# A Chronological Probabilistic Production Cost Model to Evaluate the Reliability Contribution of Limited Energy Plants



by

Tommy Leung

B.S. Engineering, Harvey Mudd College (2005)

Submitted to the Engineering Systems Division
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE IN TECHNOLOGY AND POLICY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2012



Author ....................................................................
Engineering Systems Division
May 11, 2012

Certified by ....................................................................
Ignacio J. Pérez Arriaga
Visiting Professor, Engineering Systems Division
Thesis Supervisor

Certified by ....................................................................
Carlos Batlle
Visiting Scholar, MIT Energy Initiative
Thesis Supervisor

Accepted by ....................................................................
Joel P. Clark
Professor of Materials Systems and Engineering Systems
Acting Director, Technology and Policy Program

# A Chronological Probabilistic Production Cost Model to Evaluate the Reliability Contribution of Limited Energy Plants

by

Tommy Leung

Submitted to the Engineering Systems Division
on May 11, 2012, in partial fulfillment of the
requirements for the degree of
MASTER OF SCIENCE IN TECHNOLOGY AND POLICY

## Abstract

The growth of renewables in power systems has reinvigorated research and regulatory interest in reliability analysis algorithms such as the Baleriaux/Booth convolution-based probabilistic production cost (PPC) model. However, while these traditional PPC algorithms can reasonably represent thermal plant availabilities, they do not accurately represent limited energy plants because of their generic treatment of time. In particular, in systems with limited energy plants, convolution-based PPC models tend to underestimate the loss-of-load probability and expected nonserved energy. This thesis illustrates the chronological challenges of the traditional convolution-based PPC, proposes a modification that improves the representation of chronological elements, explores the reliability contribution of LEPs using the new algorithm, and demonstrates two regulatory applications by calculating a capacity payment for an LEP and the expected-load-carrying-capability metric for any generator. To the best knowledge of the author, the introduction of multiple hydro plants with different capacity constraints and the calculations for marginal probabilities, prices, and revenues to a chronological PPC model are novel.

Thesis Supervisor: Ignacio J. Pérez Arriaga
Title: Visiting Professor, Engineering Systems Division

Thesis Supervisor: Carlos Batlle
Title: Visiting Scholar, MIT Energy Initiative

# Acknowledgments

"The only reason for time is so that everything doesn't happen at once."

-Albert Einstein

I humbly thank Pablo, Carlos, and Ignacio for their great mentorship, excellent advice, and long hours (especially with the time difference!); Ernie and Melanie for the fantastic opportunities over the last two years; and my parents for their amazingness!

THIS PAGE INTENTIONALLY LEFT BLANK

# Contents

# List of Figures

# List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

# Chapter 1

# Introduction

## 1.1 A brief overview of electricity markets

With exception to popular historical anecdotes such as Thomas Edison's light bulb, history books today contain few electricity-related events or catastrophes. Modern societies view access to low-cost and dependable electricity as a right, and until the early 1990s, vertically integrated utilities delivered reliable electricity to their consumers in a remarkably steady but otherwise unremarkable fashion.

### Vertically integrated utilities

Before the early 1990s, vertically integrated utilities operated electric power networks as regulated monopolies. These monolithic entities owned all of the generation and transmission assets within their networks. They made short-term decisions about how to operate their power plants on a day-to-day basis, as well as long-term decisions about future network and generation investments. Because they operated under cost-of-service (and therefore recouped all of their investment costs), vertically integrated utilities also tended to overinvest in their network and generation assets to ensure against the physical, political, and social impacts of system failures.

Despite the greater protection that overinvestment afforded against system failures, it also decreased economic efficiency and raised the average cost of electricity for consumers. To address these inefficiencies, in the early 1990s, power systems around the world began transitioning to electricity markets. These transitions, often referred to as market "liberalization," "deregulation," or "restructuring," separated monolithic, vertically integrated utilities into four main businesses: generation, transmission, distribution, and retail. Because of the economies of scale associated with transmission and distribution networks (for example, a high-capacity line loses less power than two lines that sum to the same capacity), network operations remained regulated monopolies. The most notable change for systems that liberalized occurred at the generator level with the creation of new power and reserve markets: individual generators and generation companies could now compete against each other for the right to sell electricity, and new entrants could enter the market with any technology of their choosing.[1]

### Liberalization challenges

However, liberalization also eliminated the central planning role of the vertically integrated utility, leaving only individual agents to make both short- and long-term decisions about plant operations and investments in their own best economic interests. Although economic theory dictates that under perfect competition and information markets will drive individual agents to make economically optimal choices, in reality the individual actor (for example, an investor) does not have perfect information. Unable to forecast demand and future prices with certainty, investors in power systems with markets will likely underinvest in new power plants because they face greater

[1]Electricity markets, as with other markets, operate on the economic principle that perfect competition will produce efficient outcomes. With enough ("perfect") information about electricity prices, load trends, generation technologies, and other pertinent aspects of the power system, private entities should be able to make prudent investment and operation decisions that lead to an economically efficient generation mix and electricity supply.

Concretely, in the United States, integrated system operators (ISOs) encourage competition by running auctions for the many electricity products that they need. Largely, these products are either energy- or reserve-related. Marginal prices for all products emerge from the auctions. To participate in these auctions, generators must submit bids consisting of quantity and price pairs that they are willing to sell electricity at. Additionally, because ISOs in the United States utilize complex bids, generators must also submit information about their plants' physical constraints (for example, their ramping, start-up, and shut-down capabilities). An ISO will take all bids and constraints into account to determine the economic merit order, which ranks plants from least to highest cost. Then, the ISO awards bids starting with the least expensive plants until all demand is met. In each hour, the last bid that the ISO accepts sets the marginal system price. This is the marginal system price that signals to investors the potential value of investing in new capacity.

risks associated with recovering their costs and fewer direct consequences from failure events—directly in contradiction to the expected behavior of vertically integrated utilities.

Because private entities are free to make their own investment decisions, and because they are likely to underinvest in capacity given their risk aversion, the electricity industry by and large concedes (implicitly and explicitly) that regulators still hold the responsibility for ensuring the availability of adequate capacity to prevent system failures. More generally, when designing market rules, regulators have the responsibility of internalizing important political and social concerns (e.g., $CO_2$ prices and climate change) that consumers cannot explicitly express preferences for because electricity markets remain immature. As explained in [15, Rodilla 2010], most demand-side consumers have not learned how to, or cannot, respond to electricity spot prices. Consumers typically also do not how to, or cannot, express their preferences for products such as supply security because historically, the vertically integrated utilities managed these types of concerns. Although economic theory dictates that consumers, if left to market forces, would adapt and eventually learn how to express preferences for these products after enough black- and brown-outs, electricity failures impose great burdens on a society. The public at large, politicians, and governments find these failures untenable (for example, consider the Californian government's response to its 2001 blackouts) and are unlikely to allow enough time for markets to mature on their own. Consequently, because consumers are unable to directly express their preferences, regulators intervene in electricity markets to address a variety of market failures, including the problem of an inadequate security of supply.

### Regulatory tools for reliability

To design market incentives and rules that guide electricity markets toward publicly desirable and economically efficient outcomes, regulators use a wide range of support tools (such as mathematical optimizations and simulations) to analyze the physical and market operations of power systems. Among the types of questions that regulators ask, questions about system reliability frequently surface. For example, how reliable is a power system? What is the contribution of a particular plant to system reliability? What percentage of a plant's capacity should the regulator consider "firm?"[2] With the prevalence of public subsidies for various generation technologies, questions about the intermittency challenges that renewables present, and concerns about climate change, these types of reliability questions and the tools that can be used to analyze them have once again piqued the interest of power systems researchers and regulators.

This thesis focuses on limited energy plants (LEPs)[3] as a source of system reliability for power systems and the convolution-based probabilistic production cost (PPC) algorithms that regulators have used to evaluate the contributions of various generators (and in particular, LEPs) to reliability. Countries such as Panama, Spain, and Ireland have used PPC models in the past to determine the contribution of thermal and hydro generators to the reliability of their power systems. For example, in Panama, the system regulator used to calculate the reliability contribution of individual generators by first, benchmarking the its entire system's loss-of-load probability (LOLP) using the PPC model; second, removing individual plants and rerunning the PPC model; and third, crediting the removed plant for reducing the probability of failure based on the net change in LOLP between the first and second steps. Although the traditional convolution-based PPC models (hereafter referred to as "traditional PPC models") that countries such as Panama have used for their reliability analyses can model thermal availability well, they have difficulty accurately representing limited energy plants. This thesis proposes a modification to the traditional PPC algorithm that better addresses the representation of chronological elements. Given the recent renewed interest in regulatory instruments to encourage investments and

[2]Although many definitions for "firm" capacity exist, for the purposes of this thesis, a plant's "firm" capacity refers to the fraction of its nameplate capacity that the regulator considers reliable. For example, a combined-cycle-gas-turbine plant with a nameplate capacity of 400 MW and a forced outage rate of 5% might have a firm capacity of 380 MW. Many definitions exist for firm capacity (and as many metrics for evaluating a generator's firm capacity; Chapter 4 explores one such metric). Broadly speaking, a generator that has a firm capacity close to its nameplate capacity should be more reliable than a generator with the same nameplate capacity, but a lower firm capacity.

[3]LEPs are generators that store a limited amount of energy/fuel; for instance, a reservoir-based hydro generator is a well known example of an LEP.

behaviors that increase system reliability, the results from this thesis can make an immediate impact on discussions about the reliability contribution of LEPs and the firm capacity of various generation technologies.

## 1.2 Reliability in electric power systems

As noted by [14, Rodilla 2010], the physical task of delivering electricity to end consumers consists of a complex set of coordinated actions between multiple actors. At every time instant, physical laws require balance between generation and demand. Supplying electricity without interruption resembles an intricate dance between large, synchronous machines that are connected across thousands of miles and constantly converting mechanical energy into electrical energy (and vice versa). The success of this complex machine requires decisions that span multiple timescales, from building generation plants and transmission networks (processes that may take multiple years) to physically operating individual generators on a minute-by-minute and second-by-second basis. [14, Rodilla 2010] describes in detail the different temporal scales that the reliability problem can be broken down into.

For this thesis, the discussion about a power system's reliability will focus on the existence of enough installed capacity and its availability to supply demand. Within this scope, a power system's reliability can be characterized using many different metrics. Because regulatory analyses frequently discuss reliability in terms of a system's LOLP and expected nonserved energy (ENSE), the remaining discussion about reliability will focus on these two metrics. Chapter 2 reviews the formal definition of LOLP and ENSE. Broadly, LOLP is the expected fraction of hours over the time period of analysis (for example, one year) in which demand will exceed available generation, and ENSE is the total expected amount of unmet demand over the time period of analysis. Using these definitions, systems with lower LOLP and ENSE values are more reliable than systems with higher LOLPs and ENSEs.

## 1.3 Limited energy plants as a reliability resource

Reservoir-hydro resources and storage technologies—more generally, LEPs—can contribute to solving the power system reliability problem. LEPs store energy in one time period for dispatch in a future time period; for this reason, LEPs face a different cost of dispatch than thermal plants. For a thermal plant, the value of generation in one hour is the difference between its operating costs and the marginal system price. For an LEP, the value of dispatching energy in one hour is the opportunity cost of not saving that energy for use in a future hour. Deciding how to best dispatch LEP resources requires a careful treatment of uncertainties that affect thermal plants to a significantly lesser degree.

### Thesis objectives

Using reservoir-hydro (hereafter referred to as hydro) plants as proxies for LEPs, this thesis examines the PPC algorithms that researchers and regulators have used to evaluate system reliability. In the past, regulators have used PPCs to analyze system reliability and the reliability contributions of individual plants because PPC algorithms require relatively little computational effort and can reasonably approximate thermal availability. However, these evaluations do not hold as well for LEPs. Because of their treatment of time, traditional PPCs assume one dispatch behavior for LEPs for a unit time and then scale this behavior up for the entire simulation period. Consequently, these models implicitly consider dispatch scenarios for LEPs that may violate their energy limits. Additionally, because of their treatment of time, traditional PPCs discard chronological information that may be particularly useful in the representation of certain types of LEP plants; e.g., hydro resources. To overcome these two limitations while also preserving the advantages of the convolution-based methodology, this thesis develops a chronological PPC model that can more accurately represent LEPs. To the best knowledge of the author, the representation of multiple hydro plants with capacity constraints and the calculation of hourly marginal probabilities, prices, and revenues in the chronological algorithm are novel contributions.

In addition to presenting the proposed chronological algorithm, Chapter 2 examines how key reliability metrics are computed in both the traditional PPC and the proposed chronological model. The last section in Chapter 2 illustrates differences between the two models by applying them to a real-size case study power system. The results highlight the optimistic nature of traditional PPC models in their treatment of hydro dispatch and reliability estimates. Chapter 3 explores in depth the calculation of total system costs, marginal prices, revenues, and reliability metrics under the new chronological PPC algorithm. Chapter 4 discusses regulatory issues related to the reliability problem by applying the chronological algorithm to (1) design incentives for LEPs to improve system reliability and (2) calculate the reliability contribution of an individual generator. Chapter 5 concludes with suggestions for future research.

## 1.4 A broader context: renewables in power systems

Academically, the representation of chronology in PPC algorithms poses an interesting challenge because the trade-offs between computational effort and model accuracy impose real constraints that might have clever-yet-undiscovered workarounds. More practically, for regulatory and policy purposes, the recent growth in renewable generation technologies also motivates the study of chronological PPC algorithms because LEPs can help power systems integrate larger fractions of renewables into their generation mixes. The remainder of this chapter offers background about the challenges of integrating renewables to explain the greater motivation behind studying and developing new algorithms (such as the PPC models presented in this thesis) for power systems.

