operated within the storage limits, which are equivalent to the minimum and maximum physical storage requirements for the reservoir.

2.2.3.5 *Power Generation*

Power generation in a reservoir is a function of the reservoir's elevation. This model calculates a coefficient describing the average power generated per unit release at each starting storage state.

2.2.3.6 *Generation Limit*

Generation for each reservoir in each time step must be within the maximum and minimum generation limits.

2.2.3.7 *Load-Resource Balance*

The system load must be equal to the sum of energy generated or traded to outside markets. Energy can be bought or sold during each time step at current market prices.

2.2.3.8 *Transmission Limit*

Energy bought and sold to and from external energy markets is limited to the capacity of the tie lines transporting the power outside the province.

2.3 Hydrology in British Columbia

British Columbia has two key river basins important to hydropower generation – the Peace and Columbia River watersheds. Over seventy percent of the energy generated in the province is by hydro plant facilities on these two river systems.

Each basin has its own characteristics and behavior; however they are similar in that the hydrology is dominated by seasonal winter snow accumulation and melting. The Peace River drains its catchment in British Columbia into Alberta in the northeast while the Columbia River basin is located in the southeastern part of the province, and it drains its catchment into US territories. Snow accumulates in basin mountain ranges from late November until early April when temperatures rise and the snowpack begins to melt. The melting period, or freshet, continues through August with high flows exhibited in May or June (Eaton and Moore 2010). The timing of the freshet is dependent upon the size of snowpack and climate conditions. Flows in fall between September and November originate primarily from precipitation events. Low flows occur in winter when precipitation accumulates on the ground as snow.

Size of remaining snowpack is an indicator for the remaining flow volume to occur during the freshet. Therefore, snowpack size, along with weather information, is used to forecast the total volume of remaining flow. The forecasted volume of remaining flow contains information about previous months' flows and captures some serial correlation within a sequence, so that high inflow from snowmelt in early months is followed by low inflow from snowmelt in later months, and vice versa.

2.4 Inflow Generation and Transition Probabilities

Sampling Stochastic Dynamic Programming uses a number of intact inflow scenarios to capture the uncertainty of inflows while also recognizing temporal and spatial relationships that exist during an annual cycle and between watersheds. When used with a hydrologic state variable, such as volume forecast, it is possible to determine the probability of switching from the current scenario to another scenario in the next stage. The following sections describe inflow scenario generation and transition probability calculations for three inflow scenario types.

2.4.1 Inflow Generation

Inflow sequences used with SSDP may be from observed data or synthetically generated realizations of annual streamflow. In this study, inflow scenarios generated by three different methods, using 60 years of historical data, are employed with SSDP algorithm and model performance is evaluated.

2.4.1.1 *Historical*

Actual streamflow sequences from eleven years in the historical record, 2003-2013, were used in the Historical SSDP model. The corresponding actual seasonal volume forecast is used as the hydrologic state variable. An advantage to using observed data from the historical record is that each instance actually occurred; therefore spatial and temporal relationships within and between scenarios are perfectly correlated. However, the number of inflow sequences is limited to the data available.

2.4.1.2 *Inflow Model - AR-1 with Principal Component Approach*

Inflow scenarios used in the Inflow Model (IM) SSDP model were generated using a separate autoregressive model with lag-1 correlation together with the principal component approach. This model generates thousands of representative annual inflows to each reservoir along the Peace and Columbia Rivers where flows are correlated with seasonal volume forecast, previous inflows, and physical (Guan et al. 2013). The Inflow Model considers volume forecasts generated for the months January through April to develop annual inflow scenarios with deterministic inflows in the months December through March. For this study 10,000 scenarios were generated, and thirty-six scenarios were sampled for use with the IM SSDP model.

2.4.1.3 *Ensemble Streamflow Prediction*

Ensemble streamflow prediction (ESP) forecasts generated by BC Hydro are used with the ESP SSDP model in the last test case. BC Hydro generates twelve-month ESP forecasts at the beginning of months December through August. The ESP SSDP model is run using the most recent ESP forecast scenarios in each time step to determine optimal decisions for the current time period.

2.4.2 Transition Probability Calculation

In SSDP, each of the algorithm's two steps requires the calculation of transition probabilities. To perform the calculation of value functions for each streamflow scenario (Step 1), the model uses the probability of flow from scenario j occurring given a forecast is made equal to the forecast for scenario i. To make an optimal decision from the functions derived in step one for a given hydrologic state (Step 2), the model considers the probability of flow from scenario j occurring given the current hydrologic state which in this study is seasonal volume forecast. The transition probabilities are calculated for each model given the information in the inflow data.

2.4.2.1 *Historical*

For both steps, linear least square fitting is used with Bayes theorem to calculate the probability of a future flow given the current condition (Faber 2001; Kelman et al. 1990). For months when a forecast is not available (September – December), the transition probability is equally likely for all scenarios (1/n, where n is the number of years of inflow sequences).

2.4.2.2 *Inflow Model - AR-1 with Principal Component Approach*

Multiple linear least square fitting was used to find the conditional probability of a scenario occurring given a forecast and realized inflow for months April to August. For September through November, the probability of flow from another scenario occurring is equally likely for all scenarios. In December through March, inflow is deterministic. Step 2 transition probabilities are derived similarly.

2.4.2.3 *Ensemble Streamflow Prediction*

Transitions between ESP scenarios were not modelled in this study. Faber and Stedinger (2001) found that performance did not improve when transition probabilities were calculated among ESP traces. Therefore, the probability of flow from scenario i occurring given a forecast from scenario j has occurred is equal to one when $j = i$, and the probability is equal to 0 when $j \neq i$. Because the ESP sequences are all generated from the same current hydrologic conditions, the probability of occurrence given the current state in Step 2 is equally likely.

2.5 Results and Discussion

Four models using inflow data from different methods of inflow scenario generation were tested—three as described in Section 2.4 and one deterministic case to help illustrate a fair comparison of model performance. Table 1 summarizes the four different models tested and their methods used to calculate Step 1 and Step 2 transition probabilities where “1/n” indicates all scenarios are equally likely to occur, “FC” indicates transition probabilities were developed from forecast information, and “I” indicates no transitions occur.

Table 1. SSDP Transition Probability Calculation Methods

	Deterministic		Historical		Inflow Model		ESP	
	Step 1	Step 2	Step 1	Step 2	Step 1	Step 2	Step 1	Step 2
October	I	I	1/n	1/n	1/n	1/n	I	1/n
November	I	I	1/n	1/n	1/n	1/n	I	1/n
December	I	I	1/n	1/n	I	I	I	1/n
January	I	I	FC	FC	I	I	I	1/n
February	I	I	FC	FC	I	I	I	1/n
March	I	I	FC	FC	I	I	I	1/n
April	I	I	FC	FC	FC	FC	I	1/n
May	I	I	FC	FC	FC	FC	I	1/n
June	I	I	FC	FC	FC	FC	I	1/n
July	I	I	FC	FC	FC	FC	I	1/n
August	I	I	FC	FC	FC	FC	I	1/n
September	I	I	1/n	1/n	1/n	1/n	I	1/n

The Deterministic SSDP model uses one observed historical monthly inflow scenario and corresponding monthly hydrologic states. The preliminary results shown in this study represent model outputs where the hydrologic state is equal to the hydrologic state in the Deterministic model. In this way, it is possible to compare release decisions for a single hydrologic state vector.

In order to understand model behavior, it is necessary to discuss inflow characteristics associated with each inflow generation method case. Figure 2 shows the average and standard deviation (error bars extend one standard deviation above and below average) of monthly inflows to the Williston Reservoir on the Peace River. During fall and winter months (October through March), the averages of each case are similar. Flows from the ESP scenarios are the most similar to the observed flow. This similarity is expected since ESP scenarios contain current hydrologic information. It can be noted that the observed flows of the deterministic case are below the Hist, IM, and ESP monthly averages for most months, including May where flows are significantly below average. Low May flows are followed by higher than average flows in June. The standard deviation of the monthly inflows varies throughout the year. Both the Inflow Model and the ESP scenarios do not contain variable inflows in the winter months (December –

March), and their standard deviation is zero. Standard deviation in the historical scenarios and ESP traces in winter months are near zero as well. The historical scenarios show the highest standard deviation over all months except in May when flows generated by the Inflow Model have the highest standard deviation.

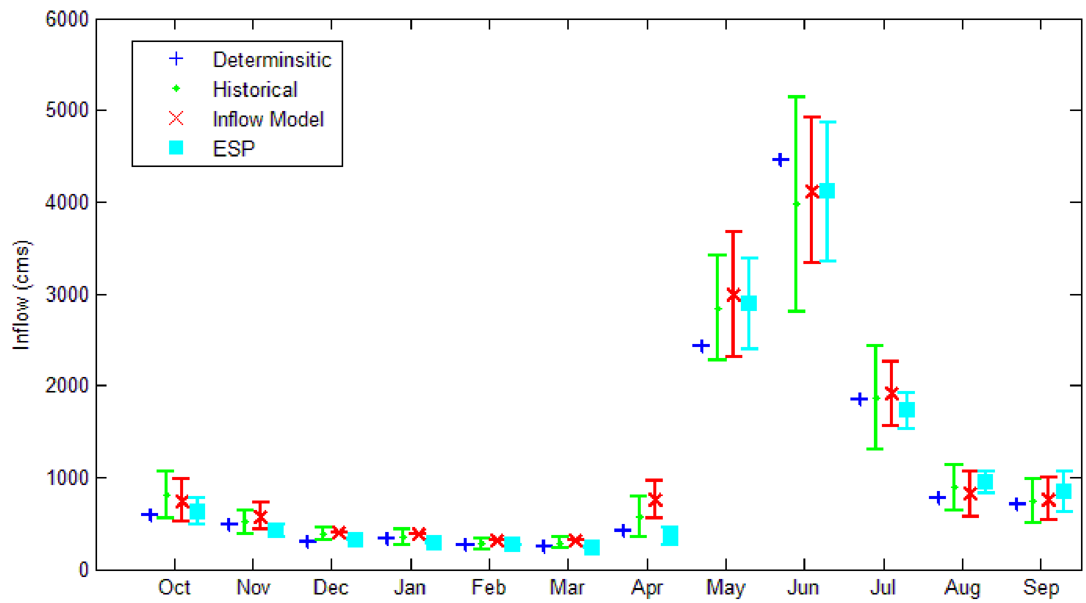

Figure 2. Average and standard deviation of monthly inflows to Williston reservoir for each inflow generation method case

.

It is possible to compare the performance of models with different inflow scenario generation methods by examining the real-time release policies given the observed hydrologic state for the deterministic case.

In general, one would expect that release policies resulting from the Historical model with the highest inflow variance and the Deterministic model with no inflow variance to be the most different, while release policies from the IM model and ESP model to fall somewhere in between. Results in Figure 3 show release policies from the historical model are the highest and policies from the Deterministic model are the lowest in July.

