Noname manuscript No.
(will be inserted by the editor)

A Scalable Solution Framework for Stochastic Transmission and Generation Planning Problems

Francisco D. Munoz · Jean-Paul Watson

Received: November 24, 2015 / Accepted: date

Abstract Current commercial software tools for transmission and generation investment planning have limited stochastic modeling capabilities. Because of this limitation, electric power utilities generally rely on scenario planning heuristics to identify potentially robust and cost effective investment plans for a broad range of system, economic, and policy conditions. Several research studies have shown that stochastic models perform significantly better than deterministic or heuristic approaches, in terms of overall costs. However, there is a lack of practical solution approaches to solve such models. In this paper we propose a scalable decomposition algorithm to solve stochastic transmission and generation planning problems, respectively considering discrete and continuous decision variables for transmission and generation investments. Given stochasticity restricted to loads and wind, solar, and hydro power output, we develop a simple scenario reduction framework based on a clustering algorithm, to yield a more tractable model. The resulting stochastic optimization model is decomposed on a scenario basis and solved using a variant of the Progressive Hedging (PH) algorithm. We perform numerical experiments using a 240-bus network representation of the Western Electricity Coordinating Council in the US. Although convergence of PH to an optimal solution is not guaranteed for mixed-integer linear optimization models, we find that it is possible to obtain solutions with acceptable optimality gaps for practical applications. Our numerical simulations are performed both on a commodity workstation and on

Francisco D. Munoz
Sandia National Laboratories, Analytics Department, P.O. Box 5800, MS 1326
Albuquerque, NM 87185-1326, USA
Tel.: +1-505-284 3787
E-mail: fdmunoz@sandia.gov

Jean-Paul Watson
Sandia National Laboratories, Analytics Department, P.O. Box 5800, MS 1326
Albuquerque, NM 87185-1318, USA
E-mail: jwatson@sandia.gov

a high-performance cluster. The results indicate that large-scale problems can be solved to a high degree of accuracy in at most two hours of wall clock time.

Keywords Transmission Planning · Generation Planning · Stochastic Mixed-Integer Programming · Progressive Hedging

1 Introduction

The US electricity transmission grid is the world's largest machine. It is composed of 200,000 miles of high-voltage transmission lines that transport electricity from over 19,000 generators to 145 million customers, with annual revenues of $364 billion for the year 2012 [17]. Operations for this extremely complicated system are managed using advanced optimization techniques in both deregulated markets (e.g., California ISO, PJM, MISO, and ISO New England) and vertically integrated utilities (e.g., SPP and BPA). Advances in both mixed-integer optimization techniques and computing technologies have allowed utilities to more accurately represent the physical characteristics of power plants and obtain lower-cost operations plans, which have yielded cost savings of approximately $300 million per year [48]. New developments in stochastic unit commitment techniques are receiving increasing attention from academics, vendors, and the government, as they may further increase the economic efficiency of software for operations planning under increasing penetration of renewables [49]. In contrast, long-term planning tools – particularly for transmission investment – have received much less attention. Estimates indicate that the US electric power industry will require $2 trillion of investments in new transmission and generation infrastructure by 2030 in order to meet forecasted demand and policy objectives [14]. Nearly half of these anticipated investments correspond to upgrades of aging transmission and distribution systems.[1]

Optimal transmission investment planning is an extremely challenging problem that has been extensively studied [35]. Although some commercial tools are capable of optimizing transmission system topology in order to minimize system costs (e.g., NetPlan [1]), most transmission planning studies employ heuristic decision rules. One standard approach is to use production cost simulation packages (e.g., PROMOD [3] and PLEXOS [2]) to assess the performance of exogenously created – typically via domain experts – transmission and generation investment plans. Unfortunately, these simulation tools do not recommend potentially better investment alternatives than those being tested. Therefore, they can only be considered as screening tools [34, 47].