Although supporters often cite clean emissions and free fuels as key reasons to promote the adoption of renewable technologies for electricity generation, the variability of renewable generation creates new load-balancing challenges for power systems. A generator's variability depends on the intermittency and predictability of its generation.[4] The lack of predictability for highly intermittent sources of electricity, such as wind and solar generators, requires power systems to make frequent supply adjustments over shorter timescales to balance load and supply. These frequent supply adjustments can create reliability problems and impose additional costs onto other generators in the system. For example, in power systems that allow renewables to dispatch first in violation of the economic merit order, the nonrenewable plants hold the responsibility for balancing generation and demand. Yet, the intermittency of renewables often increases the difference between a system's minimum and maximum net load[5]. Consequently, to accommodate excess wind generation on a low demand night in the spring, a coal plant might have to ramp or shut down. Ramping and cycling operations are generally uneconomic because plants incur more physical wear than usual (but are only paid for their generation), operate under decreased efficiency (consume more fuel per unit of electricity produced), and emit more greenhouse gasses. The inverse problem also exists: on hot summer days with peak demand for electricity, if the wind stops, the power system may not have enough thermal capacity to cover remaining demand. If the system does not have enough thermal capacity, who should be held responsible for the resources that are needed to maintain system reliability? These inversely correlated generation/demand examples show how the variability of renewables and supportive policies such as priority dispatch can create short-term externalities for nonrenewable generators and consumers, despite the benefits of renewable electricity.

In the long-term, public subsidies for renewables and priority dispatch rules may also discourage investment in other technologies that are needed to maintain reliable

---

[4]"Intermittency" refers to uncontrolled changes in the output of a generating resource, and "predictability" refers to the ability to estimate a resource's intermittency.[1] Wind turbines, concentrated solar power systems, and photovoltaic solar systems all exhibit high intermittency because the amount of electricity that they generate changes with the availability of wind and sunlight. However, the intermittency of these three technologies are not equally predictable. Generally, because weather forecasters can forecast cloud coverage with greater accuracy over longer periods of time than they can forecast wind, wind generation tends to be more difficult to predict than solar generation.

[5]The "net load" of a power system refers to its demand after subtracting out generation from nondispatchable sources such as wind and solar generators. The "net load" is the amount of demand that dispatchable generators—typically thermal and hydro plants—must provide generation for.

power systems. As capacity from renewable technologies continues to grow in power systems, despite the greater generation variability, private investors will have fewer incentives to invest in new conventional projects (such as flexible combined-cycle-gas-turbine units) because larger renewable generation mixes reduce the amount of electricity that thermal plants can sell. Market rules such as reserve capacity payments can incentivize investment in flexible generation, but ambiguity surrounds the determination of who should pay for these incentives; additionally, these new market rules may not lead to the least-cost power system for the consumer. These physical, economic, and regulatory questions represent the types of integration concerns that exist for renewables, and the tools required to analyze these problems—for example, the PPC models presented in this thesis—remain an active area of research in power systems.

THIS PAGE INTENTIONALLY LEFT BLANK

# Chapter 2

# A new proposal for evaluating power system reliability

As described in Chapter 1, regulators have used PPC algorithms to evaluate the reliability of power systems based on their LOLP and ENSE. This chapter briefly reviews the history of PPC algorithms, describes the challenge of properly modeling LEPs with traditional PPC models, and surveys the current body of literature on chronological PPC models that offer improved representations of LEPs. The end of this chapter contains a proposal for a new chronological PPC algorithm that divides every hour into its own reliability problem, as well as a comparison of the LOLP and ENSE results between the proposed algorithm and a traditional PPC model.

## 2.1 Past work: production cost models

Historically, the simplest production cost models deterministically approximated thermal plant failures by representing their output levels as fractions of their maximum capacities. The forced outage rate (FOR) of a plant determined the specific fraction of its total capacity that counted as firm capacity. Under this representation, thermal plants could never fail at their FOR-reduced capacities. Deterministic models treated hydro plants as thermal plants with FORs of zero, and they set hydro output levels to perfectly consume all available water in a given simulation period. This early genre of production models did not take into consideration the fact that when plants fail, their outputs drop to zero, leading to an underestimation of the need for generation from more expensive units. [7, Finger]

The probabilistic production cost (PPC) models that followed treated thermal plant outages more realistically. Primarily developed by [2, Baleriaux] and reintroduced in English by [3, Booth], PPCs represented thermal plants using a two-state model. In the first state, the plant is available to generate electricity at its full capacity with probability $(1 - p)$. In the second state, the plant is not available to generate electricity (due to a forced outage) with probability $p$. By considering these two potential probabilistic states, Baleriaux/Booth created a new class of production cost models that were able to more accurately capture the effect of thermal plant outages.

In the following decades, many authors proposed iterations and refinements to the Baleriaux/Booth PPC model. Of these refinements, notably [6, Conejo] developed an approach to incorporate hydrothermal coordination. More generally, the PPC techniques proposed by Conejo for optimal charging and discharging, as well as to determine the optimal merit order position to minimize system cost, applied to all LEPs (e.g., batteries, flywheels, compressed air storage)—not only hydro plants. As regulatory tools, derivatives of the Baleriaux/Booth PPC model have remained useful because they require relatively little computational effort, capture the discrete nature of plant failures, and directly convey a system's reliability in terms of its ENSE and LOLP metrics.

### 2.1.1 PPC assumptions

Despite the many advantages of traditional PPC models, their abstraction of time results in less accurate representations of power system. Most traditional PPC models represent demand over the time period of interest, $T$, with a single cumulative probability distribution function (CDF). Every small $t$ in time period $T$ looks identical. As such, traditional PPCs calculate results for one generic unit of time and then extrapolate those results out to longer time scales. The extrapolation holds for thermal plants because the dominant characteristics of a thermal plant are its capacity and failure rate. These characteristics do not change much with time. However, the extrapolation does not hold for LEPs because an LEP's energy constraint (how much energy it has stored) changes through time and affects its dispatch actions. Traditional PPCs unrealistically assume that an LEP's dispatch will remain constant through the simulation period $T$ because they cannot capture how variables change with $t$. (A graphical explanation of this follows in section 2.2.2.)

Additionally, traditional PPC models also discard chronological information. As an example of the importance of chronology, consider hydro plants. In traditional PPC models, hydro energy targets are only enforced on the average because every $t$ is identical in the simulation period $T$. If a hydro plant has 100 MWh of energy, the PPC will perfectly place every drop of water to use up all 100 MWh over time period $T$. However, in reality, due to inherent demand and plant availability uncertainty, in some hours hydro plants will generate less than what they should (and in others, more than they should). In the hours when hydro generation is short of the optimum, thermal plants will make up the difference. Conversely, in the hours when hydro generation exceeds the optimum, hydro operators will sell electricity that they could saved for more expensive hours. Both of these scenarios result in higher actual total system costs than those predicted by traditional PPC models because of the loss of chronological information.

As noted by [11, Maceira & Pereira 1996], chronology has not always posed a problem for PPCs. Specifically, in a thermal-dominated power system, PPC models can reasonably approximate the behavior of thermal plants. However, as the number of LEPs in a power system increases, the system's total energy constraints from LEPs take on greater importance and the assumption of a predominately thermal system breaks down. In nonthermal-dominated systems, chronological elements can cause material deviations between a PPC models' predictions about system reliability and the actual metrics for that system.

### 2.1.2 Existing literature on chronological PPCs

To address the limited representation of energy constraints and the loss of chronological information in traditional PPC models, [11, Maceira & Pereira 1996] proposed an algorithm that decomposes the Baleriaux/Booth PPC into a series of chronological reliability problems. The algorithm consists of a power system with multiple thermal plants and one energy-limited hydro plant. The system always dispatches its hydro plant last. Production simulations are run for time period $T$ chronologically, hour by hour. Random variables represent inflow, demand, thermal generation, and turbine capacity. As the simulation runs, the model can probabilistically track the initial reservoir level in each hour (storage) and the hourly outflow (demand minus thermal generation). To calculate hourly reservoir levels, the simulation convolves the hourly inflow, storage, and outflow variables. Unused water carries over to the next hour, and hourly water deficits represent ENSE. The chronological algorithm produces marginal costs and reliability metrics for ENSE and LOLP in each hour. Using these outputs, [11, Maceira & Pereira 1996] compared a traditional PPC model with their proposed chronological model and determined that the traditional PPC underestimated costs for the particular system that they analyzed (as noted in their paper, the comparison is always system-dependent).

The chronological model in [11, Maceira & Pereira 1996] refined the probabilistic treatment of LEPs in PPC models by incorporating important chronological parameters such as inflows, outflows, and reservoir storage levels. The work in this thesis builds off these results, as well as the works of other authors that have developed chronological PPC models to analyze hydro dispatch and other time-dependent elements in power systems:

- [4, Borges 2008] uses a chronological Markov chain model to analyze stochastic river inflows and obtain steady state probabilities for the total available energy of a small hydro power plant. Their approach covers a timescale of months and tries to improve estimates of total available hydro energy from small hydro power plants to aid long-term capacity planning. The model separately describes the turbine portion of a hydro plant with a two-state Markov model and produces reliability metrics for the expected amount of energy available taking into consideration generator failure. Instead of explicitly treating demand, Borges presents energy-availability reliability metrics for small hydro power plants. These metrics include a small hydro power plant's installed energy, expected available energy, expected generated energy, capacity factor (considering only the energy source), and generation availability factor.

- [8, Gonzalez 2005] describes a water dispatch policy that optimizes economic benefit, taking into consideration both cost minimization and the reliability objective described by [13, Nabona 1995]. The algorithm divides a year into equal subperiods of months and combines all hydro plants into a single, monolithic hydro plant of equivalent capacity and energy. Demand is represented by a load duration curve (LDC), and the authors assume that price and demand are directly correlated. The algorithm splits hydro usage explicitly for peak shaving and to cover thermal failures. Hydro energy used to cover thermal failures is always sold in the reserve market. In scenarios that have more water than ENSE for the entire year, the algorithm decides what the best allocation of water is for each multiweek subperiod. In each subperiod, an amount of hydro energy equal to the ENSE is dispatched for reliability. Any remaining hydro energy is dispatched for peak-shaving. The simulation sets the reservoir level at the beginning and does not consider inflows.

- [9, Gonzalez 2002] describes a water dispatch policy that minimizes cost, taking into consideration the stochastic nature of inflows and outflows. The simulation breaks the total time period of a year into a series of smaller periods, such as months (the algorithm generalizes well to even shorter time periods). In each interval, the simulation considers water storage, discharge, pumping, and spillage. Excess water carries from one interval to the next. Using an LDC to represent demand, the algorithm randomly samples dispatch scenarios bounded by hydro balance and feasibility rules and evaluates each scenario based on cost. Gonzalez saves "winning" scenarios and evaluates them by simulation to determine a mean total cost and its probability distribution.

- [10, Gonzalez 2000] describes a hydro dispatch procedure that convolves hydro generation and its unavailability distribution with the LDC to capture the stochastic elements of hydro generation, based on (Nabona 1995). The simulation time period is one year, subdivided into months. Optimal dispatch values are found for each subperiod.

- [13, Nabona 1995] proposes a method for optimizing long-term hydrothermal usage by splitting the amount of hydro available to serve the deterministic (economic) goal of peak shaving the LDC and the stochastic (reliability) goal of covering thermal failures.

## 2.2 A new chronological PPC algorithm

### 2.2.1 Overview

To capture the chronological information that PPCs discard when they create a single aggregate LDC for all of time period $T$, the proposed chronological algorithm deconstructs each hour in $T$ into its own reliability problem. To the best knowledge of the author, the introduction of multiple hydro plants with different capacity constraints and the calculations for marginal probabilities, prices, and revenues to a chronological PPC model are novel.

#### Thermal dispatch

In each hour, the algorithm represents demand as a discrete CDF and dispatches generation plants by convolution. For every thermal plant, the algorithm performs a single convolution between two demand CDFs. The first CDF represents demand if the thermal plant fails (i.e., the demand CDF remains unchanged). The second CDF represents demand after removing a portion of demand equal to the capacity of the thermal plant. The FOR of the thermal plant determines the weight for each CDF (the plant fails with probability $p$; the plant works with probability $1 - p$). Thermal plant dispatch in this chronological PPC model closely resembles thermal plant dispatch in the traditional PPC.

### Hydro dispatch

Hydro plant dispatch also largely resembles thermal dispatch, with the key distinction that this algorithm also tracks the energy CDF for each hydro reservoir. In each hour that the algorithm dispatches a hydro plant, it performs two convolutions. The first convolution modifies the demand CDF, much like the thermal convolution described above. In place of thermal FORs, these demand convolutions substitute the y-intercept of the hydro reservoir CDF (representing the probability that the reservoir is empty). The y-intercept of the hydro reservoir CDF is analogous to the FOR of a thermal plant. The FOR of a hydro plant changes with time, based on the amount of energy stored in the reservoir—the more energy, the more reliable the hydro plant.

The second convolution modifies the hydro reservoir CDF to reflect the amount of water released in each hour. In the hydro reservoir convolution, the scenario that hydro is needed is convolved with the scenario that hydro is not needed. If hydro is not needed, the reservoir CDF remains unchanged. If hydro is needed, then the algorithm removes an amount of energy equal to the capacity limit of the hydro plant. The intermediate LOLP values for demand determine the weights for each scenario (i.e., the probability that ENSE is strictly positive and that the system requires water is the current LOLP; the probability that ENSE is zero and that the system does not require water is the complement of the current LOLP). Modeling hydro reservoirs in this fashion allows the algorithm to uniquely distinguish water usage and availability between hours.