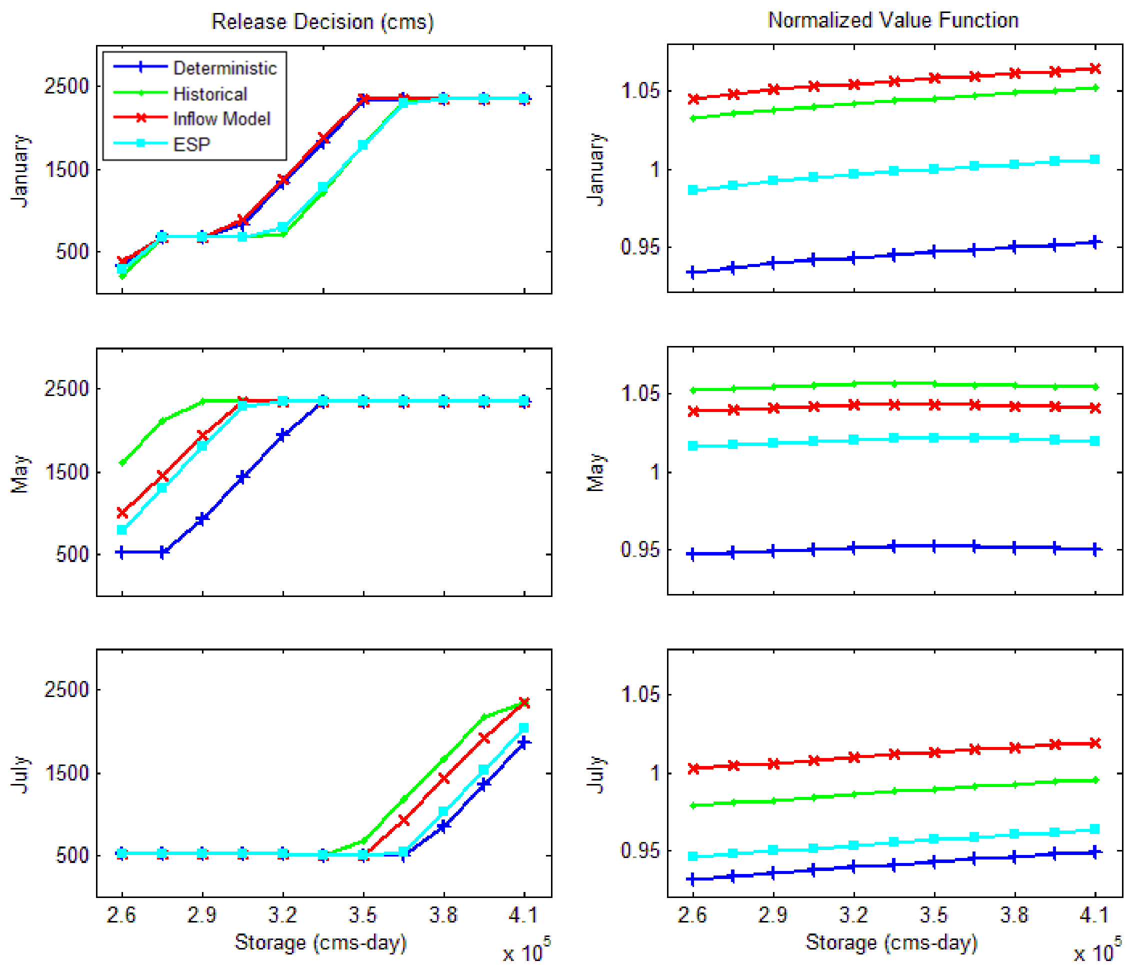

Figure 3. Comparison of release decisions and future value functions by inflow generation method for selected months.

In January, release policies from the IM case and the Deterministic model are similar while the ESP and Historical models are similar. This occurrence may be a result of calculated expected flow. In the month of January, flows from the Deterministic model and IM case only contain one value. Therefore, they are independent of the current hydrologic state. However, the Historical and the ESP models are likely affected by the current hydrologic state. A high volume forecast produced in January (83 %-tile) is not highly correlated with high flows in January (R-squared of regression equals 0.4). As a result, scenarios where January flow is relatively low are given a high probability of occurrence and the model makes low release decisions. Flows from

ESP scenarios generated for January are much lower than the deterministic flows from the base case and the Inflow Model, leading to lower release policies similar to the historical case.

High releases in May are typically unexpected as a reservoir has drafted to its lowest position and starts to refill. However, in this water year, high forecasted seasonal volume was coupled with lower than average May flows. Models that incorporated hydrologic state information were prepared for future high flows and therefore, made releases to contain these high future flows and minimize spill penalties, while the deterministic model with no knowledge of the current hydrologic state did not behave as such. Stronger relationships between the forecasted flow volume and inflow scenario volume resulted in a more aggressive release policy. The standard error of forecasted volume regressed on actual volume was the lowest for the historical inflow scenarios, and the model made the largest releases when the historical data sets were employed.

Release policies derived by the SSDP algorithm reflect the point where the tradeoff between releasing stored water to achieve current benefits and storing water to achieve future benefits is equal; therefore, the slope of the value of water function, or the marginal value of water, is of particular interest. For the months discussed, slopes of the future value functions are similar for equivalent storage. Small differences occur at low storage levels in January, and middle storage levels in May and July, which is when deviations in release policy are observed. In all cases, as the marginal value of water decreases, releases are made at lower storage levels. While the slopes are similar, the magnitudes of future value functions differ with each case and we plan to investigate this in the future.

2.6 Conclusions

The Sampling Stochastic Dynamic Programming algorithm, using several different inflow generation methods, can be successfully employed to solve the operations planning problem for the BC Hydro system. Inflow characteristics including the relationships between forecasts and realized flow and standard deviation of monthly flows affect how the model is formulated and in turn, how the model behaves. The model tended to behave with more caution as variance of flows increased. Information about the current hydrologic state of the basin (whether as a state variable or implicit in the scenarios) allowed the model to adapt to the expected flow and enabled improved decision making. The magnitudes of value functions are different for each model, although the marginal values of water were similar. Future work will include the analysis of a multi-reservoir system and statistical comparison of model performance among inflow generation methods.

3 Water Value in Reservoir Planning using Various Forecast and Inflow Generation Methods

This Chapter continues the evaluation of various inflow generation methods with the SSDP algorithm. It describes the investigation of how differences in forecasts and inflow generation methods affect the calculation of the marginal value of water.

3.1 Marginal Value of Water

When operating a hydroelectric generation system to maximize the value of resources, the fundamental problem is deciding how much water should be released from reservoir storage in each planning period. That decision depends on the tradeoff between benefits gained from the immediate release of water and the expected future benefits gained from storing water to release in the future. The benefits can be measured by the contribution of an additional unit of water to the objective function, which is called the marginal value of water (Tilmant et al. 2008). An optimal decision is made when the value of releasing that additional volume of water is equal to the value of storing the water for future use, i.e. the marginal values are equal. This is the point where the total value of system resources, the sum of the immediate and future benefits, is maximized.

When the optimization problem's state variables are reservoir storage levels, the marginal value of water in storage represents the change in water value with respect to a change of water in storage. In mathematical terms, the marginal value of water in storage corresponds to the LaGrange multipliers, or shadow prices, associated with the water storage constraints (Tilmant et al. 2008).

Knowing the marginal value of water and this point of equilibrium is valuable to reservoir operators as it is a driving force in policy decisions. The flexibility inherent to hydroelectric operations allows generating energy for export until the marginal value of water stored in reservoirs is equal to the price of trading it in the market. In contrast, operations policies dictate storing water for future trading if the current market prices are lower than the marginal value. Therefore, the investigation of marginal values with the SSDP algorithm and BC Hydro's generation system is of particular interest.

3.2 Model Description and Simulation Procedure

The SSDP models' calculation of the marginal value of water is examined at every time period. Model performance was evaluated by simulating reservoir operations using three historical data sets representing different hydrological conditions—average, dry, and wet. Marginal values of water were calculated and compared for each case.

The objective of each SSDP model is to make optimal decisions that maximize revenues using inflow and forecast data. The model formulation is explained in Chapter 2.2. Inflows and forecasts for each model are described in Chapter 2.4.1. Figure 4 provides additional detail in the description of forecasts used by illustrating forecasts available to each model at each month. Months shaded in yellow indicate no forecasts are considered, and months in green indicate no forecasts are used. No ESP forecasts are made during the months of September through November, and the model uses the most current forecasts available, which are produced in August.

	Historic	Synthetic	ESP
October			Aug ESP
November			Aug ESP
December			ESP
January	Seasonal FC	Seasonal FC	ESP
February	Seasonal FC	Seasonal FC	ESP
March	Seasonal FC	Seasonal FC	ESP
April	Seasonal FC	Seasonal FC	ESP
May	Seasonal FC	April Seasonal FC & Realized Flow	ESP
June	Seasonal FC	April Seasonal FC & Realized Flow	ESP
July	Seasonal FC	April Seasonal FC & Realized Flow	ESP
August	Seasonal FC	April Seasonal FC & Realized Flow	ESP
September			Aug ESP

Hydrologic Forecasts NOT Considered

Hydrologic Forecasts Considered

Figure 4. Monthly forecasts used with SSDP algorithm. Months shaded in yellow indicate no forecasts are considered, and months in green indicate no forecasts are used. The ESP model uses forecasts made in August for the value function calculations in September through November.

Various SSDP models can be compared by simulating the performance of each as hydrologic data is varied. The four models described in Table 1 were considered along with two additional model variations that used Historical and Synthetic inflows, but no forecasts. The procedure involved solving Eqs. (1) and (2) in Step 1 to find value functions as described in Chapter 2.2.1 for each SSDP model. Then the Eq. (3) is solved to find the expected value of water at every storage level for each month given the forecast available, H_t. The resulting future value functions are used to calculate the marginal value of water (MVW) at each month using Eq. (4), where HK is the water to power conversion coefficient.

$$MVW_t = \frac{\delta\, f(S_t, H_t)}{\delta\ S_t * HK} \qquad (4)$$

The marginal values of water are evaluated for each model. This procedure is repeated using several years of annual inflow and forecast data.

3.3 Results and Discussion

Marginal values were calculated for each of the five stochastic models described as well as the deterministic model. The latter model assumes perfect foresight and is used as a base for comparison. Table 2 lists average absolute deltas by month for each model studied. Deltas approach zero as the calculation of marginal value improves.

Table 2. Marginal Value of Water -- Average Delta from Perfect

	Hist	Syn	ESP	Hist no FC	Syn no FC
October	**0.020**	0.028	0.049	0.020	0.028
November	**0.010**	0.024	0.034	0.010	0.024
December	0.021	**0.013**	0.036	0.021	0.013
January	0.029	**0.017**	0.040	0.028	0.018
February	0.027	**0.015**	0.034	0.026	0.016
March	0.047	**0.038**	0.072	0.049	0.046
April	**0.106**	0.149	0.200	0.215	0.218
May	**0.224**	0.475	0.533	0.711	0.607
June	**0.283**	0.560	0.325	0.879	0.772
July	0.024	0.052	**0.019**	0.050	0.053
August	0.020	0.039	**0.013**	0.024	0.036
September	**0.022**	0.036	0.033	0.022	0.036

During fall and early winter months when inflows are low and primarily results of precipitation events, the marginal value of water did not vary drastically among model variants. The average delta from perfect remained below five percent for all models. Since no forecasts were available in the months of September through December, there was no difference in the results of the Hist and Hist no FC models. The same is true for the Syn and Syn no FC models. The ESP model uses the latest forecast available for its calculation, and performs worse on average than the other models.