Beyond the complex combinatorial structure of optimal transmission and generation investment planning problems, factors such as increased environmental concerns, rapidly changing economic conditions (e.g., due to shale gas

[1] The referenced report [14] estimates that $298B in transmission and $582B in distribution investments are needed by 2030.

extraction), and the increased penetration of variable generation from renewables introduce significant new risks for electric utilities. Planning under uncertainty requires identifying a single investment portfolio that would minimize a certain metric (e.g., expected cost or Conditional Value-at-Risk) when tested against a broad range of scenarios – which is significantly more complicated than planning for a single future. The price of disregarding uncertainty, e.g., by only planning for the most likely scenario, includes stranded or inefficiently sited transmission and generation assets, or higher-than-necessary future investment and operations costs.

In practice, planners account for uncertainty by performing scenario analysis, which involves finding an optimal (or near optimal) investment plan for each of a set of scenarios independently. These plans are combined using heuristic decision rules that result in a single, implementable investment recommendation.[2] However, investment plans that result from the application of these heuristics can perform significantly worse (e.g., in terms of high system costs and stranded assets) than truly stochastic approaches that seek to minimize the expected system cost [42]. Unfortunately, most stochastic-programming based implementations are only demonstrational due to their high computational complexity, and as a result have only been applied to small test systems.

In this paper, we develop a stochastic mixed-integer transmission and generation investment planning model that explicitly accounts for the variability of time-dependent parameters such as load and wind, solar, and hydro power. We propose a practical scenario reduction framework to select representative scenarios and an application of the Progressive Hedging algorithm to solve the resulting reduced stochastic optimization model. We compute a lower bound on the expected optimal system cost by solving a linear relaxation of the stochastic mixed-integer linear planning problem. An upper bound is then calculated by testing each trial investment plan against a large-scale simulation of the operations problem considering the full (sampled) distribution of stochastic parameters. Although convergence of the Progressive Hedging algorithm to a global optimum is not guaranteed in the mixed-integer case, we are able to find solutions with a 2.55% optimality gap for a 240-bus representation of the Western Electricity Coordinating Council (WECC) in the US.

The rest of this paper is structured as follows. In Section 2 we review the state-of-the-art in transmission and generation planning, together with algorithms to solve large-scale investment planning problems. In Section 3 we introduce a stochastic investment planning model that co-optimizes transmission and generation infrastructure investment decisions. In Section 4, we propose a scenario reduction framework based on the k-means clustering algorithm, to facilitate more tractable computation. Section 5 describes our implementation of the Progressive Hedging decomposition algorithm to solve our stochastic model, and discusses as a relaxation of the soft constraints that link variables across scenarios. Section 6 describes the WECC 240-bus test case and our as-

[2] A common rule is to recommend investments in transmission lines that are part of the optimal investment plan for most or all of the scenario-specific solutions [42].

sumptions regarding transmission and generation build limits. In Section 7 we present results from different numerical experiments on the WECC-240 case. Finally, we conclude in Section 8 with a summary of our contributions.

2 Literature Review

Some of the earliest work on transmission investment planning was based on linear programming [22, 63]. Advances in branch-and-bound algorithms for mixed-integer linear programming allowed researchers to later model transmission investment options using discrete variables [6, 51, 58]. This represents an improvement over linear programming models because a) the modular or "lumpy" characteristic of transmission investment alternatives can only be accurately modeled using binary variables and b) Kirchhoff's Voltage Law for candidate lines can be enforced using linear disjunctive instead of nonlinear constraints [24,45]. Further advances in the formulation of investment planning problems account for linear [66] and piecewise linear losses [5]. A comprehensive summary of transmission planning models is provided in [35].

Transmission planning models that explicitly account for uncertainty often rely on stochastic programming [4, 42–44, 50, 60, 66] or real options [27]. Some of the most common sources of uncertainty addressed in these models include demand growth, fuel prices, technology costs, regulatory conditions, and the quantities of time-dependent resources such as load and wind, solar, and hydro power. Although stochastic models are significantly more difficult to solve than their deterministic counterparts, modeling uncertainty explicitly in the decision process can yield significant cost savings. For example, [42] finds that by ignoring uncertainty (e.g., by using a standard deterministic planning model) one can increase total system cost relative to a stochastic model by more than the cost of the transmission lines themselves. However, due to computational limitations, solving such models directly (specifically, via direction solution of the extensive form or deterministic equivalent) for large-scale power systems can be extremely slow or even infeasible due to time and/or memory limits.