Lastly, the algorithm dispatches hydro plants from least to greatest capacity for maximum system reliability. Sorting hydro plants by their capacity limits acknowledges the fact that aside from energy limitations, large-capacity hydro plants can supply energy in all of the situations that small-capacity hydro plants can; the reverse is not true. (Although the energy constraint is not unimportant, it falls outside the scope of this paper). As hydro reservoirs run low on water, their availability/behavior resembles the behavior of thermal plants because ENSE from previous hours reduces the certainty of water availability (i.e., increases the hydro plant's FOR) for future hours.

### 2.2.2 Step-by-step illustration

The following section graphically explains the proposed chronological algorithm.

First, create an initial inverted load duration curve (ILDC) that describes a single hour of demand with absolute certainty by taking the LDC for a specific hour, inverting the x- and y-axis, and dividing the time axis by $T$. The ILDC resembles a CDF[1] and takes on two probabilities: all demand values less than or equal to the demand at time $t$ have probability 1, and all other demand values have probability 0. Without any loss of generality, this step can incorporate demand uncertainty by modifying the cumulative distribution probabilities attached to each discrete demand value. Figure 2-1 shows the chronological demand for this example, and Figure 2-2 shows the initial ILDC.

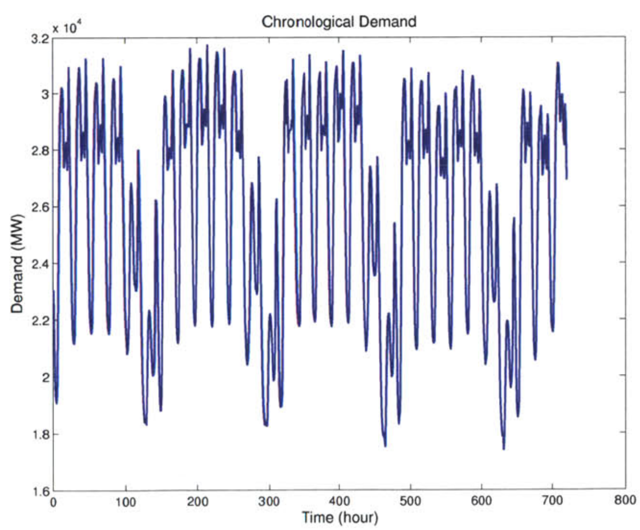


Figure 2-1: Chronological demand

[1]The ILDC is "CDF-like" because it actually describes $P(d \geq D)$, not $P(d \leq D)$.

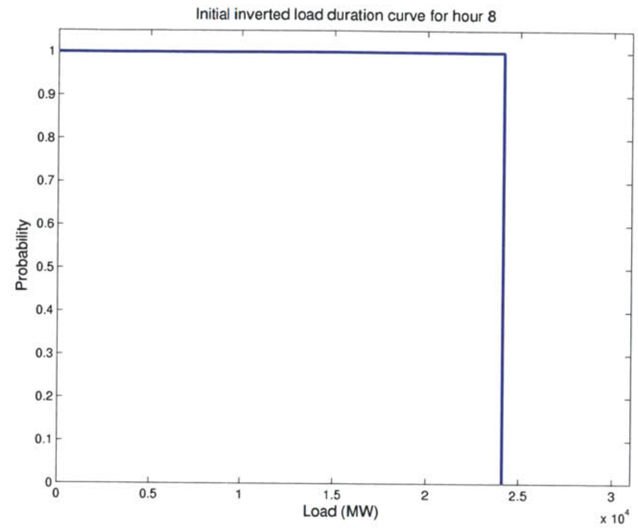


Figure 2-2: The initial ILDC for hour 8

After creating the ILDC, the algorithm dispatches thermal plants using the same convolution as traditional PPCs. As described in the overview above, for each thermal plant, the convolution combines demand CDFs with and without thermal generation after weighting each CDF using the FOR (and the FOR's complement) of the thermal plant. To create the demand CDF with thermal generation, the algorithm removes an amount of demand equal to the capacity limit of the current thermal plant. Figure 2-3 shows an example thermal convolution.

Dispatching a hydro generator resembles dispatching a thermal generator, but requires two convolutions: one for the ILDC and one for the hydro reservoir. This step describes the hydro analogue to the thermal convolution shown in Figures 2-3 and 2-4. The y-intercept of the hydro reservoir CDF, $HCDF(0)$, represents the probability that the reservoir has water and functions much like a thermal plant's FOR. $(1 - HCDF(0))$ represents the probability that the reservoir is empty. Figure 2-5 shows a sample ILDC convolution for hydro dispatch.

Updating the hydro reservoir CDF (HCDFs) completes the hydro dispatch step. HCDFs are CDF-like functions that describe $P(w \geq W)$, where $w$ represents a specific amount of hydro energy. In every hour, a hydro plant faces two scenarios: either its water is needed or not. If there is no ENSE after the thermal dispatch step, then the

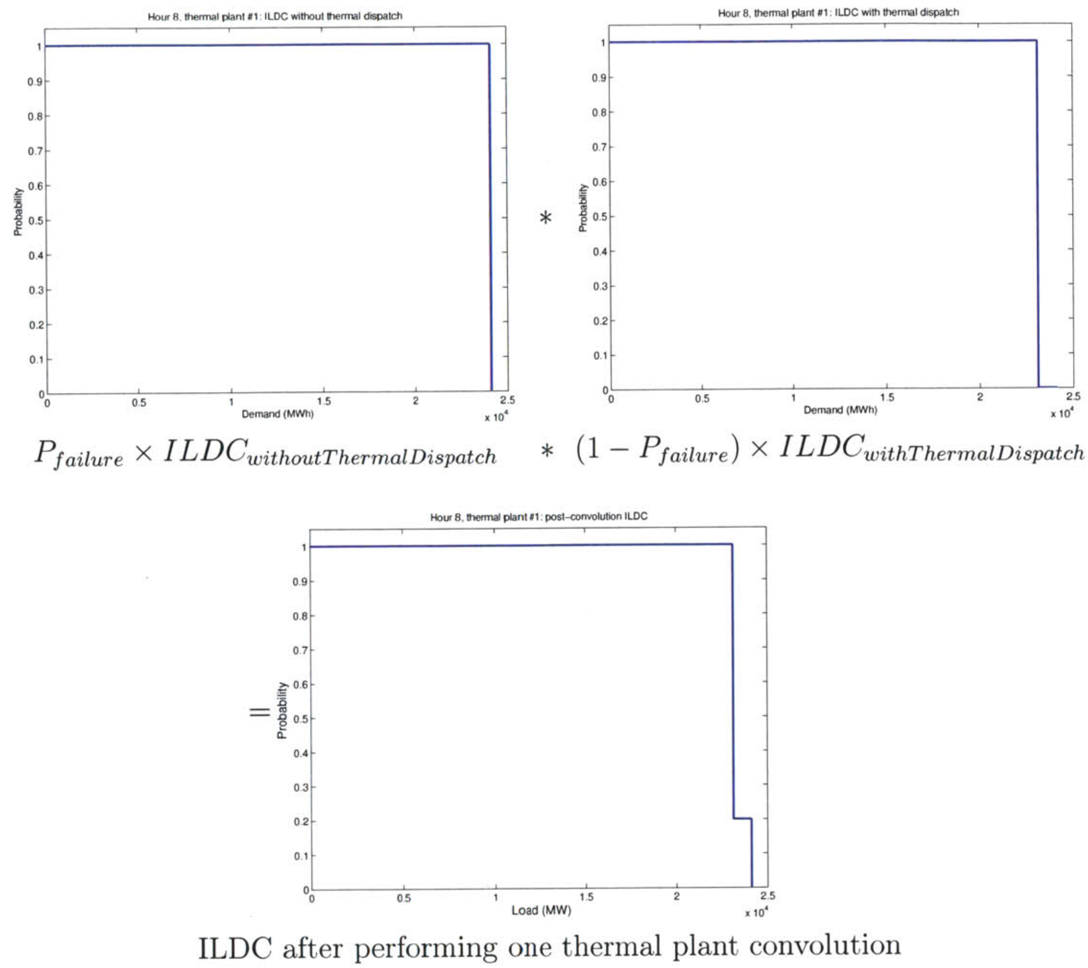


Figure 2-3: Use traditional convolution to dispatch thermal plants

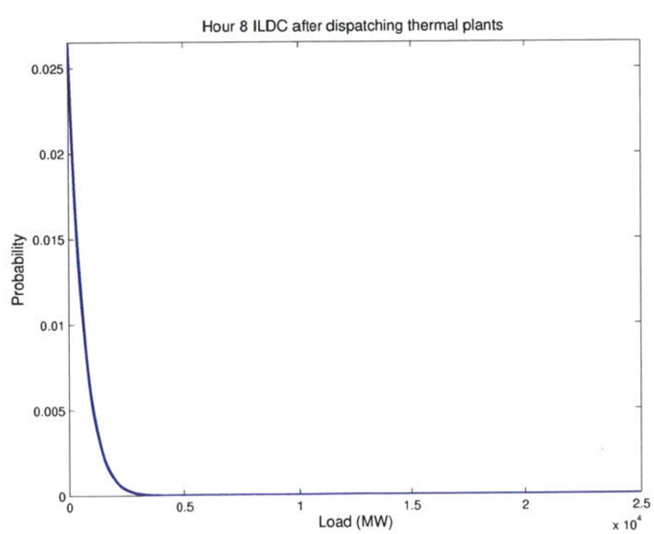


Figure 2-4: Post-thermal-dispatch ILDC for hour 8

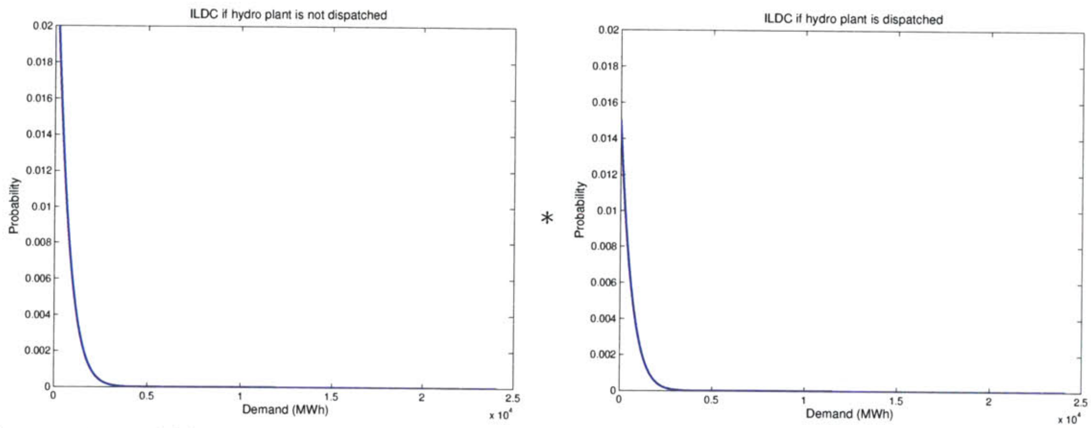


$$(1-HCDF(0)) \times ILDC_{withoutHydroDispatch} * \qquad HCDF(0) \times ILDC_{withHydroDispatch}$$

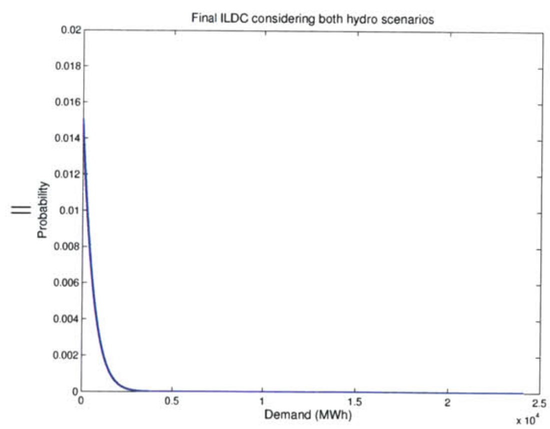


Figure 2-5: Dispatch hydro resources by convolution to cover the remaining ENSE

system does not need to dispatch any hydro generation. If ENSE is strictly positive after dispatching all thermal plants, then the hydro plant should be dispatched at the lesser of (1) its maximum capacity, or (2) the peak system demand. For most power systems, because hydro capacity is a small fraction of the peak demand and therefore also the limiting constraint, most hydro plants will be dispatched at their full capacity. The probability that water is needed is equal to the intermediate LOLP, $ILDC(0)$. The complement is as expected: the probability that water is not needed is $(1-ILDC(0))$. The two hydro scenarios are combined by convolution to obtain the new HCDF, as in the traditional PPC approach for dispatching thermal plants. As the algorithm consecutively dispatches hydro plants within the same hour, because each hydro plant dispatch modifies the ILDC, each successive hydro plant sees a different (and lower) intermediate LOLP. Figure 2-6 illustrates the reservoir convolution graphically. Figure 2-7 shows the final ILDC for hour 8. In each hour, hydro plants are dispatched sequentially, from lowest to highest capacity, to maximize reliability.

After completing thermal and hydro dispatch for one hour, the algorithm saves the final ENSE and LOLP values and iteratively continues onto the next hour in time period $T$. Because the algorithm treats every hour as its own individual reliability problem, this approach preserves the advantages of traditional PPCs while also addressing the chronological problem of representing all hours of demand with one generic cumulative distribution.