Forecasts made during the months of January through March did not greatly improve model performance over using no forecasts at all. The Syn model output the best results during the winter months of December through March. This result is primarily due to the overestimation of the marginal value of water by other models during the simulation of wet hydrology, while models using synthetic inflows were not affected (uncertainty is not considered by Syn models in these months (see Section 2.4.1.2)). It is difficult to conclude that it is good practice to ignore uncertainty during winter months, but it is clear that forecasts do not provide great additional value in calculating immediate water value in January and February. This is because there is little variation in operations during these months as high demand almost always requires turbines to be run at peak capacity.

However during the freshet, models that did consider forecasts consistently perform better than models that did not. The largest deltas in the calculation of the immediate marginal value occurred in April, May, and June. In each of these months, the historical model produced the best approximation of the marginal value of water. The ESP model produced the best results for the late freshet period in July and August when much of the seasonal flows have been realized and forecasters have high confidence.

Figure 5 shows the percentage difference from the deterministic marginal value calculation for the three hydrologic water years simulated (from top to bottom—average, dry, and wet). Deltas for each model and simulated year are displayed by month and approach zero as the calculation of marginal value improves. When the hydrologic simulations are considered separately, the largest deltas in the calculation of marginal value occurred in April, May, and June which is consistent with the average results discussed previously.

It is interesting that during these months, the simulation of wet hydrology resulted in the overestimation of the marginal value of water, while in the dry simulation, marginal value was underestimated; however, this result is attributed to the calculation of transition probabilities of this particular data set and should not be considered a rule. Because the transition probability calculation method is based on forecasts, scenario-to-scenario transition probabilities may be high when scenario inflows and water values are quite different. In this simulation, the dry and wet simulated years were extremes in hydrologic conditions (8th and 83rd percentile), yet the calculated expectation of flow was less extreme. The fact that this occurred for both the dry and wet simulations is coincidence.

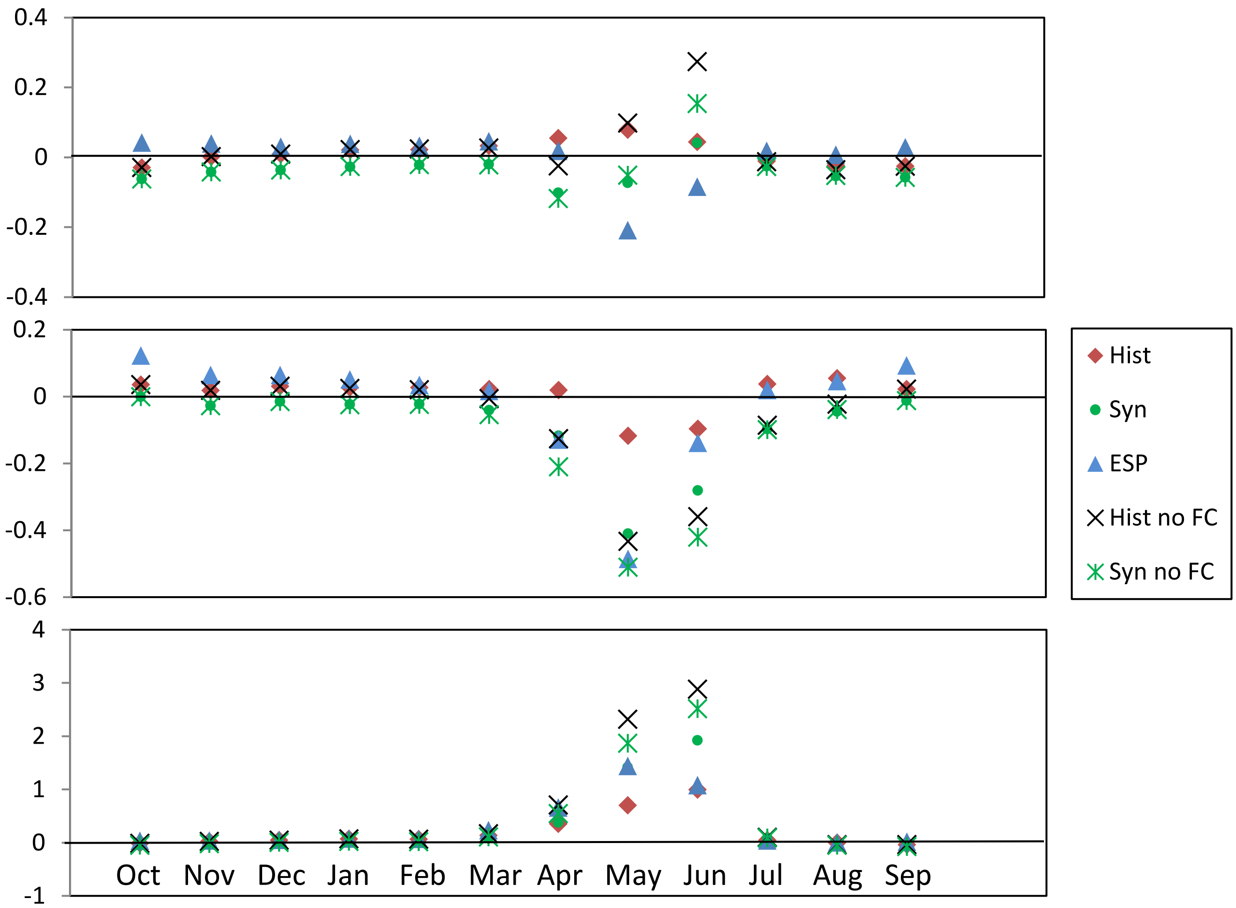

Figure 5. Comparison of marginal value of water (percent delta from perfect) by SSDP model variants for three hydrological water years—average (top), dry (middle) and wet (bottom).

3.4 Conclusions

At any time period, the value of water in a system may be approximated by an optimization model using the SSDP algorithm. A description of the change in water value with change in the amount of water in storage is valuable for reservoir operations planning and influences release policy decisions. This Chapter focused on finding the marginal value of water for the current time period given a most recent forecast. Results indicated that the value of forecasts in approximating the immediate marginal value of water in fall and winter months was relatively small due to lack of forecast information (fall) and firm winter operations strategies, but increased into the freshet as forecasts improved and operations became more flexible. In April through June, the model using historical inflow and forecast scenarios resulted in MVW approximations closest to perfect, while ESP forecasts performed best in July and August.

This study limits the approximation of the marginal value of water to the current time period, immediately following the forecast. However, understanding how forecasts affect the approximation of the value of water into future months is beneficial to reservoir operators when policy decisions may occur more than one month in advance. For example, forecasts made in winter that indicate low summer marginal values may result in operations decisions to increase releases during the spring. The following chapter investigates how future marginal values and reservoir operations are influenced by the evolution of forecasts leading into the spring freshet.

4 The Evolution of Forecasts and Water Value in Reservoir Planning Using an SSDP Algorithm

4.1 Introduction

A hydroelectric power system's operation policy is typically based on the results of solving an optimization problem where the objective is to maximize system benefits by making decisions to store and release water over the planning horizon. The problem becomes more complicated when model inputs, such as future streamflows are uncertain.

Handling the uncertainty of streamflows that contribute to reservoir inflows is one of the main challenges in modelling. Yet when uncertainty is considered, models are better able to provide estimates of the value of system resources and provide better projections of expected values of revenues, energy generation, and market transactions given potential future conditions (Abdalla et al. 2013). The sampling stochastic dynamic programming (SSDP) optimization method is an extension of dynamic programming that uses a number of intact inflow scenarios to capture the uncertainty of inflows. In considering inflow sequences, the model is also able to recognize temporal and spatial relationships that exist during an annual cycle and between watersheds (Kelman et al. 1990).

Stochastic model performance may be improved when hydrologic information about the current state of a system is provided to the model in the form of forecasts. Hydrologic forecasting improves the model's understanding of future inflow uncertainties, and its use in reservoir operations modelling has been widely shown to improve reservoir operations policies (Faber and Stedinger 2001; Stedinger et al. 1984; Tejada-Guibert et al. 1995).

Forecasts have been used with dynamic programming models to make immediate policy decisions; however, information about the future contained in forecasts extends farther into the planning horizon than just the immediate time period and may be used to make not only immediate decisions, but also decisions in future time periods. Often in reservoir planning, decisions about future events must be made in advance of the actual event occurring due to planning expectations or policy requirements. For this reason, it is useful in decision making to understand the degree of improvement in expected future operations resulting from additional information gained in the form of an updated forecast. This paper investigates how valuation of BC Hydro system resources evolves as forecasts are updated using the SSDP algorithm.

4.1.1 Reservoir Inflow Generation Methods

Reservoir inflows play a large part in determining operations policies, so in modelling a hydroelectric system, particular attention must be given to the streamflow data supplied to the model. Inflows within a system are typically spatially and temporally correlated, and the timing and volume of reservoir inflows are an important, yet these characteristics usually are not known in advance. A major challenge is providing quality data to the model that captures spatial and temporal relationships while also describing uncertainties in hydrologic conditions.

Reservoir inflows are a stochastic process that may be approximated by a distribution. However, many decision models cannot be solved using continuous representations of inflow distributions. Therefore, it is often necessary to represent inflows as discrete random outcomes of the process. Computational capability limits the number of discretized random outcomes that may be employed, and inflows may be represented by sets of random outcomes in the form of scenario trees (Kaut and Wallace 2003) or sequences.

Numerous modelling methods capture uncertainty using sets of random outcomes including stochastic dynamic programming (SDP) (Little 1955), stochastic dual dynamic programming (SDDP) (Pereira and Pinto 1991), and sampling stochastic dynamic programming (SSDP) (Kelman et al. 1990). In SSDP a set of intact inflow sequences are considered simultaneously. Inflow sequences (or scenarios) used with SSDP are realizations of annual inflows, observed or otherwise synthetically generated (Kelman et al. 1990).

Generally, inflow scenarios used by multi stage stochastic models are generated using the same basic procedure (Di Domenica et al. 2007). Historical data is evaluated to choose an appropriate model and calibrate model parameters. The model generates spatially and temporally correlated streamflow scenarios (often in the form of a scenario tree or sequences), and then those scenarios are sampled to build the data set for the decision model. There are many methods that have been used to generate streamflow scenarios. The type of methodology chosen depends on the natural hydrologic behavior of the system.

When there is no processing of the historical data set by a scenario generation model, the outcome is the set of intact historical inflow scenarios. The quantity of scenarios is limited by the data in the historical record, yet an advantage of using this data set for the decision model is that potential error introduced as a result of statistical manipulation is eliminated—inflows within scenarios are perfectly correlated since they actually occurred.