An alternative to directly solving a stochastic program is the application of decomposition algorithms. One of the first and most widely used decomposition strategies is based on Benders decomposition [7]. The specialization of this approach for stochastic programs is known as the L-shaped method [10]. The L-shaped algorithm separates the investment problem (i.e., the master problem) from the operations problems (i.e., subproblems), which is particularly advantageous for models that account for discrete investment alternatives through the use of binary decision variables. The quality of the trial investment solutions that are found in the master problem is progressively improved through the addition of Benders' cuts that are computed from the subproblems. Asymptotic convergence is guaranteed for subproblems that are convex with respect to the investment variables [23]. Examples of implementations of Benders decomposition to solve deterministic transmission planning problems include [9,24,56]. Stochastic applications include models that co-optimize transmission and generation investments simultaneously [15, 44] and formulations that focus solely on generation expansion planning [11, 12, 31, 33].

Other decomposition approaches subdivide stochastic problems on a scenario basis, i.e., stochastic investment programs are divided into a collection of linked deterministic investment problems. Lagrangian relaxation [13] and Progressive Hedging [55] have been utilized to solve large-scale stochastic unit commitment [49, 57] and generation expansion planning [33] problems. To the best of our knowledge, only Progressive Hedging has been used to solve stochastic transmission planning problems [54]. However, this application was restricted to a 14-bus test case.

Researchers have also proposed heuristic algorithms to find high-quality solution to large-scale transmission investment problems. These include greedy randomized methods [8], heuristic solutions to the investment master problem in Benders decomposition [46], search algorithms based on sensitivity analysis [36], genetic algorithms [59], and sequential approximation approaches to account for wind and load variability [50]. Although some of these heuristics can obtain potentially good solutions, none of them provide a metric of solution quality, such as a bound on the optimal system cost.

3 Model Description

We now propose a stochastic investment planning model that co-optimizes transmission and generation investment portfolios. To maintain linearity of the operations model, we consider *independent* hourly economic dispatch models and ignore all ramping and unit commitment constraints, i.e., constraints that link operational hours across time and determine unit on/off state. These relaxations could potentially yield solutions that over-estimate the flexibility of the corresponding candidate generation portfolio, and therefore yield higher production costs than those predicted by a full-fidelity operations model with commitment and dispatch. However, in practice the additional costs yielded by these relaxations represent only a small fraction of the total system cost.[3] Thus, all variables are assumed to be continuous except for those used to capture transmission investment alternatives. We also assume that generation marginal costs are constant.

We now introduce the primary notation used to define our investment planning model. Additional parameters will be introduced subsequently as needed.

Sets and Indices

[3] For instance, [32] reports that although the generation investments recommended by planning models with and without unit commitment are different, the generation fleet selected by a model based strictly on economic dispatch yields a total system cost that is only 0.02% higher than those resulting from a model based on unit commitment.

A	:	Interfaces, indexed by a;
B	:	Buses, indexed by b;
G	:	Generators, indexed by i;
G_b	:	Subset of generators at bus b;
G_E	:	Subset of existing generators;
G_C	:	Subset of candidate generators;
G_R	:	Subset of renewable generators;
H	:	Hours, indexed by h;
L	:	Transmission lines, indexed by l;
L_C	:	Subset of candidate transmission lines;
L_E	:	Subset of existing transmission lines;
T	:	Number of years in the horizon, indexed by t;

Parameters

D_{bh}	:	Load [MW];
D_{bh^*}	:	Peak load [MW];
$ELCC_i$	:	Effective load carrying capability factor [%];
F_l	:	Line limit [MW];
FG_a	:	Interface limit [%];
FOR_i	:	Forced outage rate;
GC_i	:	Capital cost of generator [$/MW];
LC_l	:	Capital cost of transmission line [$];
MC_i	:	Marginal cost of generator [$/MWh];
P_h	:	Normalized duration (probability) of dispatch scenario;
LT	:	Construction lead time [years];
M_l	:	Large positive number;[4]
POR_i	:	Planned outage rate;
RM	:	Reserve margin [%];
RPS	:	Renewable target [%];
$VOLL$	:	Value of lost load [$/MWh];
Y_i^{max}	:	Maximum resource potential [MW];
Y_i^{inst}	:	Installed capacity [MW];
δ	:	Discount factor;
γ_l	:	Line susceptance;
Ψ_{al}	:	Element of interface-line incidence matrix;
Φ_{bl}	:	Element of node-line incidence matrix;