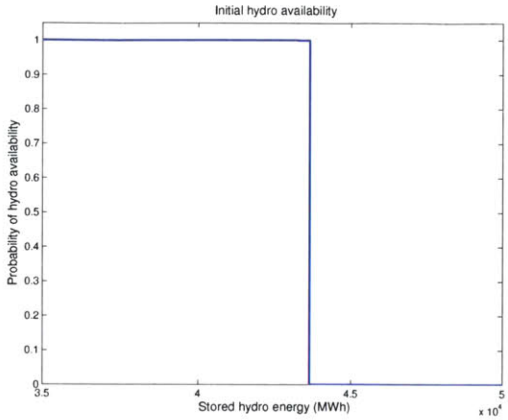


$*$

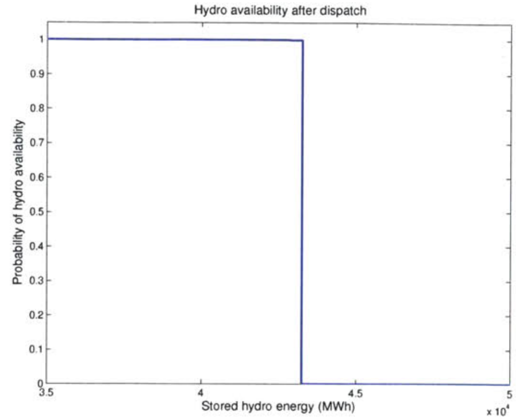


$(1-ILDC(0)) \times HCPD_{withoutHydroDispatch} * \quad ILDC(0) \times HCPD_{withHydroDispatch}$

$=$

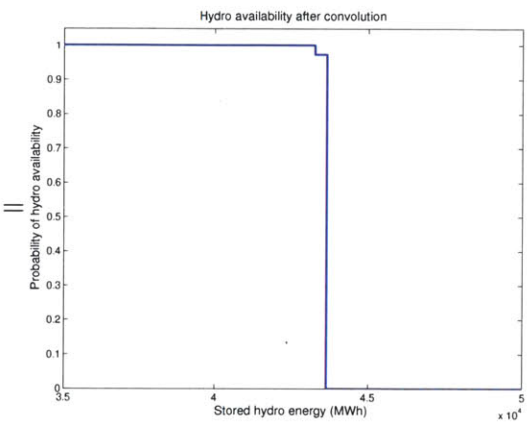


Hour 8, reservoir #1, post-convolution water availability

Figure 2-6: Hydro reservoir convolution

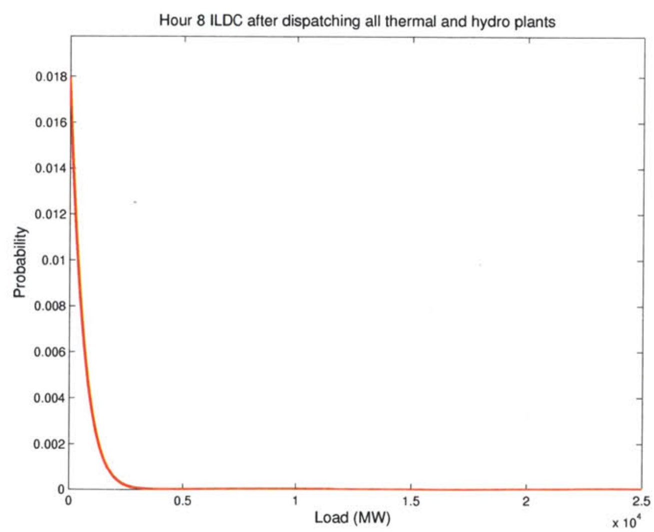


Figure 2-7: Final ILDC for hour 8, after dispatching all thermal and hydro plants

## 2.3 Calculating reliability metrics

### 2.3.1 ENSE and LOLP

Using traditional PPC algorithms, regulators and system operators can evaluate a power system's reliability by calculating its ENSE and LOLP. ENSE represents the expected amount of nonserved energy, and LOLP represents the expected fraction of hours in $T$ that will have a nonzero value for ENSE. The ENSE and LOLP for traditional PPC algorithms are calculated using the following equations:

$$NSE_{ppc} = \sum_{i=0}^{m} ILDC(i) \tag{2.1}$$

$$LOLP_{ppc} = ILDC(0) \tag{2.2}$$

where $i$ represents a point on the demand axis of the ILDC, and $m$ is the peak demand of the ILDC. Given these equations, for the traditional PPC algorithm, the ENSE is the area under the ILDC curve, and the LOLP is the y-intercept of the ILDC.

For the proposed chronological algorithm, the calculations for ENSE and LOLP must take into consideration the individual ILDCs in every hour. Accordingly, the chronologically equivalent formulas are as follows:

$$NSE_{chrono} = \sum_{t=0}^{T} \sum_{i=0}^{m} ILDC_t(i) \tag{2.3}$$

$$LOLP_{chrono} = \frac{\sum_{t=0}^{T} ILDC_t(0)}{T} \tag{2.4}$$

$LOLP_{chrono}$ has a normalizing term, $T$, that $LOLP_{ppc}$ lacks because the PPC algorithm only has one ILDC for $T$, whereas the chronological algorithm has an ILDC for each hour in $T$.

The definitions presented for ENSE and LOLP are valid for all hydro merit order positions.

### 2.3.2 Probability of at least one failure

Lastly, these algorithms can also produce values for the probability of having at least one hour of failure in time period T. The derivation is as follows. In each hour of the chronological dispatch model,

$$ILDC_t(0)$$

is the probability of ENSE exceeding zero in that hour. The complement of this probability,

$$1 - ILDC_t(0)$$

is the probability that there is no ENSE in that hour. Taking the product of these complementary probabilities for all hours in $T$,

$$\prod_{t=0}^{T}(1 - ILDC_t(0))$$

gives the probability that no failures occur for time period $T$. The complement of this probability, Equation 2.5, is the probability that at least one failure will occur.

$$P(NSE_T > 0)_{chrono} = 1 - \prod_{t=0}^{T}(ILDC_t(0)) \tag{2.5}$$

The PPC equivalent simply uses the same LOLP value for every hour:

$$P(NSE_T > 0)_{ppc} = 1 - \prod_{t=0}^{T}(1 - LOLP_{ppc}) = 1 - (1 - ILDC(0))^T \tag{2.6}$$

The definitions presented for the probability of at least one failure are only valid for (1) completely thermal scenarios and (2) the special case of dispatching LEPs after all thermal plants. Calculating the probability of at least one failure when hydro plants are not dispatched at the end of the merit order requires determining whether a drop of water used in hour $t$, even if it isn't dispatched at the end of the merit order, contributes to reliability. (In turn, this step requires complicated conditional probabilities.) Dispatching hydro at the end of the merit order places an upper bound on a system's reliability estimates.

## 2.4 Calculating costs and prices

In the proposed chronological algorithm, the amount of electricity that each generator produces is probabilistic. Consequently, the costs and profits (or losses) that a plant owner incurs are also probabilistic, and the comparison of costs and prices for all algorithms requires calculating expected values.

### 2.4.1 Calculating the expected generation cost

In every hour $t$, each plant $p$ produces

$$E[\text{generation}]_{t,p} = \sum_{i=0}^{m}[ILDC_{t,p-1}(i) - ILDC_{t,p}(i)] \tag{2.7}$$

where (as in Equation 2.1) $i$ represents a point on the demand axis of the ILDC, and $m$ is the peak demand of the ILDC. $ILDC_{t,p}$ indicates the current ILDC for hour $t$ after dispatching plant $p$; $ILDC_{t,0}$ represents the original ILDC for hour $t$. The difference on the right-hand side of Equation 2.7 is the difference of the areas under the ILDC curves, pre- and post-dispatch of plant $p$. Combining plant $p$'s expected generation and variable cost, $cost_p$, gives the expected cost for plant $p$ in hour $t$:

$$\begin{aligned} E[\text{hourly cost}]_{t,p} &= E[\text{generation}]_{t,p} \times cost_p \\ &= \sum_{i=0}^{m}[ILDC_{t,p-1}(i) - ILDC_{t,p}(i)] \times cost_p \end{aligned} \tag{2.8}$$

Summing Equation 2.8 over all hours gives the total expected cost for a single plant $p$ in time period $T$:

$$E[\text{total plant cost}] = \sum_{t=1}^{T}\left[\sum_{i=0}^{m}[ILDC_{t,p-1}(i) - ILDC_{t,p}(i)] \times cost_p\right] \tag{2.9}$$

And, summing over all plants gives the total expected system cost:

$$E[\text{total system cost}] = \sum_{p=1}^{P}\left[\sum_{t=1}^{T}\left[\sum_{i=0}^{m}[ILDC_{t,p-1}(i) - ILDC_{t,p}(i)] \times cost_p\right]\right] \quad (2.10)$$

where $P$ represents the last plant. The total expected system cost, as illustrated in Equations 2.7 through 2.10, only depends on the evolution of the ILDC after each thermal plant dispatch in every hour.

### 2.4.2 Calculating the expected revenue for each plant

#### Marginal probabilities

Calculating a plant's expected revenue requires considering the scenario that the plant is the marginal unit, as well as all of the scenarios that another plant later in the merit order is marginal. First, for a given hour $t$,

$$P(\text{NSE is the marginal technology}) = ILDC(0)_{t,P} = LOLP_t \quad (2.11)$$

The probability that a thermal plant is marginal is the complement of Equation 2.11:

$$P(\text{any generating plant is marginal}) = 1 - LOLP_t$$

Working backward, the last generating plant, $P$, has the following probability of being marginal:

$$P(\text{the last plant, } P\text{, is marginal}) = ILDC(0)_{t,P-1} - LOLP_t$$

More generally, the probability that any plant $p$ is marginal in hour $t$ is:

$$P(\text{plant } p \text{ is marginal in hour } t) = ILDC(0)_{t,p-1} - ILDC(0)_{t,p} \quad (2.12)$$
$$ILDC(0)_{t,0} = 1$$

Lastly, because either a generating plant or NSE sets the marginal price, these probabilities must sum to 1:

$$\sum_{p=1}^{P}[ILDC(0)_{t,p-1} - ILDC(0)_{t,p}] = 1$$

### Marginal prices

The marginal unit probabilities calculated in Equation 2.12 represent the likelihood that plant $p$ sets the marginal price in hour $t$. Each plant observes a unique marginal price because if plant $p$ is generating electricity, then no plant beneath plant $p$ in the merit order can set the marginal price. Therefore, the expected marginal system price in each hour for each plant $p$ is:

$$E[\text{marginal system price}]_{t,p} \tag{2.13}$$

$$= E[\text{marginal system price} \mid \text{plant } p \text{ is generating}]_t$$

$$= \frac{\sum_p^P [(ILDC(0)_{t,p-1} - ILDC(0)_{t,p}) \times cost_p]}{ILDC(0)_{t,p-1}} \tag{2.14}$$

### Expected revenues

Combining the generation for each plant (Equation 2.7) and the expected marginal system price (Equation 2.13) gives the expected revenue for plant $p$ in hour $t$:

$$E[\text{revenue}]_{t,p} \tag{2.15}$$

$$= E[\text{marginal system price}]_{t,p} \times E[\text{generation}]_{t,p}$$

$$= \frac{\sum_p^P [(ILDC(0)_{t,p-1} - ILDC(0)_{t,p}) \times cost_p]}{ILDC(0)_{t,p-1}} \times \sum_{i=0}^{m}[ILDC_{t,p-1}(i) - ILDC_{t,p}(i)]$$

Lastly, summing across all hours gives the expected revenue for each plant:

$$E[\text{total revenue}]_p \tag{2.16}$$
$$= \sum_{t=1}^{T} E[\text{revenue}]_{t,p}$$
$$= \sum_{t=1}^{T} \left[ \frac{\sum_p^P [(ILDC(0)_{t,p-1} - ILDC(0)_{t,p}) \times cost_p]}{ILDC(0)_{t,p-1}} \times \sum_{i=0}^{m} [ILDC_{t,p-1}(i) - ILDC_{t,p}(i)] \right]$$

### 2.4.3 Calculating expected profits

Trivially, the difference between the revenue and cost equations (Equations 2.16 and 2.9) gives the expected profit (loss) for each plant:

$$E[\text{total profit}]_p = E[\text{total revenue}]_p - E[\text{total cost}]_p \tag{2.17}$$

## 2.5 Comparing reliability estimates between algorithms

This section highlights the differences between the traditional PPC model and the proposed chronological PPC model by comparing reliability estimates from both for a case study power system. The case study power system contains 87 thermal plants, 19 hydro plants, and 720 hours of demand data. Figure 2-8 shows the system's LDC and optimal hydro coverage, assuming no thermal failures. This system has 31888 MW of thermal capacity, 9649 MW of hydro capacity, and a peak demand of 31728 MW.

For this reliability study, both the traditional and chronological PPC algorithms dispatch all thermal plants first (starting with the least expensive unit), followed by all hydro plants (starting with the lowest capacity plant). Abstracting hydro plants to LEPs, a comparison of the pre- and post-LEP dispatch reliability metrics reveals the contribution of LEPs to system reliability.