A scenario generation model may incorporate current hydrological conditions to produce forecasts of future inflow sequences that are conditioned on current state of the catchment, as called ensemble streamflow prediction (ESP) traces. The use of forecasts is discussed in Section 4.1.2.

4.1.2 Forecasts as Hydrologic State Variables

Information about the current state of a system can be used by models to improve reservoir operations policies (Faber and Stedinger 2001; Tejada-Guibert et al. 1995; Stedinger et al. 1984; Côté et al. 2011). As inflows are a product of the hydrological process and thus serially and spatially correlated, information about the current condition of the hydrologic process can be used in describing future flows into reservoirs through the calculation of conditional probabilities of inflows occurring in the future.

The choice of a hydrological state variable depends on the characteristics of the system and the information available. Previous month's inflow and current period's inflow are common choices. Stedinger et al. (1984) showed that the best forecast of the current period's inflow showed improvements in reservoir operation policies over using previous month's inflow. In areas where hydrology is dominated by the seasonal events of snowmelt accumulation and subsequent melting, the snowmelt runoff forecast has shown to be a useful indicator of future flows (Tejada-Guibert et al. 1995; Kelman et al. 1990; Faber and Stedinger 2001).

Many snowmelt runoff models have been developed. In most models, snowmelt runoff is forecasted by simulating the snow accumulation and melting process then routing the runoff to streams and reservoirs (World Meteorological Organization 1986). Often the outputs of snowmelt runoff models are deterministic outcomes; however, because of the inherent uncertainties in forecasting, each forecasted outcome must be made with an associated level of certainty. When this level of certainty is quantified, the user becomes aware that a forecast is evolving over time as the level of certainty increases or decreases even though the deterministic forecast may not change. A probability distribution is required to fully express certainty for continuous variables such as snowmelt runoff (Krzysztofowicz 2001). For example, a forecast may provide values for mean and standard deviation assuming a normal distribution, where the standard deviation represents the level of certainty of the forecast. When forecasting snowmelt runoff and reservoir inflows, one would expect the confidence in forecasts to increase as the snowmelt season progresses since knowledge of past months' runoff can be considered. This is especially true in the Pacific Northwest, where during the middle and end of the snowmelt season, there is little precipitation and most of a reservoir's inflow originates in mountain snowpack (Eaton and Moore 2010).

One method of handling hydrologic forecasting uncertainty is by generating an ensemble of hydrographs representing possible realizations of future flow. These forecasts consist of a set of possible future inflow sequences that are developed using current conditions and historical meteorological conditions (Day 1985). ESP forecasts are particularly attractive for use with SSDP since the algorithm uses intact sequences in its calculations. ESP forecasts have been used with the SSDP algorithm in Faber and Stedinger (2001) and Kim et al., (2007). These studies show improved performance when employing ESP streamflow forecast scenarios for short term planning with an SSDP algorithm. In this work, we investigate the evolution of ESP streamflow forecasts leading up to the freshet in reservoir operations planning.

4.1.3 Marginal Value of Water

When operating a hydroelectric generation system to maximize the value of resources, the fundamental problem is deciding how much water should be released from reservoir storage in each planning period. The decision driver is the marginal value of water, which is the incremental benefit associated with the change of the amount of water in reservoir storage (Tilmant et al. 2008). An optimal decision is made when the value of releasing an additional volume of water is equal to the value of storing water for future use, i.e. the marginal values are equal. This is the point where the total value of system resources—the sum of the immediate and future benefits is maximized.

Knowing the marginal value of water is valuable to reservoir operators as it is a driving force in policy decisions. When deciding whether to release or store water, the marginal value is compared to current market prices. For example, water may be released when its marginal value is low or stored to be released in the future when its marginal value is above current market value. Therefore, this study examines the marginal value of water with various implementations of the SSDP algorithm.

The paper is organized as follows: Sections 4.2 details the SSDP algorithm. Section 4.3 contains the results of a case study of the BC Hydro system including a description of the system and application of the SSDP model. Conclusions are drawn in Section 4.4.

4.2 Methods in Dynamic Programming

4.2.1 Stochastic Dynamic Programming

The reservoir operations planning problem is often solved using the dynamic programming technique. In dynamic programming, a multistage planning problem is broken into a series of smaller one stage problems that are solved successively. The problem is described in each stage by the state of the system which is represented by reservoir storage, S_t, and a hydrologic state, H_t. The algorithm optimizes a decision, release of water, R_t, that maximizes the sum of current benefits to the system and expected future benefits achieved when making that decision, R_t (Bellman 1957). The recursive equation is solved at every stage and state of the problem starting at the last stage and moving backwards in time. When flows, Q_t, are unknown, the conditional expectation of current and future benefits is calculated using a stochastic dynamic programming methodology as in Eq. (5). Future benefits depend on the state of the system in the next stage, S_{t+1} and depends on the decision made, R_t, inflows to the reservoir, Q_t, the starting storage state, S_t, and evaporation from the surface of the reservoir, $e(S_t, S_{t+1})$ shown in Eq. (6).

$$f_t(S_t, H_t) = \max_{R_t} \underset{Q_t|H_t}{E} \{B_t(S_t, Q_t, R_t) + \alpha \, [f_{t+1}(S_{t+1}, Q_t)]\} \tag{5}$$

$$S_{t+1} = S_t + Q_t - R_t - \mathrm{e}(S_\mathrm{t}, S_{\mathrm{t+1}}) \tag{6}$$

If current benefits do not depend on current flows, then the expectation is dropped from the first term and Eq. (5) becomes Eq. (7). This is a reasonable assumption in a monthly model since operation decisions are modified throughout the month as information about realized inflow becomes available (Stedinger et al. 1984) and short term forecasts are updated.

$$f_t(S_t, H_t) = \max_{R_t} \left\{B_t(S_t, Q_t, R_t) + \alpha \underset{Q_t|H_t}{E} [f_{t+1}(S_{t+1}, Q_t)]\right\} \tag{7}$$

4.2.2 Sampling Stochastic Dynamic Programming

The SSDP formulation extends the SDP algorithm. Like SDP, an optimal decision is selected for each stage and state combination that maximizes the expected future benefits of making that decision. The hydrologic state variable used in SSDP is an individual streamflow scenario, i. However, after a decision, R_t, is optimized in each sub-problem, SSDP undertakes an additional step to update the future value function that reflects the value of making a release decision, R_t, on the scenario i. The model developed by Faber and Stedinger (2001) is shown in Eq. (8) and Eq. (9).

For each scenario *i* and all discretized S_t at each time t in the planning horizon:

$$\max_{R_t} \left\{B_t(S_t, Q_t(i), R_t) + \alpha \underset{j|i}{E} [f_{t+1}(S_{t+1}, Q_t(j))]\right\} \tag{8}$$

$$f(S_t, i) = B_t(S_t, Q_t(i), R_t) + \alpha \, [f_{t+1}(S_{t+1}, Q_t(i))] \tag{9}$$

B_t	benefit function at stage t
S_t	reservoir storage at stage t
$Q_t(i)$	inflow to reservoir at stage t for scenario *i*
R_t	power release for a given scenario *i* from reservoir at stage t

i scenario *i*, representing flows of a particular inflow sequence

j scenario *j*, representing flows of a particular inflow sequence

α discount factor

$\underset{j|i}{E}$ expectation assessed using transition probabilities of the remainder of scenario *j* starting in period *t+1* given scenario *i* in period *t*

Calculation of the expected future benefits in Eq. (8) requires the conditional probability, $P_t(j|i)$, describing the likelihood of the remainder of scenario j occurring in period t +1 following scenario i in period t. If no transitions occur, $P_t(j|i)$ is expressed by Eq. (10). In this case, future flows are assumed known, and the problem reduces to a deterministic dynamic program.

$$P_t(j|i) = \begin{cases} 1, & j = i \\ 0, & j \neq i \end{cases} \tag{10}$$

When uncertainty exists, but information about the future is not used to in the calculation of the conditional probability, all scenarios are equally likely to occur and $P_t(j|i)$ is expressed in Eq. (11), were N is the number of scenarios considered.

$$P_t(j|i) = \frac{1}{N} \tag{11}$$

For cases where a hydrologic state variable is useful indicator of future flows, Kelman et al. (1990) developed a methodology to calculate the conditional state to state transition probabilities using Bayes Theorem:

$$P_t(y_{t+1}(j)|H_t) = \frac{P_t(H|y_t(i))p(i)}{\sum_{n=1}^{N} P_t(H|y_t(j))p(j)} \tag{12}$$

Where y is defined as sum of actual flows, Q, between the current time, t, and the end of the snowmelt season for the inflow scenario j, and H_t is the forecasted snowmelt runoff between the current time and end of snowmelt season. The probability $P_t(H|y_t(i))$ is found by regressing forecasted seasonal flow on actual seasonal flow. It is assumed that the probability distribution is normal around the calculated expected actual flow. The standard deviation is equal to the standard error of the regression. When it is assumed that the probability of scenario j following scenario i is equivalent to the probability of scenario j following the forecast $H_t(i)$ from scenario i, Eq. (13) is substituted into Eq. (12), and this Bayesian approach may be used to find the conditional scenario to scenario transition probabilities, $P_t(j|i)$.

$$P_t(j|i) = P_t(y_{t+1}(j)|H_t(i)) \quad (13)$$

A simple extension is made in the application to a multi-reservoir model as suggested by Faber (2001) where a scenario is considered the array of inflows and forecasts for the set of reservoirs. The transition from i to j is calculated where y(j) is the sum of flows into all reservoirs, k = 1..K, and H(i) is the sum of forecasted seasonal flows from all reservoirs, k = 1..K. This application is valid for the BC Hydro system since for each scenario, flows at both reservoirs are based on coincident historical weather and retain spatial correlation.

In the calculation of one-reservoir model value functions, Faber and Stedinger (2001) found no significant difference in model performance between using the more sophisticated transition probability calculation method and not modelling transitions at all. They judged this was because streamflow persistence was effectively captured with the single trace, and uncertainty was captured in the re-optimization procedure for real time decision making (Section 4.2.4).

4.2.3 Value Function Approximation

The future value of the system is required to solve Eqs. (8) and (9). Because the function describing the future value is typically not known, the problem is solved at discretized values of the continuous variables S_t and i throughout the state space. The solution to the problem may not fall on a pre-defined grid point in the discretized state space, so the future value must be approximated. Many methods of value function approximation have been used with SDP including multidimensional linear, polynomial, and spline interpolation as described in Johnson et al. (1993). A multi-linear interpolation method using a convex hull algorithm was used in this multi-reservoir model and is discussed with the case study in Section 4.3.1.1.

4.2.4 Re-optimizing to Calculate Future Marginal Value of Water

The outcome of Eq. (8) yields the optimal decision and the optimized value of the system at specific grid points in the state space. The actual state of the system seldom falls on one of these discrete points, and the operator must find an optimal decision given current conditions. This may be accomplished by interpolating between solutions, $f(S_t, i)$, or using the set of solutions in a re-optimization procedure. Tejada-Guibert et al. (1993) showed higher estimated average annual benefits are achieved when decisions are made through a re-optimization procedure compared to benefits from interpolated decisions.