Decision variables

[4] The selection of disjuntive parameters is extremely important to ensure a tight mixed-integer formulation. Here we follow the approach described in [42].

f_{lh}	:	Power flows [MW];
g_{ih}	:	Generation dispatch level [MW];
r_{bh}	:	Load curtailment level [MW];
x_l	:	Transmission investment decisions [binary];
y_i	:	Generation investment decisions [MW];
θ_{bh}	:	Phase angles;

The objective function in our model minimizes the sum of capital investment cost plus the present worth of expected operations cost for a 50-year horizon ($T = 50$). As in [42], we assume that all new transmission and generation infrastructure has a 10-year permitting and construction lead time ($LT = 10$).[5] The transmission system is modeled using a linear DC power flow approximation, which is a standard assumption for long-term investment planning studies.[6] This formulation could be easily extended to account for linear [66] or quadratic [5] transmission losses, but at the expense of a more complex model.

Costs in our model are formally defined as follows:

Capital cost

$$CC = \sum_{i \in G_C} GC_i y_i + \sum_{l \in L_C} LC_l x_l \tag{1}$$

Present worth of operations cost for each dispatch scenario

$$OC_h = \left(\frac{1}{1+\delta}\right)^{LT} \sum_{t \in T} \left(\frac{1}{1+\delta}\right)^{t} 8760 \left(\sum_{i \in G} MC_i \; g_{ih} + \sum_{b \in B} VOLL \; r_{bh}\right) \tag{2}$$

The objective function is then formally given as:

$$min \; CC + \sum_{h \in H} P_h OC_h \tag{3}$$

The core model constraints are given as follows:

Generation build limits

$$y_i \leq Y_i^{max} \quad \forall i \in G_C \tag{4}$$

[5] Although permitting and construction lead times for renewable generation technologies and small conventional generators could be much shorter in the US, transmission and large generation projects can take up to 10 years to complete.

[6] Modeling AC power flows in a long-term investment planning model is an ongoing subject of research.

Installed reserves The installed capacity of renewable generation technologies is derated using Effective Load Carrying Capability (ELCC) factors based on historical data.[7] For simplicity, we enforce a single constraint on available reserves in the system. However, this basic construct could be easily enforced on subsets of buses (e.g., reliability regions), as is often done in large interconnected systems. The peak-load hour is denoted h^*.

$$\sum_{i \in G \setminus G_R} (Y_i^{inst} + y_i) + \sum_{i \in G_R} ELCC_i (Y_i^{inst} + y_i) \geq (1 + RM)(\sum_{b \in B} D_{bh^*}) \quad (5)$$

Renewable target per year This constraint imposes a minimum amount of power supplied from renewable generators, defined as a fraction of the supplied demand. For power systems that span across multiple jurisdictions (e.g., the WECC or MISO), renewable targets are enforced on a state-by-state basis (e.g., via state Renewable Portfolio Standards).

$$\sum_{h \in H} \sum_{i \in G_R} P_h g_{ih} \geq RPS(\sum_{h \in H} \sum_{i \in G} P_h g_{ih}) \quad (6)$$

Kirchhoff's Current Law

$$\sum_{l \in L} \Phi_{bl} f_{lh} + \sum_{i \in G_b} g_{ih} + r_{bh} = D_{bh} \quad \forall b \in B, h \in H \quad (7)$$

Kirchhoff's Voltage Law (KVL) for existing lines

$$f_{lh} - \gamma_l \sum_{b \in B} \Phi_{bl} \theta_{bh} = 0 \quad \forall l \in L_E, h \in H \quad (8)$$

Kirchhoff's Voltage Law for candidate lines We enforce KVL for candidate lines using a disjunctive formulation that allows us to maintain linearity. To avoid weaker than necessary relaxations, we use the minimum feasible values of M_l such that the angle differences are unconstrained if a line is not selected for construction ($x_l = 0$).

$$|f_{lh} - \gamma_l \sum_{b \in B} \Phi_{bl} \theta_{bh}| \leq M_l (1 - x_l) \quad \forall l \in L_C, h \in H \quad (9)$$