Table 2.1 contains the pre-LEP and post-LEP simulation results for ENSE, LOLP, and probability of at least one failure. Of particular and immediate interest, the estimates of ENSE and LOLP after only dispatching thermal plants differ between the

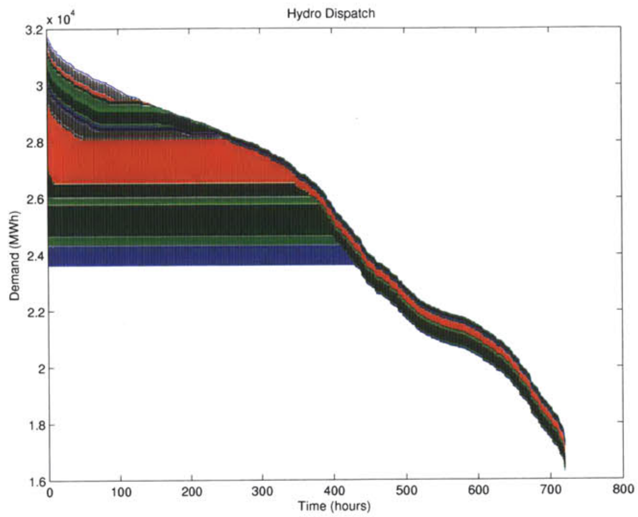


Figure 2-8: Peak-shaving operations on a load duration curve

two algorithms. This discrepancy occurs because the initial ILDCs are not identical. The chronological algorithm contains $T$ total ILDCs, and each ILDC initially describes demand in hour $t$ with complete certainty. The PPC algorithm, on the other hand, contains one ILDC that initially takes on the value of the inverse demand function. The difference between each algorithm's treatment of demand and time explains why the two algorithms report different amounts of ENSE-coverage by LEPs in the system, despite the systems having identical generation plants and capacities.

| | PPC | Chrono. | $\Delta$(PPC - Chrono.) |
|---|---|---|---|
| **Pre-LEP ENSE (MWh)** | 455680 | 452360 | 3320 |
| **Post-LEP ENSE (MWh)** | 1.9534 | 2.9245 | -0.9711 |
| **difference (MWh)** | 455678 | 452357 | – |
| | | | |
| **Pre-LEP LOLP** | 0.3434 | 0.3421 | 0.0013 |
| **Post-LEP LOLP** | 6.376e-6 | 9.2827e-6 | -3e-6 |
| **difference** | .3434 | 0.3421 | – |
| | | | |
| **Pre-LEP** $P(NSE_T > 0)$ | 1 | 1 | 0 |
| **Post-LEP** $P(NSE_T > 0)$ | 0.0046 | 0.0067 | -0.0021 |
| **difference** | 0.9954 | 0.9933 | – |

Table 2.1: PPC versus chronological dispatch results

The post-LEP dispatch results show that the traditional PPC consistently overestimates the power system's reliability (i.e., the traditional PPC consistently underestimates ENSE, LOLP, and the probability of at least one failure) compared to the chronological algorithm. The rows labeled "difference" show the amounts of ENSE, LOLP, and probability-of-at-least-one-failure reduction that can be attributed to LEP generation. Compared to the traditional PPC results, the chronological algorithm predicts that the LEPs will be able to cover 0.9711 MWh less ENSE, that the system has a 3e-6 greater LOLP, and that the system has a 0.21% greater chance of experiencing at least one failure for this particular month of demand. In summary, the traditional PPC simultaneously overestimates total reliability and underestimates LEP contributions to system reliability relative to the chronological PPC.

Lastly, Figure 2-9 plots the changes to ENSE, LOLP, and the probability of at least one failure as time progresses for the chronological algorithm. Because the traditional PPC algorithm treats every hour generically, equivalent graphs for the PPC algorithm would look like horizontal lines that take on the post-LEP values in Table 2.1. As expected, as time increases and the hydro reservoirs start to run out of water, system reliability declines.

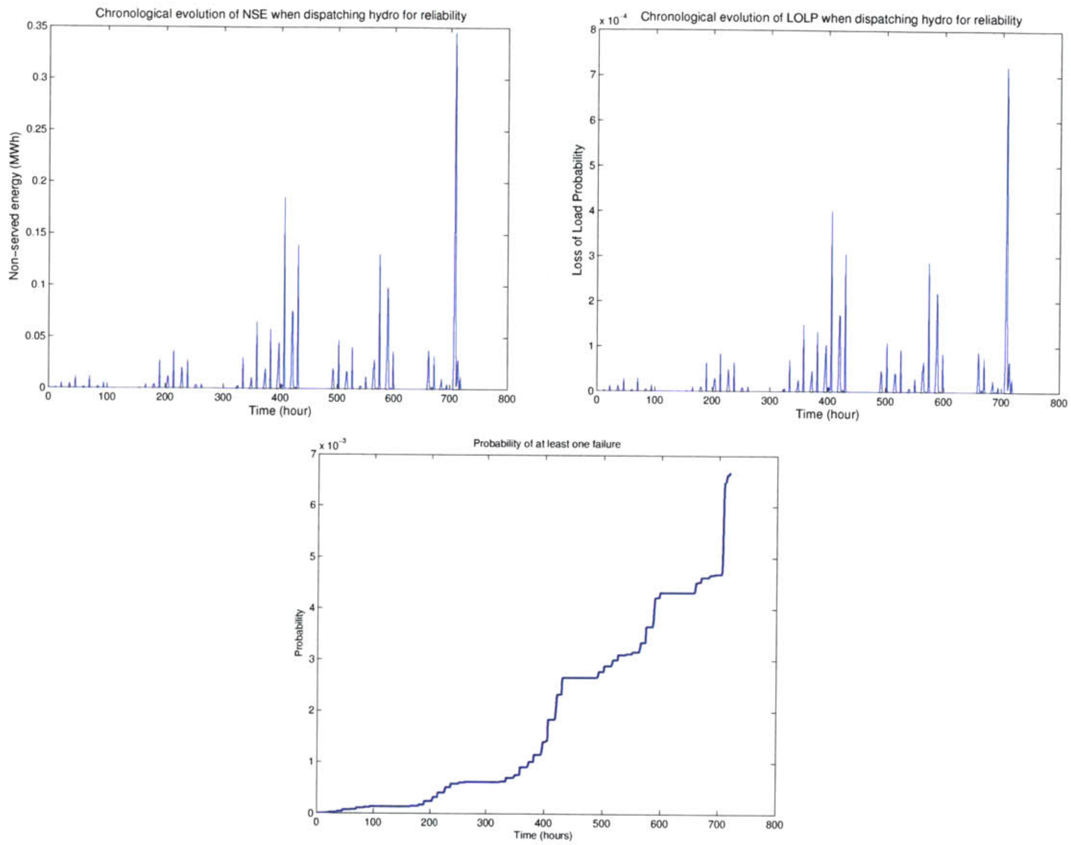


Figure 2-9: Chronological tracking of reliability metrics

## 2.6 Implications for traditional PPC algorithms

As illustrated by the traditional PPC equations for ENSE, LOLP, and probability-of-at-least-one-failure from Section 2.3, traditional PPCs treat every hour generically and then scale up results for that hour to obtain metrics for a week, month, or year. For power systems with mostly thermal units, PPCs reasonably approximate system operations because thermal availability comprises the greatest source of uncertainty, and the convolution operation adequately captures this source of uncertainty. Consequently, the assumption that all hours are the same in thermal-dominated power systems is valid to a first approximation. However, as shown by the results in the previous section, traditional PPCs optimistically overestimate reliability metrics because they cannot distinguish one hour from the next. Because every hour in a PPC model shares the same ILDC, traditional PPC models cannot consider alternative uses for resources that have chronological dependencies. This chronological challenge reduces the usefulness of non-chronological PPC models in power systems with LEPs.

THIS PAGE INTENTIONALLY LEFT BLANK

# Chapter 3

# Exploring LEP costs and contributions to system reliability

In the reliability case study from the previous chapter, the chronological PPC algorithm dispatched hydro plants as generators-of-last-resort to obtain lower-bound (best possible) estimates of the power system's ENSE and LOLP. However, dispatching hydro plants at the end of the merit order[1] also greatly increased the total system cost because the system spilled water that it could have otherwise used to displace expensive thermal generation. This chapter explores the dynamics between total system cost and reliability by varying the merit order position for hydro plants. As appropriate, comparisons are made to results from the traditional PPC model. Without any loss of generality, the conclusions about cost and reliability for hydro plants should also hold for other LEP technologies.

[1]The term "merit order" refers to the ranking that a power system follows when deciding which plants to dispatch first. In systems with electricity markets, market operators determine the merit order by sorting plants from lowest bid to highest bid. In principle, plants that are dispatched earlier in the merit order have lower variable costs than plants that are dispatched later in the merit order.

## 3.1 LEP dispatch methods

### 3.1.1 Peak shaving dispatch

Assuming that thermal plants have perfect availability, system operators can minimize total cost by dispatching LEPs when marginal prices are at their highest. In this "peak-shaving" pattern (shown in Figure 2-8), because LEPs have low variable costs, each megawatt-hour of electricity from an LEP tends to displace a more expensive megawatt-hour from another technology. Peak-shaving, however, reflects a purely economic objective. If the NSE price in a power system were set with perfect information about demand, thermal plant availability, and the desired level of reliability, then the marginal prices that emerge during times of scarcity should adequately encourage the necessary capacity investments. As most systems do not set the price of NSE absolutely correctly or do not allow generators to bid the full NSE price (for many reasons, including (1) calculating the correct NSE price is difficult, and (2) high electricity prices are politically and socially unpopular), the price signals that emerge from energy markets do not typically reflect the potential reliability premium that LEPs could command because of their ability to serve as generators-of-last-resort.

### 3.1.2 Dual objective economic-reliability dispatch

Given that market distortions such as price caps affect the price signals that LEPs receive in energy markets, if regulators want to encourage LEPs to contribute to greater system reliability, they might consider dispatching plants under a dual economic/reliability objective. Under such a scheme—for exampling, minimizing total cost given an explicit level of ENSE and LOLP—the price signals for LEPs may significantly change. However, the state-of-the-art for this category of PPC algorithms still requires significant simplifying assumptions about generation units and chronology.

For example, in [8, Gonzalez et al. 2005], the authors split a year into identical-length subperiods and use a PPC model to determine the ENSE in each subperiod. The ENSE in each subperiod directly dictates how much water to allocate for reliability dispatch. If water in excess of what is required for reliability dispatch exists, the algorithm optimally chooses the peak-shaving allocation in each subperiod that minimizes cost. To perform this optimization, the algorithm makes the following simplifications and assumptions:

1. The algorithm treats demand as a single LDC, discarding potentially useful chronological information.

2. A single, monolithic hydro plant represents all of the hydro resources in the system.

3. The authors' hydro allocation scheme assumes a direct correlation between demand and marginal price for any generation sold into the energy market.

4. The algorithm treats demand and initial reservoir levels deterministically.

5. The algorithm does not consider inflows.

As in the [8, Gonzalez et al. 2005] algorithm, most PPC algorithms make at least one of the above assumptions in exchange for computational simplicity. However, because of the chronological challenges described in Chapter 2, each of these assumptions covers an important aspect of hydro plant/LEP operations that can materially impact a model's results. The proposed chronological PPC algorithm in Chapter 2 removes the first three assumptions enumerated above. Additionally, the chronological PPC directly allows for probabilistic representations of hourly demand, hourly reservoir levels, and hourly inflows via modification of the cumulative probability distributions for demand (ILDC) and reservoir levels (HCDF) at every time $t$. By improving the representation of chronological information, the chronological PPC should produce more realistic predictions about the effects of different LEP dispatch patterns on total system cost, revenues, and reliability.

## 3.2 Results

To compare the effects of LEP dispatch on system reliability, this case study calculates the ENSE, LOLP, total system cost, and hydro revenue for every possible hydro position in the merit order using both a traditional PPC and the chronological PPC model. The power system remains the same as the system presented in Section 2.5. Intuitively, as hydro moves higher in the merit order, a power system's reliability should increase because more water remains available to serve unmet demand. However, the increased availability of water occurs at the expense of spilled water and more thermal generation; consequently, a total system cost minimum should appear within these explorations. Given that a power system's total cost depends both on its generation costs and the cost of NSE, the study in this chapter analyzes five different NSE price scenarios ($0, $150, $300, $1000, and $5000 per MWh). As before, without loss of generality, the conclusions about how hydro plants affect system reliability and total cost should also apply for other LEP technologies.

### 3.2.1 Reliability

Both the traditional PPC and chronological models predict that ENSE and LOLP monotonically decline (system reliability increases) as the system dispatches hydro later in the merit order. As expected after the comparison of algorithms in Chapter 2, the PPC algorithm always produces optimistic estimates of system reliability relative to the chronological algorithm. Figure 3-1 shows results from both models. Chapter 4 discusses regulatory instruments that regulators could design, based on these reliability metrics and the information about costs and revenues in the following sections, to motivate hydro/LEP operators to dispatch their plants in a fashion to achieve a target level of ENSE and LOLP.

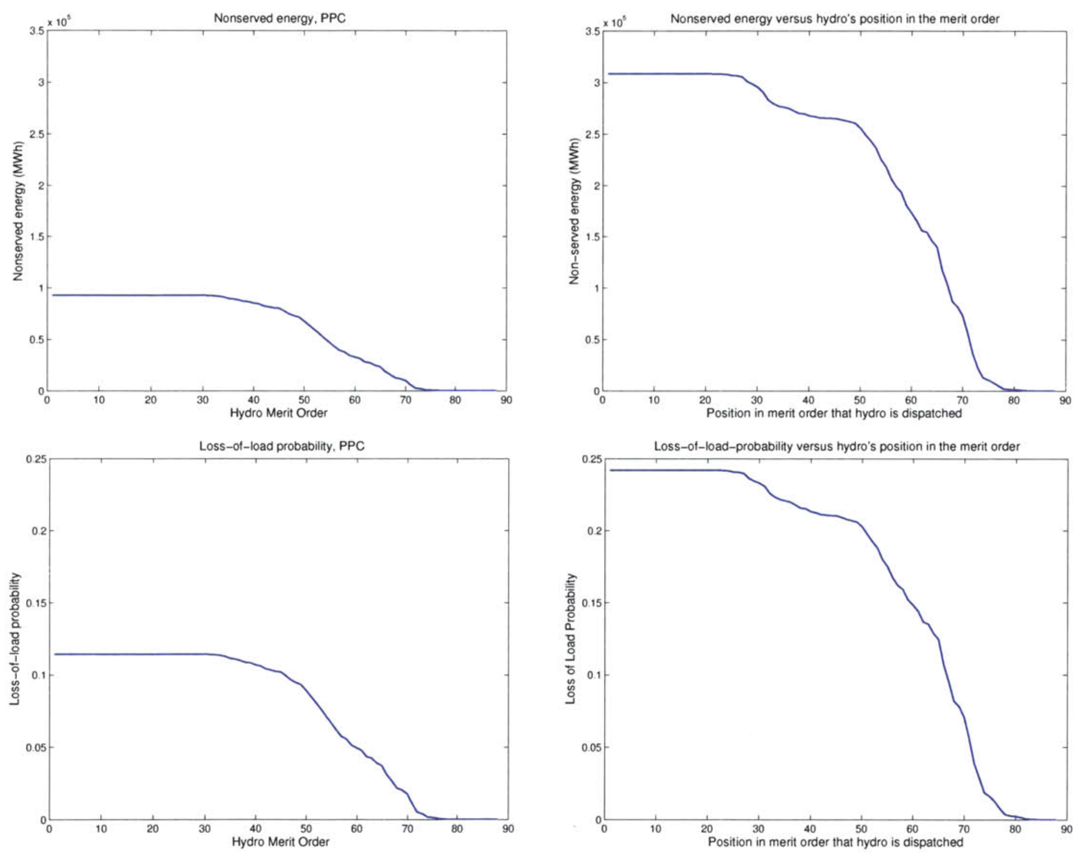


Figure 3-1: Reliability metrics for the PPC (left) and chronological (right) algorithms

### 3.2.2 Total system cost

As shown in Figure 3-2, both algorithms predict that total system costs decline as the system dispatches hydro later in the merit order. The blue trendline represents pure thermal generation costs because NSE has a price of zero. As thermal plants do not have any special chronological attributes in either the chronological PPC or the traditional PPC model, the thermal costs predicted by both algorithms closely resemble one another.