Faber and Stedinger (2001) applied this methodology with SSDP to implement real-time policies by solving for optimized release decision, R_t, using Eq. (14) with the expectation using the probability, $P_t[\text{scenario } j|H_t]$.

For the current hydrologic state, H, and current S_t at time, t:

$$f(S_t, H) = \max_{R_t} \underset{j|H}{E} \{B_t(S_t, Q_t(j), R_t) + \alpha\, [f_{t+1}(S_{t+1}, j)]\} \tag{14}$$

H	hydrologic state
$\underset{j\|H}{E}$	expectation assessed using transition probabilities of flow from scenario *j* occurring given hydrologic state *H*

The methodology can be extended to solve for an optimal release for future time periods, f where f = t + x, in Eq. (15). In this case, the expectation is assessed using the probability of an inflow scenario, j, occurring in a future time period, f, given the current hydrologic condition, H_t which is represented as $P_f(\text{scenario } j|H_f^t)$.

$$f\left(S_f, H_f^t\right) = \max_{R_t} \underset{j|H_f^t}{E} \left\{B_f\left(S_f, Q_f(j), R_f\right) + \alpha\left[f_{f+1}(S_{f+1}, j)\right]\right\} \tag{15}$$

The probability $P_f(\text{scenario } j|H_f^t)$ can be found using methodology similar to that used for the calculation of the scenario to scenario transition probabilities required in the calculation of the value function (Eq. (16)). In this case, $y_f(j)$, is defined as sum of actual flows from scenario j between time, f and the end of the freshet in August. H_f^t is the sum of forecasted flows occurring between time, f, and the end of the freshet from a forecast made in time, t.

$$P_f(j|H_f^t) = P_t(y_f(j)|H_f^t)) \tag{16}$$

After the optimal decision is found using Eq. (15), the value of optimized solution, $f(S_f, H_f^t)$, represents the expected future value of water in the system at time *f*.

This model is different than Faber and Stedinger (2001) and Tejada-Guilbert et al. (1993) in that we expand the re-optimization procedure to determine not only immediate optimal release policies, but also release policies in future stages and subsequently the future value of water. This formulation will allow us to investigate and assess the dynamics of water in time as forecasts are updated.

4.3 Case Study

4.3.1 Description of BC Hydro System and Model

BC Hydro is a provincial Crown corporation serving British Columbia, Canada that is mandated to generate power to meet the domestic load and to purchase, distribute, and sell electricity. Over ninety percent of energy generated in British Columbia is from renewable sources including hydropower.

BC Hydro's transmission network is linked with the province of Alberta and the western United States. This allows for the exchange of energy over a large market. BC Hydro is able to take advantage of the flexibility offered with a primarily hydroelectric generation system where it is able to store energy and purchase electricity from the market when prices are low and then generate energy to sell to the market when prices are high.

British Columbia has two key river basins important to hydropower generation – the Peace and Columbia River watersheds. Over seventy percent of the energy generated in the province is by hydro plant facilities on these two river systems, and virtually all of that potential energy is stored and generated at the Williston Reservoir on the Peace River and the Kinbasket Reservoir on the Columbia River System. Table 3. Reservior Characteristics provides a description of the reservoirs modeled.

Table 3. Reservior Characteristics

	Williston	Kinbasket
Generation Station	Gordon M. Shrum (GMS)	Mica (MCA)
Reservoir capacity (km^3)	74	24
Dam height (m)	186	240
Generation capacity (MW)	2,876	1,805

While the reservoirs are located in separate basins, they are similar in that the hydrology is dominated by seasonal winter snow accumulation and melting. The Peace River drains its catchment in British Columbia into Alberta in the northeast while the Columbia River basin is located in the southeastern part of the province, and it drains its catchment into US territories. Snow accumulates in basin mountain ranges from late November until early April when temperatures rise and the snowpack begins to melt. The melting period, or freshet, continues through August with high flows exhibited in May or June (Eaton and Moore 2010). The timing of the freshet is dependent upon the size of snowpack and climate conditions. Flows in fall between September and November originate primarily from precipitation events. Low flows occur in winter when precipitation accumulates on the ground as snow.

Size of remaining snowpack is an indicator for the remaining flow volume to occur during the freshet. Therefore, snowpack size, along with weather information, is used to forecast the total volume of remaining flow over the snowmelt season. Forecasted seasonal runoff volume contains information about previous months' flows and captures some serial correlation within a sequence, so that high inflow from snowmelt in early months is followed by low inflow from snowmelt in later months, and vice versa making it an attractive choice for use as a hydrologic state variable in the BC Hydro operations model.

BC Hydro's operating goals include maximizing the value of system resources while meeting customer demand. The marginal value of water in storage drives reservoir operations decisions, and therefore, having an accurate estimate of this value is helpful to BC Hydro in its evaluation of trade-offs between releasing water to achieve immediate benefits and storing water to gain benefits in the future.

An important activity of BC Hydro reservoir operations is the timing of drawing down the reservoirs in the winter and spring to prepare for filling during the freshet. The drawdown schedule is typically made between the months of February and April and depends heavily on estimates of marginal water values in February through June.

A new optimization model using a Sampling Stochastic Dynamic programming algorithm was developed to find the value of water in storage and release policies to maximize water value for the BC Hydro system. The SSDP model consists of Eqs. (8), (9), and (15), with an objective to maximize the present value of water at each monthly time step in the planning horizon and throughout the discretized state space. The stage problem was solved by a linear programming model which allowed the algorithm to search for the optimal decision, R_t, over a continuous interval, rather using a than using traditional DP search loop requiring the discretization of release decisions. Once the problem was solved at every stage in the planning horizon, a value iteration approach was used to reach a steady state solution. The model was formulated in AMPL and sub problems were solved by the Cplex solver.

In this model, the value of water is a function of immediate benefits including internal and external energy trading and potential future benefits from energy production. Revenues are generated by selling energy to meet internal demand and exporting and importing energy to and from external markets (Eq. (17)).

$$B_t(S_t, Q_t(j), R_t) = a_t * Generation_{(Load)t} \pm b_t * Trade_t \tag{17}$$

Where $\mathrm{Generation}_{(\mathrm{Load})t}$ is energy generated from release, R_t, used to meet demand at stage t, $\mathrm{Trade}_\mathrm{t}$ are the imports and exports exchanged at stage t, a_t are the benefits of meeting the domestic load, and b_t are prices at which energy is exchanged at stage t.

Model decisions are constrained by a number of physical and procedural bounds. BC Hydro must conduct operations to meet internal customer demand while also complying with environmental and other non-power constraints. The decision variables of the SSDP model are the following: reservoir outflows for power generation, power generation, energy sales to external markets, energy bought from external markets, energy sold to local market to meet demand, and spills from the reservoir.

Conservation of mass requires that the storage in each reservoir at any time step is the sum of the storage at the previous time step, current period inflows, and releases similar to as shown in Eq. (6); however, this model separates releases through turbines and releases to spillways. It is assumed that storage losses from evaporation are zero since in Canada rainfall on the reservoir generally replaces any evaporative losses.

$$S_{t+1} = S_t + Q_t - R_t - Spill_t - \mathrm{e}(S_\mathrm{t}, S_\mathrm{t+1}) \tag{18}$$

Where $\mathrm{R_t}$ is power release at stage t, and $\mathrm{Spill_t}$ is non-power release at stage t. Each reservoir must be operated within its storage limits which are equivalent to the minimum and maximum physical storage requirements for the reservoir.

$$S_{min} < S_t < S_{max} \tag{19}$$

Demand may be met by power generation through hydro, wind, fossil fuels, or purchasing energy from outside of the province. Once internal demand is met through generation or imports, excess energy may be exported at market prices which vary throughout the year. The system load must be equal to the sum of energy generated and traded to outside markets. Energy may be bought or sold during each time step at current market prices. Prices and demand fluctuate throughout the year with seasonal weather patterns and are shown in Figure 6.

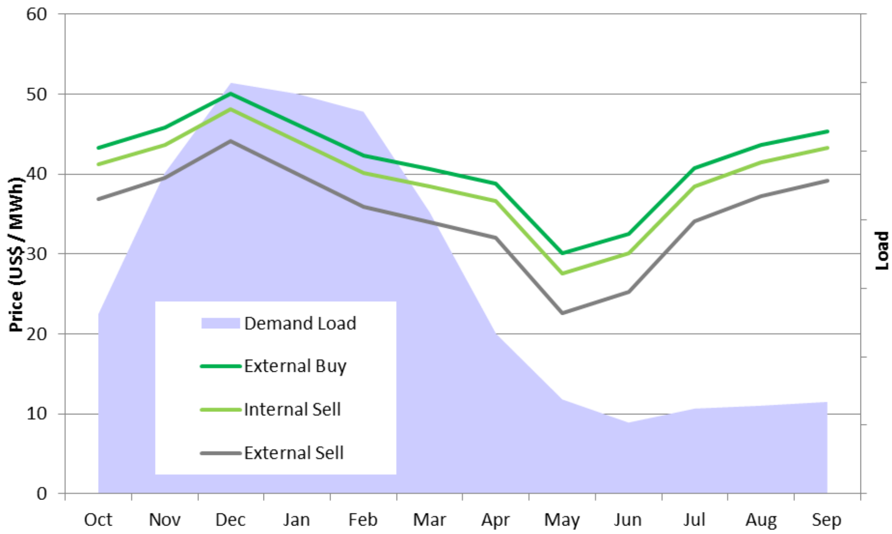

Figure 6. Economic inputs to model. Prices of trading are seasonal and shown on the left axis. Demand load is represented by the shaded region and valued on the right axis.

$$Load_t = Generation_t \pm Trade_t \tag{20}$$

Load_t is the net system energy demand at stage t, and $\mathrm{Generation}_t$ is energy generated from turbine releases, R_t, at stage t.

Energy generation in a reservoir is a function of the reservoir's elevation. A coefficient is calculated describing the average power generated per unit release at each starting storage state.

$$Generation_t(R_t, S_t) = HK(S)R_t \tag{21}$$

Where HK(S) is a generation coefficient for starting state, S. Generation for each reservoir in each time step must be within the maximum and minimum limits.

$$R_{min} < R_t < R_{max} \tag{22}$$

Energy bought and sold to and from external energy markets is limited to the capacity of the transmission system.

$$0 < Trade_t < Trans_{max} \quad (23)$$

$Trans_{max}$ is the maximum energy that may be transmitted in time t.

4.3.1.1 *Value Function Approximation with Convex Hull Algorithm*

As discussed in Section 4.2.3, the function, $f(S_t, i)$ is solved for a set of defined points. Therefore when solving Eqs. (8) and (9), the future value function, $f_{t+1}(S_t, i)$, must be approximated. If we assume that the value function is convex, then the use of the mathematical concept of a convex hull may be used to approximate the value function.