[7] In reality, the capacity value of some renewable resources decreases as the amount of installed capacity of the resource increases. For instance, in a system where most of the peak-load hours occur in the late afternoon, when the probabilities of lost load are highest and when solar resources are often fully available, the first few MWs of solar will directly reduce the need of power supply from peaking generation units. In that case, solar resources should be assigned a capacity value that is close to their rated capacity (i.e., $ELCC \approx 1$). However, higher penetrations levels of solar could result in a shift of peak net load hours to the early evening, when the resource is no longer available. In this second case, the marginal contribution of an additional MW of solar capacity towards a reduction of the peak net load (and a reduction of the probability of lost load) is nearly zero (i.e., $ELCC \approx 0$) [39]. In practice, electric power utilities rely on estimates of ELCC from historical data for new power plants. In some cases, these values are updated year by year based on the actual contribution of the resources towards reductions of peak net loads [40].

Power flow limits for existing lines

$$|f_{lh}| \leq F_l \quad \forall l \in L_E, h \in H \tag{10}$$

Power flow limits for candidate lines

$$|f_{lh}| \leq F_l x_l \quad \forall l \in L_C, h \in H \tag{11}$$

Maximum generation limits We assume that all generators are dispatchable. In particular, this implies that renewables resources can be (partially) curtailed. Variability of renewable resources is modeled through coefficients $W_{ih} \in [0, 1]$ based on historical or forecasted data. Coefficients W_{ih} are set to 1 for all conventional generation technologies.

$$g_{ih} \leq W_{ih}(1 - FOR_i)(1 - POR_i)(Y_i^{inst} + y_i) \quad \forall i \in G, h \in H \tag{12}$$

Interfaces Some system operators impose limits on the sum of power flows across group of lines (also known as flowgates). We assume that the interface limits correspond to a fixed fraction of the installed transmission capacity of the lines that belong to the interface.

$$\sum_{l \in L} \Psi_{al} f_{lh} \leq FG_a(\sum_{l \in L_E} |\Psi_{al}| F_l + \sum_{l \in L_C} |\Psi_{al}| F_l x_l) \quad \forall a \in A, h \in H \tag{13}$$

Integrality and non-negativity

$$x_l \in \{0, 1\} \quad \forall i \in L_C \tag{14}$$

$$y_i, g_{ih}, r_{b,h} \geq 0 \quad \forall i \in G, b \in B, h \in H \tag{15}$$

4 Scenario Reduction Framework

The investment planning model described in the previous section can result in an extremely large mixed-integer linear optimization problem, depending on the system size and the granularity of (e.g., number of hours represented in) the operations model. A common approximation to reduce the computational complexity of such large-scale stochastic programs involves the use of load duration curves, in which a small number of representative hours are selected across an entire year (e.g., peak, shoulder, and off-peak load hours) [28]. The papers by [18] and [49] describe more sophisticated scenario selection techniques applied to long and short-term power systems planning problems, respectively. Other generic scenario reduction frameworks include importance sampling [31], latin hypercube sampling [61], and moment matching [30]. Convergence of these methods to the optimal solution considering the full distribution of scenarios is guaranteed in the limit, for large sample sizes [16]. Unfortunately, these approaches do not provide bounds on the optimal system cost of the full-resolution problem.

Here we extend the clustering-based scenario reduction method developed in [44], which can be used to both reduce the size of the stochastic investment planning problem and to derive bounds on the optimal system cost. Other researchers have applied this approach to reduce the computational complexity of large stochastic programs, but in the context of production cost models [19, 29, 67] and not investment planning.