The remaining trendlines illustrate the combined cost of thermal generation and ENSE at different NSE prices. Both algorithms predict that a minimum total cost appears as hydro moves further down the merit order for all but the purely thermal scenario, indicating that rising thermal costs balance declining ENSE costs for all scenarios with a nonzero NSE price. Lastly, in both algorithms, total system costs converge when the system dispatches hydro at the end for all NSE price scenarios because the system has enough combined thermal and hydro generation to meet demand if the system uses hydro's limited energy as a last resort.

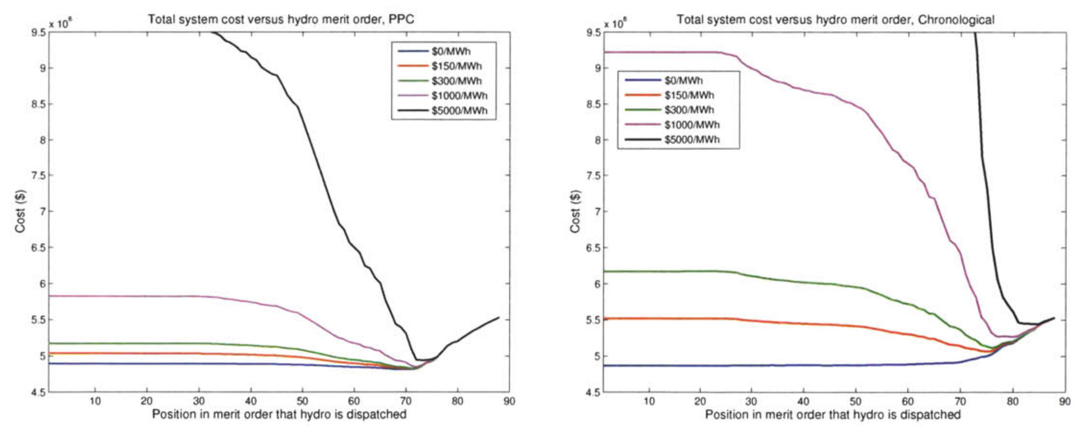


Figure 3-2: Total system cost for the PPC (left) and chronological (right) models

However, the chronological algorithm contains an interesting and different result regarding the position of the optimal hydro position that minimizes total system cost. In the traditional PPC model, the minimum total system cost appears at hydro merit order position 71 regardless of the price of NSE. Additionally, the PPC model predicts that after position 71, the system dispatches enough thermal generation to cover any remaining ENSE with its limited hydro resources. As such, in the PPC model, when the system dispatches hydro after position 71 in the merit order, ENSE costs decline; thermal costs play a dominant role in the total cost; and the total cost for all NSE scenarios quickly converges.

In the chronological model, two important distinctions appear for total cost. First, the chronological model predicts that the system will need to dispatch hydro resources later in the merit order to minimize total cost as the price of NSE increases—i.e., unlike the results from the PPC model, the same hydro merit order position does not minimize total cost for all NSE price scenarios. Second, in the chronological model, total costs do not converge as quickly as in the PPC model. This slower convergence suggests that the amount and cost of ENSE continues to have a nontrivial impact after the system begins to spill water. Lastly, in agreement with the PPC model, the chronological model predicts that dispatching hydro as a generator-of-last-resort allows the system to fully cover its demand with thermal and hydro generation. Consequently, at the last merit order position, thermal costs play a dominant role and ENSE costs decline significantly, resulting in convergence for total cost predictions across all NSE scenarios.

### 3.2.3 Thermal generation

Because the total cost minimums that appear in Figure 3-2 result from balancing declining ENSE costs with rising thermal generation costs, this section explains the effects of moving hydro through the merit order on thermal generators. Although this section shows results from the chronological algorithm, the same analysis and conclusions apply to the PPC algorithm.

Figure 3-3 graphically illustrates the total energy from a baseload thermal plant (left graph) and an intermediate-load thermal plant (right graph) in the system with each change in the merit order for hydro. The generator on the left is the least expensive thermal generator in the system, and the generator on the right is 72th least expensive thermal generator (of 87 thermal plants). The baseload generator experiences no change in energy output as the system moves hydro through the merit order because regardless of the amount of hydro generation, the system has enough demand for the baseload generator to operate at its full output.

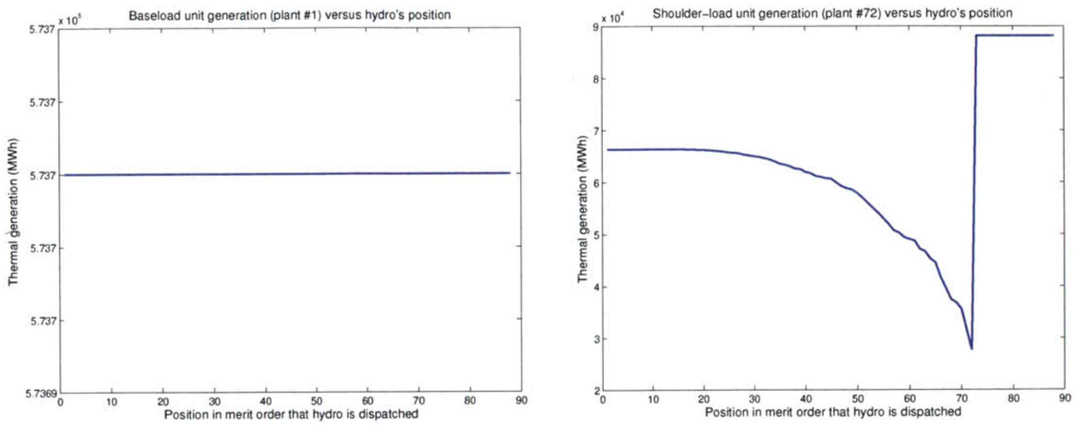


Figure 3-3: Baseload versus intermediate-load thermal unit generation for the chronological model

However, hydro plants do affect the total output of non-baseload plants. As the system moves hydro further in the merit order, the full amount of energy in the hydro reservoir (constrained by capacity limits for hydro) covers demand that otherwise would have been supplied by thermal units above the hydro unit. As such, an intermediate-load unit such as the plant shown on the right in Figure 3-3 will generate less energy as the system moves hydro further up the merit order until the hydro plant overtakes the thermal unit. When this happens, the plant's thermal output returns to depending only on the amount of demand remaining in the system and the plant's capacity limits.

### 3.2.4 Hydro generator revenue

Using the marginal probabilities and expected revenue formulas that were developed in Chapter 2, this section discusses changes in hydro generation and revenue as the system moves hydro units later in the merit order under both the traditional PPC and chronological model. The revenue-related graphs in this section contain many interesting features. First, this section discusses the common peak in revenues around hydro merit order position 71 for both the PPC and chronological algorithm. Then, the remainder of this section analyzes the discrepancies between the two algorithms.

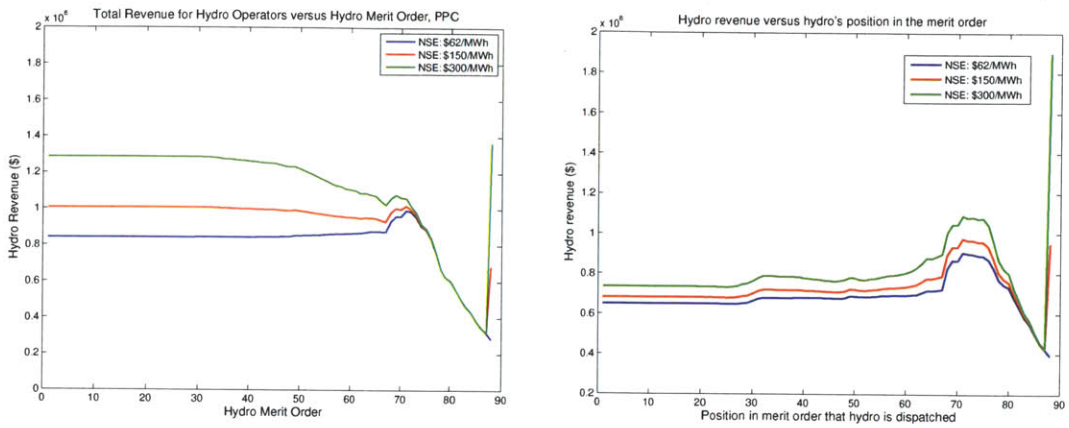


Figure 3-4: Hydro revenues for the traditional PPC (left) and chronological (right) models

### Revenue peaks when the system begins to spill water

Figure 3-4 contains the expected hydro revenues for both algorithms. Both algorithms predict that revenues will reach a local maximum around hydro position 71, when the system begins to spill water. (The spike in revenues at the end of the merit order will be explained shortly.) The increase in revenues around position 71 refers back to the discussion about total costs: as the system spills water, it must rely on more expensive thermal plants to meet remaining demand. As such, marginal prices begin to rise. Hydro merit order position 71 reflects a balancing point for the case study power system. To the left of this position, the price of ENSE dominates the expected marginal system price; to the right of this position, ENSE costs begin to decline as thermal costs rise. As such, hydro plants tend to earn more money when the system dispatches as much of their stored energy as possible.

### Hydro generation discrepancies

An important discrepancy appears between the hydro generation predictions from each algorithm. Figure 3-5 plots the expected hydro generation from the traditional PPC model (left) and the chronological model (right). The PPC model predicts

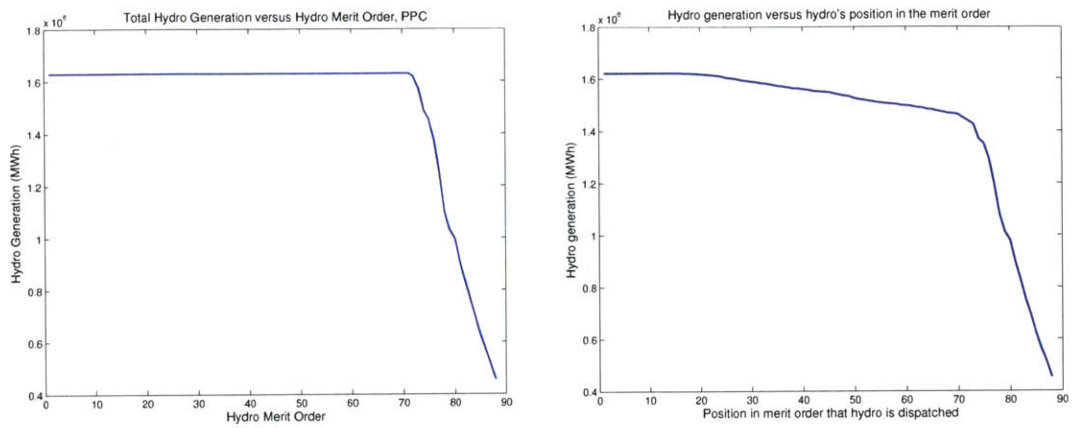


Figure 3-5: Hydro generation for the traditional PPC (left) and chronological (right) models

that hydro generation will remain constant until the system begins to spill water. The chronological algorithm, however, predicts that hydro generation incrementally declines with each increase in merit order position. The rate of decline for hydro generation is initially almost zero over the first 23 positions as hydro displaces baseload plants, then increases from position 23 to position 71 as hydro displaces intermediate-load plants, and finally significantly increases from position 72 to the end as the system spills water. Both algorithms, as shown in the generation charts, predict that the system will spill water after position 71. However, only the chronological algorithm predicts that the opportunities for hydro generators to fully use their water probabilistically declines with each position increment in the merit order.