A convex hull for a set of points is defined as the smallest convex polygon that contains all points in the set (Cormen 2001). There are many algorithms that may be used to find the convex hull for a set of points. The QuickHull algorithm is used in this study (Barber et al. 1996). Qhull software implements the QuickHull algorithm and outputs the set of hyperplanes (or facets) that comprise the convex hull. Convex hull approximation methodologies have been applied to hydropower planning models by Dias et al. (2010) where the cost to go function of two reservoirs in a cascade system were modeled using the convex hull algorithm, and a four dimensional model of hydropower generation function has been developed using the convex hull algorithm (Diniz and Maceira 2008).

The facets comprising a convex hull may be used to approximate the future value function. For this maximization problem, the upper limits of the value function are constrained by the facets of convex hull. For example, in a two-reservoir problem, the constraint corresponding to a particular facet is shown in Eq. (24):

$$V \leq \text{a}\, S_1 + \text{b}\, S_2 + \text{c} \quad (24)$$

Where V is the value of water in storage, S_1 and S_2 are reservoir storage levels, and a, b, and c correspond to the plane equation of the facet.

This methodology makes it possible to solve the problem using linear programming at each stage. Figure 7 shows an example of how facets generated by the convex hull are used as linear constraints. As the dynamic programming problem is solved recursively, the value function approximation using the convex hull is updated. In each stage, the optimization problem is solved for every point in the discretized state space. The set of points consisting of the reservoir storage levels of each reservoir and the calculated present value are used with the convex hull algorithm to generate the convex hull whose facets are used to approximate the future value function in the previous time period. The facets of the convex hull completely surround the set of points; however, in order to solve the maximization problem, only the facets that make up the upper portion of the hull are used in the approximation.

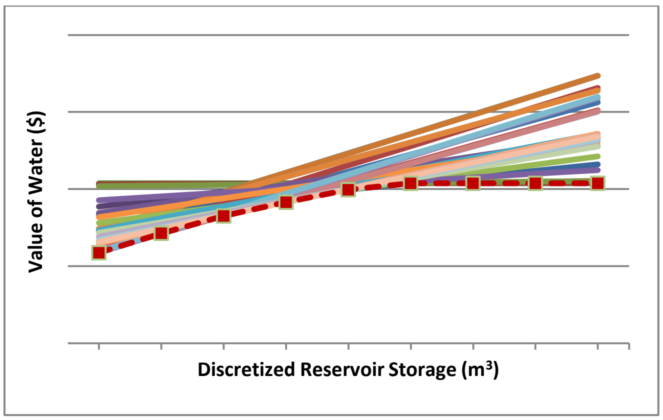

Figure 7. Two dimensional representation of facets of a convex hull used to approximate the value function. The values at the discretized grid points are represented by square points and the piecewise linear approximation is shown by the dotted line.

4.3.2 Model Variations and Cases Studied

Several models were created making use of different forecasting methods to gain a more comprehensive understanding of the benefits in updating. The following sections describe the cases evaluated and details of each model including forecast method, inflow data, and transition probability calculation.

Three model variations were studied: SSDP with no forecasts, SSDP with Historical Seasonal Runoff Volume forecasts, and SSDP with ESP forecasts. The SSDP model with no forecasts and the SSDP model with Historical Seasonal forecasts used actual streamflow sequences from eleven years in the historical record, 2003-2013. For the remainder of the paper, we refer to these models as *Hist noFC* and *Hist FC* models, respectively. The SSDP with ESP forecasts uses inflows contained in the scenarios from ensemble streamflow prediction (ESP) forecasts generated by BC Hydro and is referred to as the *ESP* model. These twelve-month inflow sequences are renewed at the beginning of months December through August.

Transition probabilities are required in the calculation of the value function (Eq. (8)) and also in re-optimization (Eq. (15)) to find optimal current and future releases and values. To perform the calculation of value functions for each streamflow scenario, the model uses the probability that flow from scenario j occurs following flow from scenario i. This probability may be calculated in one of three ways described in Eqs. (10), (11), and (12) for each time step depending on the uncertainty and the availability of forecasts. In this study, scenario to scenario transition probabilities used in the value function calculation are represented by Eq. (10) for all cases and at all time periods. This relationship captures the persistence of flows in a scenario, and simplifies the problem by eliminating the probability of switching to another scenario in a future time period, effectively assuming that future flows are known. Models making this simplification capture uncertainty in re-optimization for making future releases.

To make optimal current and future decisions by re-optimizing using the derived value functions, the model considers the probability of flow from scenario j occurring at a future time, given the current hydrologic state, H_t. The probability of a future flow given the current condition for the historical inflow scenarios were calculated using linear least square fitting and Bayes theorem as described in Eq. (12).

Assigning equally likely occurrence probabilities to ensemble streamflow prediction members is a common choice and has been used to represent ESP transition probabilities with SSDP (Faber and Stedinger 2001). More recent studies have developed new methodologies to assign weights to forecasted scenarios to improve descriptions of scenario probabilities using information known about current hydrologic or climactic patterns. These new methods are described and reviewed in Stedinger and Kim (2010). Weighting methods were not used in this study, and we assume that each member of an ESP forecast is equally likely to occur, and it follows that the transition probability for the *ESP* model, P[H|j], equals 1/n.

When no forecast is available as in the *Hist no FC* model, no new information can be provided to the model and it is assumed that all scenarios have equal probability of occurring in the future yielding, P[H|j] = 1/n.

A summary of the models investigated and details about is presented Table 4.

Table 4. Variations of Models Investigated

Inflow/ Forecast	Transition Probabilities P[j|i]	P[j|H]	Months of Forecast Updating
Hist no FC	I	1/n	none
Hist FC	I	B	Feb/Mar/Apr
ESP	I	1/n	Feb/Mar/Apr

4.3.3 Marginal Water Value Simulation Procedure

We examined the SSDP models' calculation of the marginal value of water with different forecast methods over time. Model performance was evaluated by simulating reservoir operations using historical data sets. Future marginal values of water were calculated and compared.

The simulation procedure involves solving Eqs. (8) and (9) for different variations of the SSDP model in Table 4 to develop future value functions. Then Eq. (15) is solved to find the expected value of water at every storage level in the system for several future months of the freshet using the hydrological information available at various months preceding the freshet. Expected future value functions for the months of April through June are calculated from forecasts updated in February, March, and April. The future value functions are used to calculate the marginal value of water (MVW) for each of the future months, f, and for each forecast month using Eq. (25). The evolution of the marginal values of water with evolving forecasts are evaluated and compared with the marginal value of water calculated by a dynamic programming model having perfect foresight.

$$\mathrm{MVW_f} = \frac{\delta\, f(S_f, H_t)}{\delta\ S_f * HK(\,S_f)} \tag{25}$$

Since input values do not change with time, value functions found from Eqs. (8) and (9) need only to be solved once for the *Hist no FC* and *Hist FC* models. However the *ESP* model is rerun using the most recent ESP forecasts in each stage.

4.3.4 Results - Marginal Water Value

The performance of the models are measured using the delta between the model's output of marginal value of water and the marginal value of water calculated from the perfect foresight model. We assume that the model having perfect foresight is able to hypothetically produce the best possible policies and the best case estimates of water value by employing the most efficient use of water. It follows that the marginal values of water are thus ideal, and we measure the delta between the SSDP model outputs and these ideal outputs in evaluating the quality of model performance.

The results shown in Figure 8 indicate that the use of forecasts improves estimates of the marginal value of water by the SSDP model. Forecasts are used by the model to gain knowledge of future flows which allows the model to use water more efficiently. As certainty associated with a forecast increases, the model is able to narrow its focus toward specific potential inflow realizations. When the forecast is accurate, the estimate of the marginal value of water improves. The magnitude of improvement depends heavily on the month being estimated and the timing of the forecast. Overall, the improvements in the marginal value calculation increases as forecasts are updated using both the ESP and Historical snowmelt runoff forecasts. The greatest improvements were seen when using the latest available forecasts (made in April) and are as high as 581 percent (at MCA using the Hist April forecast to estimate values in June). A five-fold change is partly due to the comparison with low marginal values from the model having perfect foresight in May and June causing even modest changes in improvement to comprise a large percentage.

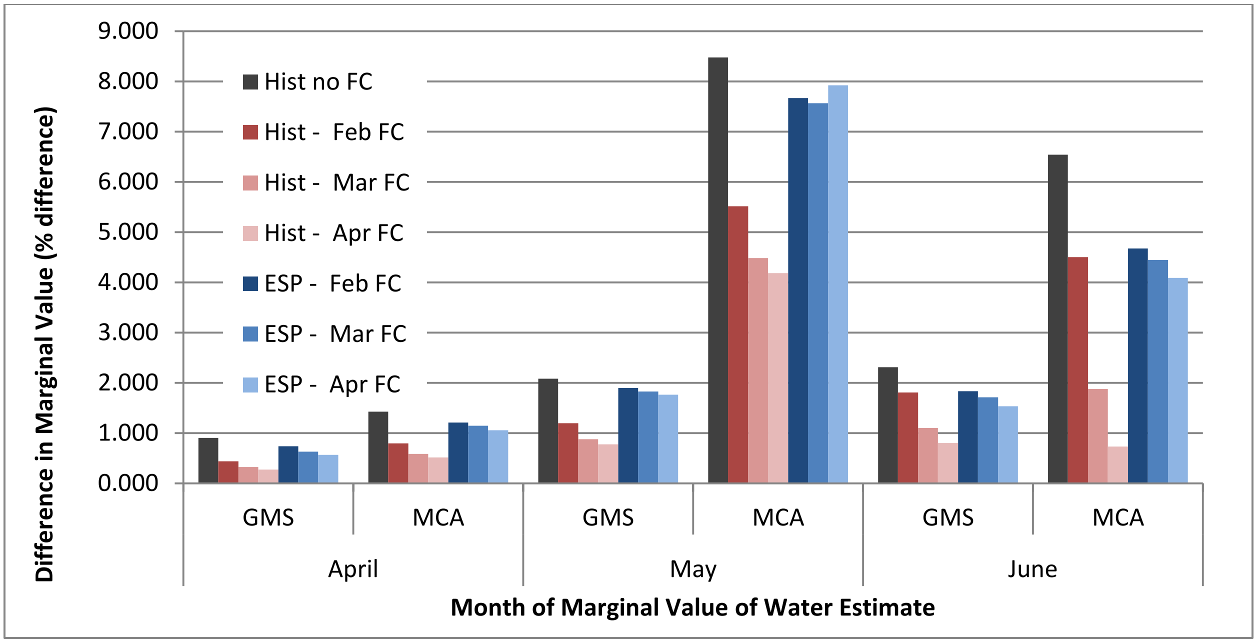

Figure 8. Marginal value of water (percent difference from model having perfect foresight) at GMS and MCA during the freshet as forecasts are updated. Models using historical seasonal runoff forecast are in shades of red, and models using ESP forecasts are in shades of blue with the timing and type of the forecast indicated on the legend.