In the remainder of this paper, we refer to representative hours h as scenarios s, making the connection to stochastic programming explicit. Note that in our planning problem all stochasticity is restricted to loads and wind, solar, and hydro power. We let the vector $\mathbf{W}(s)$ denote a realization of these parameters and denote the corresponding probability space by (P, Ω), where $s \in \Omega$. In practice, this probability space can be constructed using historical and forecasted observations of these time-dependent parameters, e.g., $N = 8,760$ observations of loads and wind, solar, and hydro power from a representative year, in which case $P_s = 1/N \ \forall s \in \Omega$. We partition the space Ω into k disjoint subsets $S_1, S_2, ..., S_k$, attempting to minimize the sum of the square differences between the elements and the mean of each subset using the k-means algorithm. We then construct the new sample space Ψ^k, composed of k events $\psi_1, \psi_2, ..., \psi_k$, and denote a realization loads and wind, solar, and hydro power by $\mathbf{W}(\psi_i) = E[\mathbf{W}(s)|S_i]$ with probability $P_{\psi_i} = P_{S_i} \ \forall i \in \{1, ..., k\}$. Thus, each vector $\mathbf{W}(\psi_i)$ represents the average conditions of load and wind, solar, and hydro power within each partition S_i. Minimization of this metric ensures that hours of similar characteristics are grouped into the same partition, e.g, low wind and low demand hours will be separated from high wind and peak load hours.

For $k = 1$ this method returns a single observation equal to the expected value of the stochastic parameters ($\mathbf{W}(\psi) = E[\mathbf{W}(s)|\Omega]$) and the model described in Section 3 is equivalent to the expected value problem [10]. For small values of k ($k > 1$) the set $\mathbf{W}(\psi_1), ..., \mathbf{W}(\psi_k)$ represents a coarse approximation of the full distribution of uncertain parameters and the investment planning problem becomes stochastic. Increasing the number of partitions improves the quality of the approximation of the stochastic parameters but at the expense of a larger and more difficult optimization model; in practice, there is an exponential relationship between the number of scenarios considered and solution time. In the limit, for $k = N$, the set of events $\psi_1, ..., \psi_N$ becomes equivalent to the original sample space Ω.

Previous numerical experiments with this method indicate that only a small number of clusters are required to capture a large percentage of the variance[8] present in the full dataset $\mathbf{W}(\omega_1), ..., \mathbf{W}(\omega_N)$. The remaining fraction of variance converges to 1 asymptotically as the cluster count is increased [29, 44, 62]. Hobbs and Ji [29] also compare the performance of different hierarchical and non-hierarchical clustering algorithms for probabilistic production cost simulation and conclude that k-means yields the best clustering efficiency.

[8] The percentage of variance captured relative to the full dataset is measured as the ratio of the between-cluster variance and the total variance.

Potential extensions to this scenario reduction framework include consideration of other sources of uncertainties such as fuel prices, capital cost of transmission and generation, renewable targets, and emissions policies [42]. As in [29], those stochastic parameters could be clustered separately from the space of load, wind, solar, and hydro observations. The main advantage of creating representative scenarios through a clustering algorithm compared to sampling approaches is the possibility of defining lower and upper bounds on the total system cost. We describe this property in the next section.

5 Scenario-Based Decomposition and Bounds Computation

In this section we describe our implementation of the Progressive Hedging decomposition algorithm to solve the stochastic MILP investment planning problem described in Section 3. We also propose a relaxation of the renewable target constraint, to make the problem separable on a per-scenario basis. In Section 5.2 we propose an approach to compute upper and lower bounds on the optimal system cost. We solve the MILP investment planning problem with clustered scenarios and find a trial investment plan using Progressive Hedging. We then compute lower bounds using a linear relaxation of the investment planning problem and upper bounds by testing the trial investment plan against an economic dispatch simulation using the full distribution of loads and wind, solar, and hydro output.

5.1 A Progressive Hedging Implementation

We solve our stochastic investment planning model using the Progressive Hedging (PH) algorithm proposed by Rockafellar and Wets [55]. Unlike Benders decomposition, which separates investments (represented in a master problem) from operations (represented in subproblems), PH decomposes the problem by scenario. Scenario-based decomposition is accomplished by introducing scenario-specific investment variables x_s and initially relaxing the non-anticipativity restrictions, i.e., the constraints that require $x_s = x_u \ \forall s, u \in \Omega$. Non-anticipativity is then restored iteratively by penalizing deviations from the average "first-stage" or investment solution (i.e., $\bar{x} = \sum_{s \in S} P_s x_s$). Thus, each sub-problem in PH corresponds to a deterministic version of the stochastic problem, with an penalty-augmented objective function – each of which can be solved independently. Although the PH decomposition framework results in several mixed-integer investment planing problems, these subproblems can be solved efficiently using off-the-shelf commercial solvers. Consequently, PH avoids the growth in size and complexity of the master problem commonly associated with Benders-based decomposition, which is often solved as a single large mixed-integer linear program [44].