Figure 3-6 offers a graphical explanation of the chronological algorithm's predictions. The magenta trendline represents hourly hydro generation when the system dispatches hydro first. The green trendline represents hourly hydro generation when the system dispatches hydro last. In the former scenario, as a baseload plant, the hydro plant probabilistically exhausts its entire reservoir approximately 250 hours into the simulation. In the latter scenario, as the absolute last generator, the hydro plant never runs out of water. As hydro moves from serving baseload to serving peak demand, it covers less and less demand. Consequently, with each step further in the merit order, hydro generation probabilistically declines. In the traditional PPC algo-

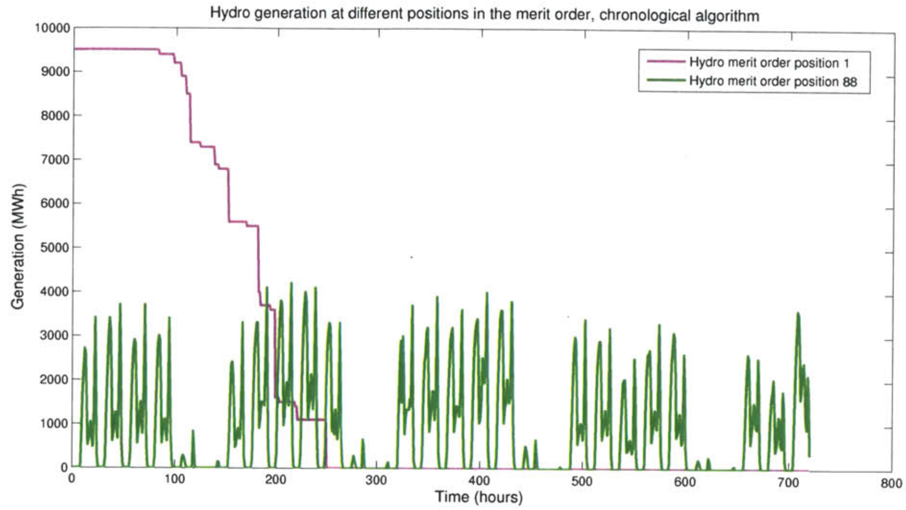


Figure 3-6: Hydro hourly generation, chronological model

rithm, because every hour appears identical, water is available with complete certainty to perfectly cover demand. Eventually, the system has more water than demand and must begin to spill water. For this unrealistic reason, the traditional PPC does not predict the same amount of hydro generation as the chronological algorithm.

### Hydro revenue discrepancies

The two algorithms also disagree about the expected revenues that hydro generators will receive. In the traditional PPC algorithm, although hydro generation stays constant when the system dispatches hydro between merit order positions 1 and 71, the expected revenues from this generation decline because the average expected marginal price declines. Figure 3-4 shows the expected hydro revenues as predicted by the traditional PPC algorithm (on the left) and the chronological algorithm (on the right). Because the traditional PPC algorithm does not distinguish between individual hours, each possible dispatch position for hydro has only one *average* expected marginal system price for hydro generation. In the chronological model, *hourly* marginal system prices exist for each possible dispatch position. Figure 3-7 shows that a direct correlation exists between the amount of ENSE at any given hour and the expected marginal system price for hydro in that hour—the larger the ENSE, the greater the expected marginal system price. These differences between how the two models treat expected marginal system prices explain why the traditional PPC algorithm predicts relatively higher revenues earlier in the merit order, as well as relatively lower hydro revenues later in the merit order, compared to hydro revenue predictions from the chronological algorithm.

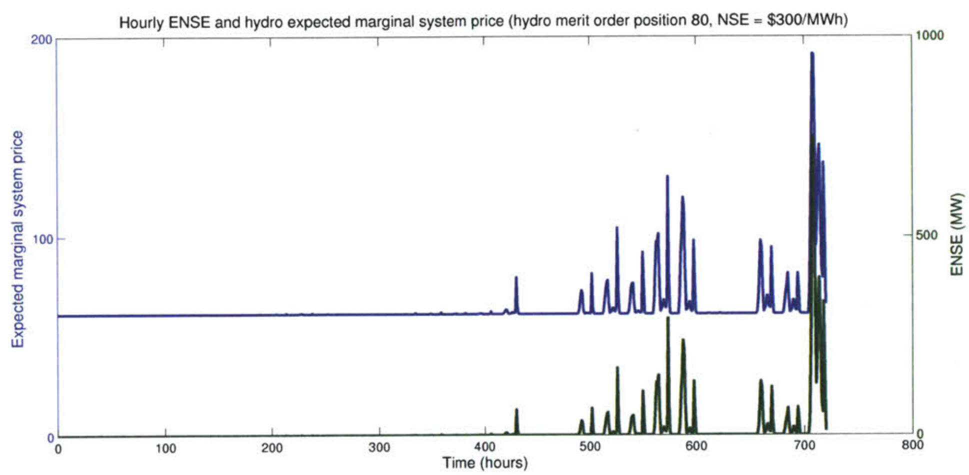


Figure 3-7: Hourly ENSE and hydro marginal system price, chronological model

### Revenue peaks at the end of the merit order

Another interesting result deserves greater explanation in the hydro revenue plots of Figure 3-4. Because hydro operators can bid the price of NSE, because the possible NSE prices are at least one order of magnitude greater than the most expensive thermal plant, and because thermal generation alone cannot cover all hours of demand in this system, the expected marginal price for hydro can dramatically increase when the system dispatches hydro last. Referring back to Equation 2.13, at the end of the merit order, the calculation of the expected price only considers the NSE price with probability 1. The magnitude of the expected revenue increase depends on the price of NSE. Figure 3-4 plots the expected hydro revenues from both models for NSE prices of \$62, \$150, and \$300 per MWh. In both models, the most expensive thermal unit has a variable cost of \$61.60/MWh. The scenarios with an NSE price of \$62/MWh serves as an intuition check for the explanation about why the last expected revenue figure can spike, but does not necessarily have to. Not surprisingly, when the price of NSE is large, both models predict that hydro generators earn the greatest profits when the system dispatches hydro generation at the end of the merit order.

However, realistically, due to competition between hydro operators, operator risk aversion, and demand response from large consumers, marginal prices for electricity in any hour are unlikely to exhibit dramatic increases that would result in the revenue spikes shown in Figure 3-4. Most hydro operators, out of the concern that the few hours of high-priced NSE may not occur exactly as predicted each year (both in quantity and time), will not willingly spill water in hopes of capturing those high prices. Because these hours of NSE represent rare tail events, if hydro operators withheld water in hopes of capturing these prices, and the amount of NSE in a particular year happened to not meet their predictions, they would lose money. As most hydro operators are risk averse to this type of business model, they will manage their reservoirs such that they can reasonably expect to sell electricity and capture most of the highest marginal system prices throughout the year. For the case study power system hydro operators, this risk-balancing, profit-maximizing behavior results in most hydro operators preferring to dispatch at the local optimum (position 71) and not at the end of the merit order.

## 3.3 Summary

The case study in this chapter demonstrates the proposed chronological PPC algorithm as a tool for evaluating the contribution of LEPs to power system reliability. As expected, reliability and the dispatch position for hydro are directly correlated, and dispatching hydro resources later in the merit order reduces a system's ENSE and LOLP. However, the relationship between total system cost/hydro operator revenue and system reliability is not linear. Because total system cost depends on both generation and ENSE, an NSE-price-dependent optimal hydro merit order position exists that minimizes total system cost. Consequently, by default, hydro operators will provide some level of reliability as a normal part of their profit-maximizing behavior. To target lower levels of ENSE and LOLP than these profit-maximizing defaults using hydro resources, regulators will need to provide hydro operators with additional economic incentives.

# Chapter 4

# Regulatory tools for reliability

The case study in the previous chapter demonstrated that a power system can increase its reliability by saving its hydro resources (and, more generally, any limited energy resources) for dispatch later in the merit order. However, in power systems with electricity markets, this type of dispatch behavior rarely maximizes profits. Owners of LEP technologies are unlikely to dispatch their plants to explicitly improve system reliability without additional economic incentives. As described in Chapter 1, the responsibility of securing the supply of electricity usually belongs to regulators because risk-averse market agents, due to market failures such as demand uncertainty and the inability to fully and properly price NSE, will most likely underinvest in capacity. If regulators want LEP owners to act as reliability resources, they will need to compensate LEP owners for their lost revenues.

This chapter updates two common regulatory tools used to address the security of supply problem based on the results from Chapter 3 and the chronological PPC model. The first tool calculates the size of the capacity payment required to compensate LEPs for deviating from their profit-maximizing behavior. The second tool estimates a generator's expected load carrying capability (ELCC), a metric that regulators and policymakers frequently use to determine how much load a specific generator can serve without affecting its overall reliability. These two applications demonstrate the cost and operational insights that the chronological PPC model can contribute to current policy discussions about renewable generation technologies, portfolio standards, public subsidies, and the impacts of different generation mixes on power systems.

## 4.1 Capacity payments

To motivate LEP operators to serve as generators-of-last-resort for hours when a power system is most likely to experience its highest ENSE and LOLP, regulators can offer capacity payments as compensation for the operators' lost revenues. Although capacity payments can take many forms, broadly, they represent stable revenue streams to plant operators in exchange for the operators' guarantees that a fraction of their capacity will remain available to generate electricity as needed. The stability of a capacity payment eliminates some of the demand and price risks inherent to selling electricity. If LEP operators believe that they can earn more money from the combined revenue of capacity payments and proceeds from the energy market, then they will willingly take the capacity payment and hold their limited energy for times of failure.

Regulators can use the chronological PPC algorithm to calculate the minimum threshold capacity payment required to encourage LEP operators to serve as generators-of-last-resort. Continuing with the case study power system, the capacity payment threshold is equal to the difference in revenues that hydro operators earn under their (1) profit-maximizing behavior and (2) reliability-dispatch behavior. To illustrate the dynamics between hydro revenues and reliability, Figure 4-1 stacks the chronological estimates from Chapter 3 for ENSE (left), LOLP (right), and expected hydro revenues on top of each other. The two graphs closely resemble one another because ENSE and LOLP are correlated. As the system dispatches its hydro resources later in the merit order, it has more stored energy to handle potential thermal plant failures. Consequently, both ENSE and LOLP monotonically decline with incremental increases in hydro's dispatch position.

By default, LEP operators will "supply" the system with an amount of reliability equal to the values of ENSE and LOLP at their profit-maximizing merit order position

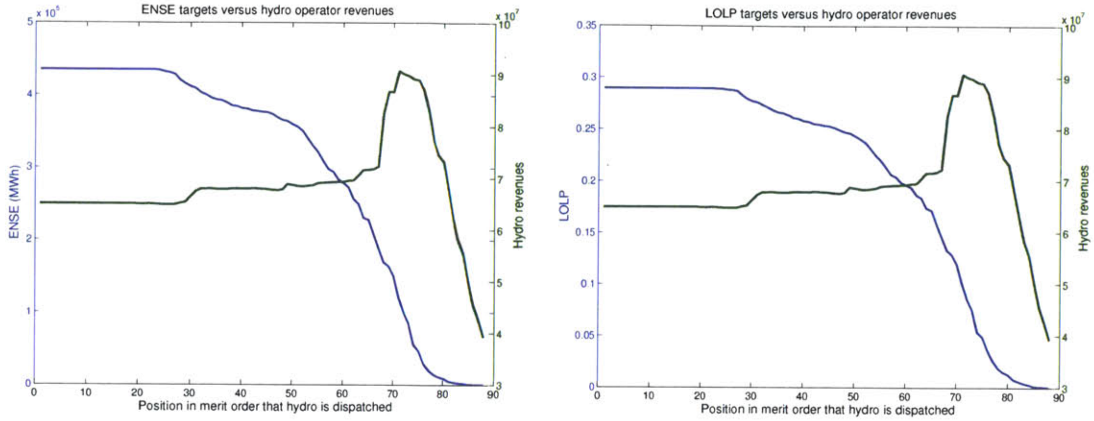


Figure 4-1: Calculating capacity payments based on ENSE, LOLP, and expected hydro revenues

for "free." In this case study, the hydro operator earns the most revenue when it dispatches its plants at merit order position 71. If the regulator decides to target a lower ENSE or LOLP value by modifying the dispatch order of its hydro plants, then it must dispatch its hydro resources after position 71.

However, because hydro revenues begin to decline after position 71, operators will unlikely dispatch their plants at later positions without additional incentives. Figure 4-2 shows the relationship between the case study system's ENSE and LOLP values versus lost hydro revenues as cost-reliability frontiers. The last point on each trendline (at hydro dispatch position 88) represents the default level of ENSE and LOLP that the algorithm predicts for the system under the hydro operator's profit-maximizing behavior. Moving from right to left on the trendlines, system reliability increases because ENSE and LOLP (tracked on the x-axes) decline. The y-axes show how much revenue hydro operators lose with each reliability improvement due to their later dispatch in the merit order. The revenue differences were directly calculated from the hydro revenue data shown in Figure 4-1. To motivate hydro operators to help reduce system ENSE or LOLP (i.e., to move toward the origin in either graphs), regulators will have to pay hydro owners the revenue differences as indicated on the y-axes.

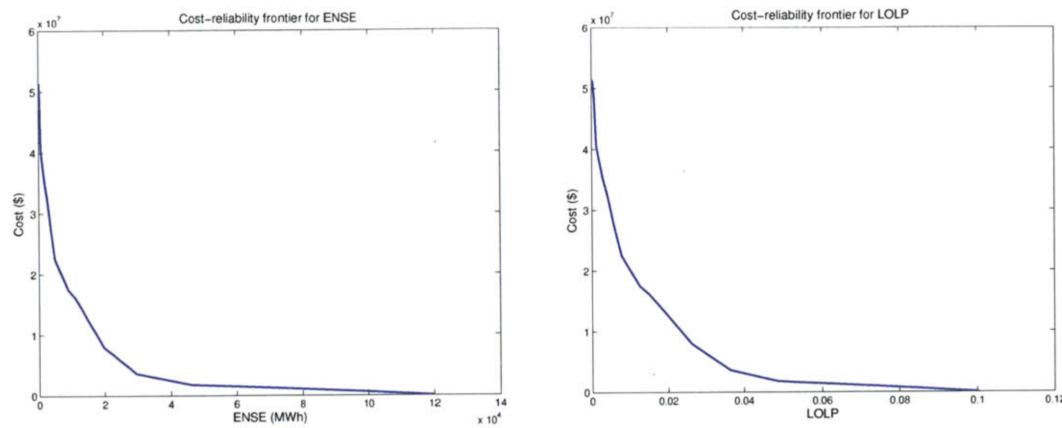


Figure 4-2: ENSE and LOLP cost-reliability frontiers

## 4.2 Calculating a generator's ELCC

To evaluate the individual contribution of a generator to system reliability, regulators can estimate that generator's ELCC. As explained in detail by [12, Milligan 2008], the ELCC metric "measure[s] the additional load that [a] system can supply with the particular generator of interest, with no net change in reliability."