When focusing on the models using historical forecasts (Hist – FC), the estimates of future marginal values further improve with time. The February forecast improved the MVW calculation by an average of 120 percent, with additional improvements of 94 and 33 percent in March and April. This result is an effect of increasing forecast certainty that is captured in the calculation of transition probabilities. The R-squared of the regression between forecasted and realized flows may be used as an indicator to describe the degree of forecast certitude and is shown in Table 5. As forecasts evolve from February to April, the R-squared value increases for every month of future flow.

Table 5. R-Squared of Regression between Evolving Historical Forecasts and Realized Inflow

Forecast Month	Month of Future Flow		
	April	May	June
February	0.731	0.708	0.682
March	0.778	0.778	0.742
April	0.828	0.827	0.815

However, ESP forecasts do not exhibit increased certainty with time. The coefficient of variance for ESP scenarios increases as the forecast is updated in March and April (for the future flow months of May and June) indicated in bold on Table 6. This lack of confidence explains the degradation of the marginal value of water estimate at Kinbasket Reservoir in May with the April forecast (seen in Figure 8).

Table 6. Monthly Coefficient of Variance with Evolving ESP Forecast

Forecast Month	Williston			Kinbasket		
	Month of Future Flow					
	April	May	June	April	May	June
February	0.242	0.184	0.141	0.211	0.199	0.120
March	0.227	**0.198**	0.137	0.199	**0.210**	0.118
April	0.178	**0.203**	0.123	0.163	**0.225**	**0.125**

4.3.5 Policy and Total Benefit Simulation Procedure

To determine how marginal values influence the benefits that may be achieved, the procedure above was modified to demonstrate reservoir operations from policies derived from each of the studied cases. Similar to the procedure above, steady state future value functions resulting from Eqs. (8) and (9) are developed, then Eq. (15) is solved to find the optimized real time releases using the hydrological data available and current storage level, rather than at all storage levels, at each step. The simulation moves forward from the beginning to the end of the planning horizon (February-July). We choose starting storage levels at 60 percent full at Williston and 62 percent full at Kinbasket which is a realistic realization of reservoir levels at this time of year. Once a policy decision is chosen, the reservoir storage level is updated based on the optimized policy, the actual flow that occurred, and expected spilling. The resulting storage value is used in the next time period and the procedure is repeated to the end of the planning horizon.

The value of different forecasts can be assessed by comparing the benefits gained in each case. Revenues from each time period are summed to determine value of sales; however we must also consider the value of water in storage at the end of the planning horizon. Since we allow reservoir storage to progress freely resulting from enacting optimal policy decisions, each simulation may arrive at different June ending storage values. For consistency among cases, this value is calculated using the marginal value of water derived from the deterministic model for all cases (Eq. (26)). The sum of revenues generated over all time periods and the ending storage value for each simulation are termed Total Benefits and compared.

For a system of K reservoirs:

$$\text{Value of Water in Storage} = \sum_{k=1}^{K} \int_{S_{k(min)}}^{S_{k(max)}} MVW * dS_k \qquad (26)$$

4.3.6 Results – Policy and Total Benefit

The simulations with forecast updating resulted in policy changes which varied with forecast timing and by reservoir, shown in Figure 9. Policies at Williston reservoir varied little during the month of February. March forecast updates resulted in changes in policy to the Historical model, but the ESP model policy remained consistent through the end of the planning horizon with no changes as a result of the updated April forecast. Differences in policy decisions were most evident at the Kinbasket reservoir. Release decisions are updated as forecasts evolve in February and March for both the *Hist FC* and *ESP* models, yet similar to the Williston reservoir, no policy changes occurred as a result of the April forecast update.

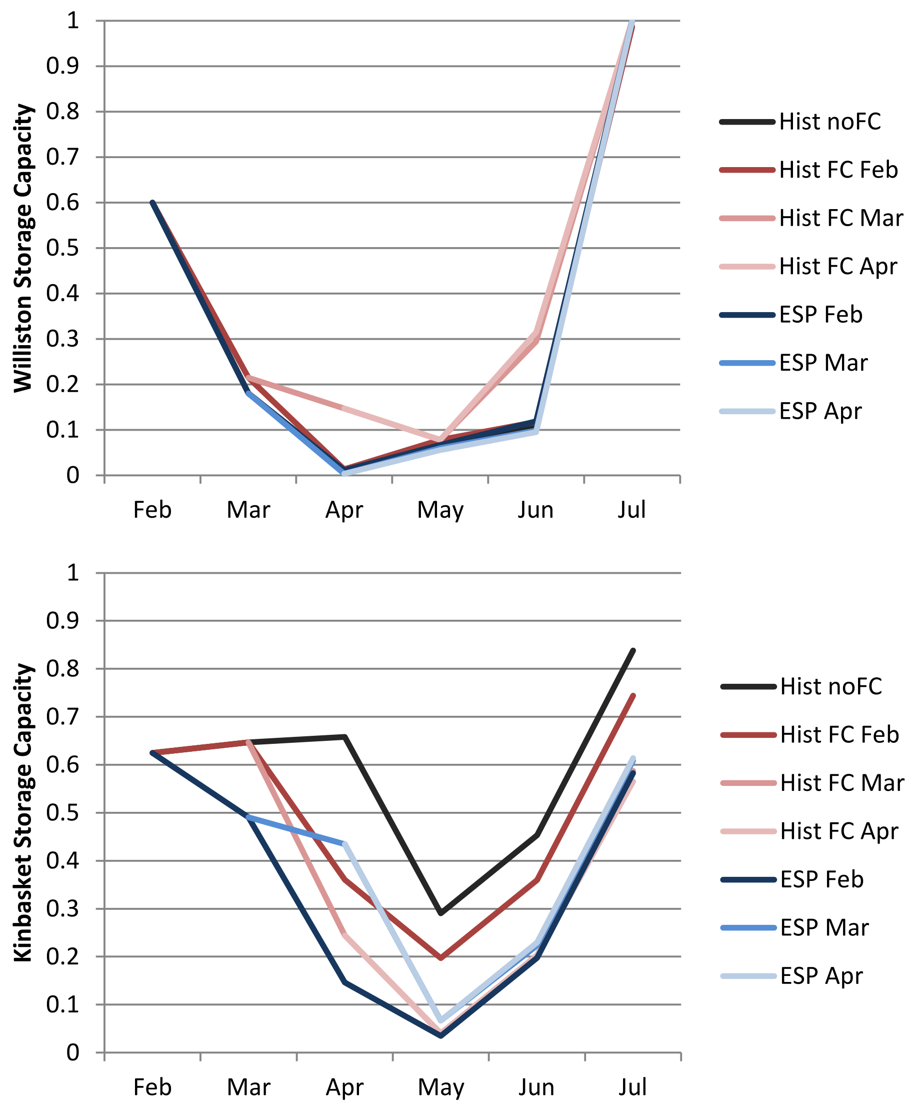

Figure 9. Projected storage at Williston and Kinbasket reservoirs. Optimized storage using no forecast is represented in black. Models using seasonal runoff forecast are in shades of red and models using ESP forecasts are in shades of blue with the timing of the forecast indicated on the legend.

Increases in total benefits (Figure 10) gained from operating the system were realized when forecasts were considered, and the timing of forecast was important. By the March forecast update, increased benefits of over 14 percent were realized for both the *Hist FC* and *ESP* models. However, the increases in total benefit occurring from the April forecast update are modest (less than 0.2 percent). Based on the lack of policy changes in April described above, this is not surprising. No gain or little gain in benefits made between February and April forecasts (*ESP* model) and March and April forecasts (*Hist FC* model) suggest that forecasts made in early spring may be sufficient to make future policy decisions.

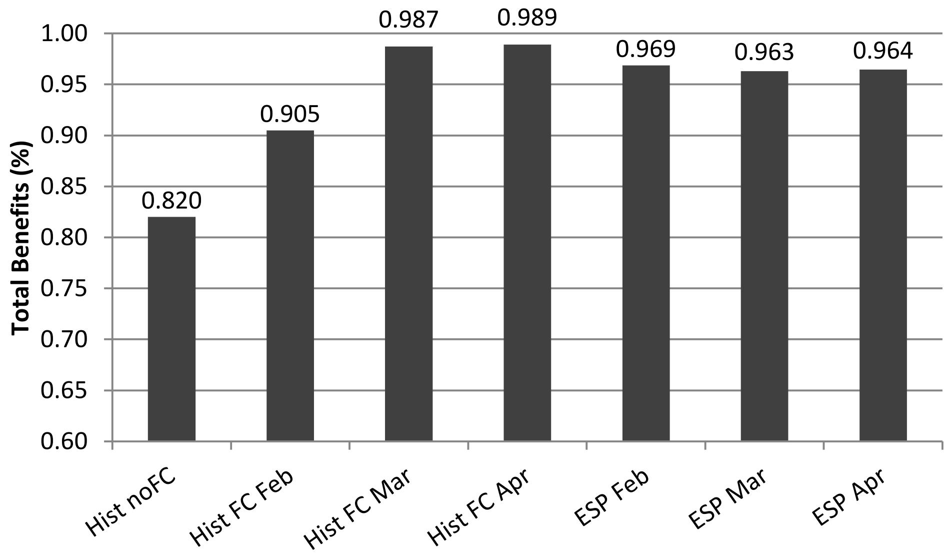

Figure 10. Total benefits (as a percentage of perfect) for evolving forecasts. Type and time of forecast are indicated on horizontal axis.

This result is interesting in that improvements are seen in the estimates of marginal values, yet they are not apparent in actual benefits. This is exemplified by clear and consistent improvement in marginal value of water estimates with the April forecast update in the *Hist FC* model (seen in Figure 8), and little improvement in the calculation of simulated total benefits in April. Little improvement in total benefits indicates that the differences in marginal values do not greatly affect results of the simulation. The marginal values shown are averages of the marginal value over the entire storage range. Therefore, overall improvements do not signify changes have occurred at every storage state. While the averaged marginal values improved during this simulation, the marginal values calculated at the storage states visited did not result in policy changes, and therefore no additional contribution to benefits was achieved.

4.4 Conclusions

In this paper we have demonstrated the use of two types of current forecasts with the SSDP algorithm to determine future policy decisions and the future value of water. We show that forecasts evolve in magnitude and level of certainty and both of these factors impact the model's estimates of the marginal value of water, which shapes future policies and benefits achieved. However, while marginal values are affected, it does not follow that operations will. The results indicate that the value of forecast updating is limited as additional benefits may diminish with time. The expected additional value associated with an updated forecast is valuable information for a decision maker. By comparing the value of increased future benefits to the value of making an early decision, a policy may be set when the most benefit is achieved.

Future work will include implementation of this model at BC Hydro as part of the Water Value Project. The model will be extended to incorporate Columbia River Treaty operational requirements that are modelled easily with the scenario based approach of the SSDP formulation.

5 Summary and Conclusions

This thesis developed a reservoir optimization model with the ability to consider the stochastic nature of inflows in maximizing the value of system resources. It investigated the use of various hydrological data inputs with a Sampling Stochastic Dynamic Programming (SSDP) model of the BC Hydro energy generation system. Several BC Hydro developed data sets consisting of inflow sequences and forecasts were used with the SSDP algorithm on complex single and multi-reservoir models.