In order to apply PH to the investment planning problem defined in Section 3, we must first relax any constraints that link variables across scenarios, e.g., expected-value constraints. In our case, the renewable portfolio standard (RPS) target (Constraint (6)) couples economic dispatch variables from all hours within a sample year. In practice, RPS targets are considered soft constraints. Most RPSs allow electric utilities to violate a pre-specified target by imposing a fixed cost per MWh of energy not supplied from qualifying renewable resources. Further, utilities that generate more power from renewables than the minimum required per year can sell those renewable credits to utilities in other regions, or bank them for compliance in subsequent years. Using this economic interpretation of potential mismatches to the annual RPS target, we dualize Constraint (6) in the objective function as follows:

$$min\ CC + \sum_{s\in\Omega} P_s OC_s + \lambda \left[RPS(\sum_{s\in\Omega} P_s(\sum_{i\in G} g_{is} - \sum_{i\in G_R} g_{is})) \right]. \tag{16}$$

The parameter λ represents a price ceiling on the cost of meeting the RPS target. If the price ceiling is too low, the planning model will recommend meeting only fraction of the annual target ($\sum_{s\in\Omega}\sum_{i\in G_R} g_{is}$) and as a consequence will pay a non-compliance fine for the shortfall (i.e., $RPS(\sum_{s\in\Omega} P_s(\sum_{i\in G} g_{is} - \sum_{i\in G_R} g_{is}) \geq 0$). If the price ceiling is too high, it becomes profitable to supply more power from renewables than the minimum required by the target. These additional revenues are expressed as a negative cost in the objective function of the planning problem (i.e., $RPS(\sum_{s\in\Omega} P_s(\sum_{i\in G} g_{is} - \sum_{i\in G_R} g_{is}) \leq 0$). In our experiments, we estimate the value of λ that supports compliance with the renewable target by performing a sensitivity analysis on the LP relaxation of the investment planning problem. The extensive form of the LP relaxation can be easily solved using an off-the-shelf package and no decomposition is required. However, market-based values for λ from enacted renewable policies can also be used as an input to the model.

Using a more compact notation we define the investment planning problem as follows:

$$TC((\Omega, P)) = \min_{\mathbf{x}}\ \mathbf{e}\cdot\mathbf{x} + \sum_{s\in\Omega} P_s g_s(\mathbf{x}) \tag{17}$$

$$\text{s.t.}\ A\mathbf{x} \leq \mathbf{b} \tag{18}$$

$$\mathbf{x} = (\mathbf{x_a}, \mathbf{x_b}),\ \mathbf{x_a} \in \{0,1\},\ \mathbf{x_b} \geq 0 \tag{19}$$

The vector $\mathbf{x}$ denotes the transmission (binary) and generation (continuous) investment variables and the vector $\mathbf{e}$ denotes capital costs associated with the investment alternatives. Constraint (18) is a vector form representation of generation build limits (Constraint 4) and installed reserves (Constraint 5). The function $g_s(\mathbf{x})$ denotes the operations problem for a given scenario $s \in \Omega$ and is defined as follows:

$$g_s(\mathbf{x}) = \min_{\mathbf{y_s}}\ \mathbf{c_s}\cdot\mathbf{y_s} \tag{20}$$

$$\text{s.t.}\ U\mathbf{y_s} \leq \mathbf{r_s} - T_s\mathbf{x} \tag{21}$$

$$\mathbf{y_s} \geq \mathbf{0} \tag{22}$$

The vector $\mathbf{y_s}$ represents generation dispatch levels, phase angles, power flows, and demand curtailment. Constraints (7)-(13) and (15) can be written in vector form, abstracted in Constraint (21). The vector of cost coefficients $\mathbf{c_s}$ represents marginal costs, load curtailment penalties, and the terms associated with the dualization of Constraint (6).