Generally, calculating a generator's ELCC requires several iterative steps. First, the ELCC model removes the generator of interest and adjusts the system's load to achieve a desired baseline reliability metric (for example, a target LOLP value). Then, in a second iteration, the model adds the generator of interest back into the generation mix and calculates a new (and lower) LOLP. Afterward, the model removes the generator and iteratively adds in a benchmark generator—for example, a peaker plant with a 5% failure rate—until the system LOLP returns to the second iteration LOLP. The total capacity of the inserted benchmark units determines the ELCC of the removed generator. Critics of this technique have expressed concerns about the assumptions required for the benchmark unit; however, as long as all generators in the same system are compared against the same benchmark unit, a fair evaluation of each generator's relative capacity credit should be possible. [12, Milligan 2008]

As ELCC remains a popular and useful (albeit not perfect) metric to evaluate firm capacity, this section explains how regulators can estimate a generator's ELCC using the chronological PPC algorithm. The specific estimation presented here, for simplicity, makes two adjustments to the approach described above: instead of initially adjusting the load to target a specific reliability metric and assuming a benchmark unit, this calculation simply compares two final ILDCs from the chronological PPC algorithm. In the first scenario, the case study power system has all of its generators. In the second scenario, the case study power system has removed a generator of interest. This example demonstrates the ELCC calculation for a single hydro plant in the case-study power system.

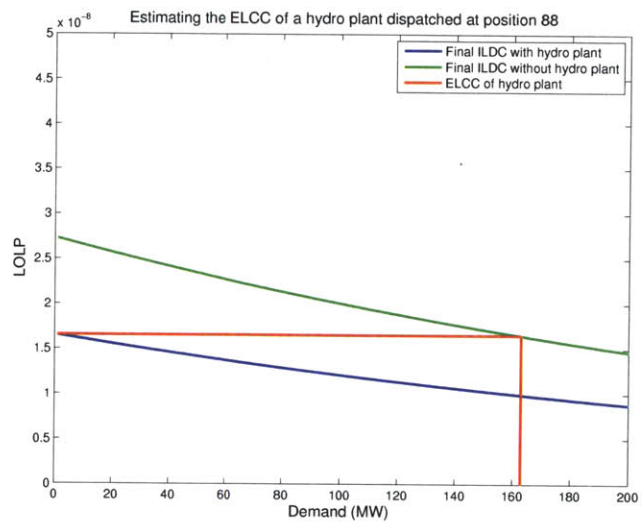


Figure 4-3: ELCC estimate using the chronological PPC model

Figure 4-3 graphs a portion of the two final ILDCs. The blue line represents the full system's final ILDC; the green line represents the system's final ILDC without the hydro unit. Intuitively, removing a generator increases the system's LOLP because it has one less generator to cover demand. Therefore, as expected, the green ILDC (which represents the system with one less hydro plant) is greater than the blue ILDC at every demand point. As before, the y-intercepts of the ILDCs represent their systems' LOLPs. The difference in LOLPs between the two ILDCs represents the change in reliability that the system can attribute to the removed hydro plant. Linking this change in LOLPs to capacity, the demand level at which the without-hydro system achieves the same LOLP value as the full system represents the ELCC of the removed hydro unit. This technique assigns the hydro plant, which has a capacity limit of 391 MW and 4364.3 MWh of reservoir energy, an ELCC of 163 MW if the system dispatches this plant at the end of the merit order. Figure 4-3 graphically shows the relationship between LOLP and ELCC in red. As the availability of hydro plants in this model depends only on their reservoir energy, the plant's ELCC of 163 MW relative to its total capacity of 391 MW suggests that the plant does not have much energy to serve later hours of ENSE.

As noted by [12, Milligan 2008], many methods exist for calculating a generator's ELCC. The chronological PPC-based method presented in this section provides a quick probabilistic estimate with a more accurate representation of limited energy plants compared to traditional PPC methods.

## 4.3 Summary

The tools developed in this chapter demonstrate regulatory applications of the chronological PPC model that allow regulators to place a value on system reliability, as well as to determine how to appropriately compensate LEP operators for their contributions to system reliability. The chronological PPC extends naturally to support the calculation of ENSE/LOLP-based metrics, such as a generator's ELCC, because the final ILDC from the chronological algorithm is a cumulative distribution function that looks identical to the final ILDC from other traditional PPC models. As such, the outputs of the chronological PPC model should generalize well as inputs for other PPC-based models while removing the former assumption about identical and generic units of time. As regulators and policymakers continue to focus on the integration of renewable resources that exhibit strong chronological behaviors in power systems, tools such as the chronological PPC model can contribute meaningful information about the impacts of different generation mixes and policies to the discussion.

# Chapter 5

# Summary & Conclusions

Before electricity markets, vertically integrated utilities controlled all of the operations and investments within electric power systems. These monolithic units ran regulated monopolistic businesses. They ensured the security of electricity supply by often erring on the side of overinvestment to avoid the political and social repercussions of electricity failures. When power systems began developing electricity markets to encourage greater economic efficiency, power systems lost the vertically integrated utility as a central planner. Risk-averse market agents, faced with market failures such as a lack of perfect information about demand, tended to err on the side of underinvestment to ensure the viability of their businesses. Consequently, in the transition to electricity markets, because electricity failures still have political and social consequences, most regulators found themselves responsible for developing market rules that promote competition and ensure the security of electricity supply.

### Traditional convolution-based PPC models

To address the reliability problem, regulators use many analytical tools to understand the physical and economic operations of power systems and markets. The traditional PPC models described in Chapter 2 allowed regulators to estimate the reliability of a power system. Historically, convolution-based PPC models reasonably approximated systems with thermal-dominated generation mixes and demand-driven uncertainty. However, these methods do not sufficiently represent nonthermal generation sources such as hydro or LEPs because of their generic representation of time. The traditional PPC model treats every hour generically and scales the hourly result to obtain weekly, monthly, or annual metrics. This implicitly requires certain plants, such as LEPs, to take on unrealistic dispatch behaviors.

To address these challenges, the proposed chronological PPC model in Chapter 2 breaks the traditional reliability problem into individual, hourly reliability problems. For the case study power system with 87 thermal units, 19 hydro units, 31888 MW of thermal capacity, 9649 MW of hydro capacity, 720 hours of demand data, and a peak demand of 31728 MW, the chronological algorithm estimated that the ENSE should be 0.9711 MWh higher than the traditional PPC model's estimate; that the LOLP should be 3e-6 higher; and that the probability of at least one failure should be 0.0021 higher. Compared to the chronological algorithm, the traditional PPC model consistently overestimates system reliability (i.e., underestimate a power system's ENSE and LOLP).

### The reliability contribution of LEPs

LEPs can serve as generators-of-last-resort to improve system reliability. However, dispatching LEPs at the end of the merit order also greatly increases a power system's total cost. In Chapter 3, the exploration of hydro dispatch position in the economic merit order revealed that (as expected) reliability and dispatch position are directly correlated. The calculation of marginal unit probabilities, prices, and revenues in the chronological PPC model are, to the author's best knowledge, novel research contributions.

Interpreting the reliability and cost results, hydro operators will provide a default level of reliability as a normal part of their profit-maximizing behavior. However, the relationship between hydro operator revenues and system reliability is not linear. As the system begins to dispatch hydro further in the merit order, its hydro plants begin to spill water (i.e., at some point, the system operator dispatches so many thermal plants before relying on its hydro resources that the hydro plants cannot possibly use up all of their water by the end of the simulation period). Hydro operators begin to lose money as they spill water. Consequently, to target lower levels of ENSE and LOLP than the profit-maximizing defaults offered by hydro plants, regulators will need to provide additional economic incentives to encourage hydro operators to dispatch later in the merit order.

The chronological PPC model also extends easily to help design these economic incentives. Chapter 4 demonstrated the calculation of capacity payments for LEPs as cost-reliability frontiers based on the hydro operator's revenues and system ENSE/LOLP, as well as how to estimate a generator's firm capacity (for any generation technology, not just hydro/LEPs) using the well-known ELCC metric.

### Future work

As investment in renewables continues to grow and regulatory and political discussions about electricity shift toward integration concerns and climate change, the chronological PPC developed in this thesis can offer useful insights about the costs and operational impacts of different generation mixes. Because the chronological model produces outputs in the same form as other traditional PPC models (ENSE in units of energy, LOLP values as proper probabilities, and ILDCs as CDF-like functions), other models that build metrics off of PPC outputs such as ENSE and LOLP (for example, the ELCC metric) can also directly benefit from the chronological model's treatment of time. Tools such as the chronological PPC can improve regulatory and political discussions about electricity and renewables by providing decision makers with a greater understanding of the economic and operational impacts of their decisions.

The analyses conducted in Chapters 3 and 4 only represent a few of the questions that the chronological model can explore. Building on the case study from those chapters, the following list details a few additional ideas for future research.

1. The current study of different merit order positions for hydro dispatched all hydro plants at the same position at every hour. This limit constrained the model to only test 88 dispatch positions. In reality, the optimal hydro dispatch

position may be different for each hour. A dynamic programming algorithm could explore different dispatch positions and their effects on reliability.

2. The analysis of the effects of hydro merit order dispatch position on reliability assumed a single, constant price for NSE. Varying the cost of NSE and introducing demand-response into the total cost calculation allows the model to consider other sources of generation and reliability, such as load-shedding.

3. Chapters 3 and 4 touched on the topic of maximum NSE prices and reliability targets. The chronological model can also directly explore the dynamics between price caps for bids and ENSE/LOLP, as well as the dynamics between price caps for bids *and* targets for ENSE/LOLP, by imposing price constraints at the dispatch level and then evaluating the resulting cost-reliability frontiers.

# Appendix A

# Acronyms

CDF: cumulative distribution function

ENSE: expected nonserved energy

FOR: forced outage rate

HCDF: hydro cumulative distribution function

ILDC: inverted load duration curve

LDC: load duration curve

LEP: limited energy plant

LOLP: loss-of-load probability

NSE: nonserved energy

PPC: probabilistic production cost

THIS PAGE INTENTIONALLY LEFT BLANK

# Bibliography

[1] Ignacio J Pérez Arriaga. Managing large scale penetration of intermittent renewables. In *Massachusetts Institute of Technology Energy Initiative Symposium Series*, Cambridge, MA, 2011. Massachusetts Institute of Technology.

[2] H. Baleriaux, E. Jamoulle, and Fr. Linard de Guertechin. Simulation de l'explotation d'un parc de machines thermiques de poduction d'éelectrictée couplée àa des statons de pompage. *Societée Royale Belge des Electriciens*, 7:225–245, 1967.

[3] R R Booth. Power System Simulation Model Based On Probability Analysis. *IEEE Transactions on Power Apparatus and Systems*, pages 62–69, 1972.

[4] Carmen L.T. Borges and Roberto J. Pinto. Small Hydro Power Plants Energy Availability Modeling for Generation Reliability Evaluation. *IEEE Transactions on Power Systems*, 23(3):1125–1135, August 2008.

[5] Jia-Yo Chiang, A.M. Breipohl, F.N. Lee, and R. Adapa. Estimating the variance of production cost using a stochastic load model, 2000.

[6] Antonio Conejo. *Equivalent Load Production Cost Models for Thermal Dominated Electric Energy Systems*. PhD thesis, Universidad Pontificia Comillas, 1992.

[7] S. Finger. Modeling Conventional and Pumped Hydro-Electric Energy Using Booth-Baleriaux Probabilistic Simulation. 1975.

[8] Camino González, Jesús Juan, José Mira, F.J. Francisco J Prieto, María J Sánchez, C. Gonzalez, and M.J. Sanchez. Reliability Analysis for Systems With Large Hydro Resources in a Deregulated Electric Power Market. *IEEE Transactions on Power Systems*, 20(1):90–95, February 2005.

[9] J. A. González. Probabilistic Production Costing Modeled with AMPL. *Power*, 17(2):277–282, 2002.

[10] José González and Narcís Nabona. Multicommodity long-term hydrogeneration optimization with capacity and energy constraints. *Top*, 8(1):73–96, June 2000.

[11] M.E.P. Maceira, M.V.F. Pereira, and Rio De Janeiro. Analytical modeling of chronological reservoir operation in probabilistic production costing [of hydrothermal power systems]. *IEEE Transactions on Power Systems*, 11(1):171–180, 1996.

[12] M Milligan and K Porter. Determining the Capacity Value of Wind : An Updated Survey of Methods and Implementation Preprint. *Energy*, 2008.

[13] N. Nabona, J. A. González, and J. Castro. Optimum Long-Term Hydrothermal Coordination with Fuel Limits. *IEEE Transactions on Power Systems*, 10(2):1054–1062, May 1995.

[14] Pablo Rodilla. *Regulatory Tools to Enhance Security of Supply at the Generation Level*. PhD thesis, Universidad Pontificia Comillas de Madrid, 2010.

[15] Pablo Rodilla and Carlos Batlle. Security of Electricity Supply at the Generation Level: Problem Analysis. 2010.