SSDP is an attractive method for solving the BC Hydro system reservoir operations problem since it uses scenarios in its solution algorithm. Operational constraints mandated in the Columbia River Treaty (CRT) are structured to work well with scenario based modelling approaches.

The SSDP algorithm and various hydrological data were first applied to a single-reservoir model of the Williston Reservoir on the Peace River. The hydrologic data and study period was refined and applied to an expanded model representing operations of the BC Hydro system by a multi-reservoir system including both the Williston and Kinbasket reservoirs on the Peace and Columbia River basins. Both system representations produced consistent results.

A summary of research conducted is listed below:

- The SSDP algorithm was adapted for use with various types of inflows and forecasts including historical streamflows and forecasts, synthetically generated inflows and forecasts using a lag-1 autoregressive model and the principle component approach, and ESP forecasts. Differences in model outputs including policies and value functions from one year of simulated operations were compared.
- Marginal values were calculated at every month using various types of hydrologic data by simulating SSDP policies with three hydrological years of historical forecast and inflow data. The simulated marginal values of water were compared with the marginal value of water calculated by a deterministic model.

- The future values of BC Hydro system resources change as new forecasts become available and operations decisions are updated. The additional benefits gained with updated historical and ESP forecasts are compared on a two-reservoir model.

The overall conclusions reached by this study are summarized:

- Inflow characteristics including the relationships between forecasts and realized flow and standard deviation of monthly flows affect model behavior. Hydrologic state variables allow the model to understand expected flows and enabled improved decision making.
- The value of forecasts in approximating the marginal value of water in fall and winter months is relatively small, but increases into the freshet. During this time, historical and ESP hydrologic data sets result in better marginal value of water approximations than forecasts and inflows generated from an autoregressive model.
- The magnitude and level of certainty of forecasts impact estimates of the marginal value of water; however, it does not follow that these factors impact operations. The value of forecast updating is limited as additional benefits gained from refining forecasts may diminish with time.
- Finally, SSDP can be successfully applied to the BC Hydro system. Operational constraints are captured with the scenario-based optimization method, and this structure will allow the future planned work of incorporating CRT constraints to the SSDP model.

Bibliography

Abdalla, Alaa, Joel Evans, Ziad Shawwash, Ziming Guan, and Amr Ayad. 2013. *Water Value Project Definition Phase Report*. Vancouver, BC: BC Hydro.

Barber, C. Bradford, David P. Dobkin, and Hannu Huhdanpaa. 1996. “The Quickhull Algorithm for Convex Hulls.” *ACM Trans. Math. Softw.* 22 (4): 469–83. doi:10.1145/235815.235821.

“BC Bulk Transmission System.” 2007. BC Hydro. http://transmission.bchydro.com/NR/rdonlyres/127363BF-8F52-4CB0-91D4-2122EB048916/0/BCTCBCBulkMapAug07.pdf.

Bellman, Richard. 1957. *Dynamic Programming*. Princeton: University Press.

Cormen, Thomas H. 2001. *Introduction To Algorithms*. MIT Press.

Côté, Pascal, Didier Haguma, Robert Leconte, and Stephane Krau. 2011. “Stochastic Optimisation of Hydro-Quebec Hydropower Installations: A Statistical Comparison between SDP and SSDP Methods.” *Canadian Journal of Civil Engineering* 38 (12): 1427–34. doi:10.1139/l11-101.

Day, G. 1985. “Extended Streamflow Forecasting Using NWSRFS.” *Journal of Water Resources Planning and Management* 111 (2): 157–70. doi:10.1061/(ASCE)0733-9496(1985)111:2(157).

Dias, Bruno H., André L. M. Marcato, Reinaldo C. Souza, Murilo P. Soares, Ivo C. Silva Junior, Edimar J. de Oliveira, Rafael B. S. Brandi, and Tales P. Ramos. 2010. “Stochastic Dynamic Programming Applied to Hydrothermal Power Systems Operation Planning Based on the Convex Hull Algorithm.” *Mathematical Problems in Engineering* 2010 (May): e390940. doi:10.1155/2010/390940.

Di Domenica, N., C. Lucas, G. Mitra, and P. Valente. 2007. “Scenario Generation for Stochastic Programming and Simulation: A Modelling Perspective.” *IMA Journal of Management Mathematics* 20 (1): 1–38. doi:10.1093/imaman/dpm027.

Diniz, AL., and M. E P Maceira. 2008. “A Four-Dimensional Model of Hydro Generation for the Short-Term Hydrothermal Dispatch Problem Considering Head and Spillage Effects.” *IEEE Transactions on Power Systems* 23 (3): 1298–1308. doi:10.1109/TPWRS.2008.922253.

Eaton, Brett, and RD Moore. 2010. “Regional Hydrology.” *Compendium of Forest Hydrology and Geomorphology in British Columbia* 1: 85–110.

Faber, B.A. 2001. “Real-Time Reservoir Optimization Using Ensemble Streamflow Forecasts.” University of Colorado, Boulder, Colo.

Faber, B.A., and J.R. Stedinger. 2001. “Reservoir Optimization Using Sampling SDP with Ensemble Streamflow Prediction (ESP) Forecasts.” *Journal of Hydrology* 249 (1–4): 113–33. doi:10.1016/S0022-1694(01)00419-X.

Fourer, Robert. 2003. *AMPL: A Modeling Language for Mathematical Programming*. 2nd ed. Pacific Grove, CA: Thomson/Brooks/Cole.

Guan, Ziming, Ziad Shawwash, Alaa Abdalla, Amr Ayad, and Joel Evans. 2013. “Assessing How Uncertainty Affects Reservoir Operations.” *HydroWorld Weekly* 32 (3): 76–81.

Johnson, Sharon A., Jery R. Stedinger, Christine A. Shoemaker, Ying Li, and José Alberto Tejada-Guibert. 1993. “Numerical Solution of Continuous-State Dynamic Programs Using Linear and Spline Interpolation.” *Operations Research* 41 (3): 484–500.

Kaut, Michal, and Stein W. Wallace. 2003. “Evaluation of Scenario-Generation Methods for Stochastic Programming.” *World Wide Web, Stochastic Programming E-Print Series*, July. http://edoc.hu-berlin.de/docviews/abstract.php?id=26731.

Kelman, Jerson, Jery R. Stedinger, Lisa A. Cooper, Eric Hsu, and Sun-Quan Yuan. 1990. “Sampling Stochastic Dynamic Programming Applied to Reservoir Operation.” *Water Resources Research* 26 (3): 447–54. doi:10.1029/WR026i003p00447.

Kim, Y., H. Eum, E. Lee, and I. Ko. 2007. “Optimizing Operational Policies of a Korean Multireservoir System Using Sampling Stochastic Dynamic Programming with Ensemble Streamflow Prediction.” *Journal of Water Resources Planning and Management* 133 (1): 4–14. doi:10.1061/(ASCE)0733-9496(2007)133:1(4).

Kim, Y. O., H. I. Eum, and I. H. Ko. 2006. “Value of Updating Ensemble Streamflow Prediction in Reservoir Operations.” In , 422–31. American Society of Civil Engineers. doi:10.1061/40875(212)42.

Krzysztofowicz, Roman. 2001. “The Case for Probabilistic Forecasting in Hydrology.” *Journal of Hydrology* 249 (1–4): 2–9. doi:10.1016/S0022-1694(01)00420-6.

Little, John D. C. 1955. “The Use of Storage Water in a Hydroelectric System.” *Journal of the Operations Research Society of America* 3 (2): 187–97.

Loucks, Daniel P., and Philp J. Dorfman. 1975. “An Evaluation of Some Linear Decision Rules in Chance-Constrained Models for Reservoir Planning and Operation.” *Water Resources Research* 11 (6): 777–82. doi:10.1029/WR011i006p00777.

Pereira, M. V. F., and L. M. V. G. Pinto. 1991. “Multi-Stage Stochastic Optimization Applied to Energy Planning.” *Mathematical Programming* 52 (1-3): 359–75. doi:10.1007/BF01582895.

Schaffer, Jennifer. 2014. “Performance of a Sampling Stochastic Dynamic Programming Algorithm Using Various Inflow Generation Methods.” presented at the CWRA

Workshop: From operational hydrological forecast to reservoir management optimization, Quebec, QC, September 17.

Schaffer, Jennifer, and Ziad Shawwash. 2014. “Performance of an SSDP Algorithm Using Various Inflow Generation Methods.” In *Proceedings of HydroVision Conference*. Nashville, TN.

Stedinger, Jery R., and Young-Oh Kim. 2010. “Probabilities for Ensemble Forecasts Reflecting Climate Information.” *Journal of Hydrology* 391 (1–2): 9–23. doi:10.1016/j.jhydrol.2010.06.038.

Stedinger, Jery R., Bola F. Sule, and Daniel P. Loucks. 1984. “Stochastic Dynamic Programming Models for Reservoir Operation Optimization.” *Water Resources Research* 20 (11): 1499–1505. doi:10.1029/WR020i011p01499.

Tejada-Guibert, J. Alberto, Sharon A. Johnson, and Jery R. Stedinger. 1993. “Comparison of Two Approaches for Implementing Multireservoir Operating Policies Derived Using Stochastic Dynamic Programming.” *Water Resources Research* 29 (12): 3969–80. doi:10.1029/93WR02277.

———. 1995. “The Value of Hydrologic Information in Stochastic Dynamic Programming Models of a Multireservoir System.” *Water Resources Research* 31 (10): 2571–79. doi:10.1029/95WR02172.

Tejada-Guibert, Jose Alberto. 1990. “Spline Stochastic Dynamic Programming for Multiple Reservoir System Operation Optimization.” Ph.D., United States -- New York: Cornell University. http://search.proquest.com.ezproxy.library.ubc.ca/docview/303851525/abstract?accountid=14656.

Tilmant, A., D. Pinte, and Q. Goor. 2008. “Assessing Marginal Water Values in Multipurpose Multireservoir Systems via Stochastic Programming.” *Water Resources Research* 44 (12): W12431. doi:10.1029/2008WR007024.

Vicuna, Sebastian, John A. Dracup, Jay R. Lund, Larry L. Dale, and Edwin P. Maurer. 2010. “Basin-Scale Water System Operations with Uncertain Future Climate Conditions: Methodology and Case Studies.” *Water Resources Research* 46 (4): W04505. doi:10.1029/2009WR007838.

Weijs, Steven V., and Nick van de Giesen. 2013. “An Information-Theoretical Perspective on Weighted Ensemble Forecasts.” *Journal of Hydrology* 498 (August): 177–90. doi:10.1016/j.jhydrol.2013.06.033.

World Meteorological Organization. 1986. *Intercomparison of Models of Snowmelt Runoff*. Geneva, Switzerland: Secretariat of the World Meteorological Organization.

Yakowitz, Sidney. 1982. “Dynamic Programming Applications in Water Resources.” *Water Resources Research* 18 (4): 673–96. doi:10.1029/WR018i004p00673.