To describe the PH algorithm, we first introduce an iteration counter v, a set of multiplier vectors $\mathbf{w_s^v}$, and a vector $\boldsymbol{\rho}$. Pseudo-code for PH is then given as follows:

Algorithm 1 Progressive Hedging

1: Initialization:
   $v \leftarrow 0$ and $\mathbf{w_s^v} \leftarrow 0, \forall s \in \Omega$
2: Iteration 0:
   $\mathbf{x_s^v} = argmin_{\mathbf{x}} f(\mathbf{x}) + g_s(\mathbf{x})$,
   subject to constraints (18), (19), (21), and (22), $\forall s \in \Omega$
3: Aggregation:
   $\bar{\mathbf{x}}^{\mathbf{v}} = \sum_{s \in \Omega} P_s \mathbf{x_s^v}$
4: Convergence check:
   If $\sum_{s \in \Omega} P_s ||\mathbf{x_s^v} - \bar{\mathbf{x}}_{\mathbf{s}}{}^{\mathbf{v}}|| < \epsilon$, stop.
5: Iteration update:
   $v = v + 1$
6: Multiplier update:
   $\mathbf{w_s^v} \leftarrow \mathbf{w_s^{v-1}} + \boldsymbol{\rho}(\mathbf{x_s^{v-1}} - \bar{\mathbf{x}}^{\mathbf{v-1}}), \forall s \in \Omega$
7: Iteration v:
   $\mathbf{x_s^v} = argmin_{\mathbf{x}} f(\mathbf{x}) + g_s(\mathbf{x}) + \mathbf{w_s^v} \cdot \mathbf{x_s^v} + \frac{1}{2}\boldsymbol{\rho} \cdot ||\mathbf{x_s^v} - \bar{\mathbf{x}}_{\mathbf{s}}{}^{\mathbf{v-1}}||^2$,
   subject to Constraints (18), (19), (21), and (22), $\forall s \in \Omega$
8: Repeat:
   Go to Step 3.

The performance of PH in practice is strongly dependent on the value of the penalty parameter $\boldsymbol{\rho}$, used to update the multipliers $\mathbf{w_s^v}$. For example, small $\boldsymbol{\rho}$ values can significantly delay convergence, as insufficiently large changes in the multipliers $\mathbf{w_s^v}$ will fail to yield perturbations to the first stage investment variables $\mathbf{x_s^v}$. Similarly, large $\boldsymbol{\rho}$ values can cause "over-shooting" by inducing radical changes in $\mathbf{x_s^v}$ values, yielding oscillatory behavior in PH. A "cost-proportional" approach to setting $\boldsymbol{\rho}$ appears to be effective in practice [64], in which the value of a specific component of $\boldsymbol{\rho}$ is proportional to the coefficient of the corresponding variable in the objective function of the deterministic scenario optimization problem. We follow this approach and define values of $\boldsymbol{\rho}$ for transmission and generation investment variables equal to $\rho_T^l = \mu_T LC_l$ and $\rho_G^i = \mu_G GC_i$, respectively. The parameters μ_T and μ_G are scaling factors over which we perform a sensitivity analysis, as described in Section 7.3.

While PH is provably convergent in the convex case [55], the presence of discrete decision variables can lead to cyclic behavior that prevents termination of the algorithm. To detect the presence of such cycles, we analyze the history of the multiplier vectors $\mathbf{w_s^v}$ at each PH iteration, using computationally and memory-efficient techniques described in [64]. If cycles are of sufficient length (as quantified in terms of the number of PH iterations between recurrent values of a multiplier), then a variable involved in the cycle is selected arbitrarily. The cycle is then broken by fixing the value of the corresponding variable in each scenario $s \in \Omega$ to the largest observed value in any scenario. This strategy, referred to as "slamming", guarantees the feasibility of all sub-problems in so-called one-sided diet problems [21] – of which our investment planning problem is an instance. In our experiments, cycles are detected and broken on an infrequent basis. Although cycle detection and breaking (among other strategies) are effective means to obtain a convergent PH algorithm for problems with discrete decision variables, no proof of optimality is available. Consequently, PH has historically been viewed as a heuristic in the mixed-integer case [37]. However, recent advances in PH have included general methods for computing rigorous bounds at any iteration of the algorithm [20].

Another technique to accelerate PH involves the introduction of an optimality gap when solving scenario sub-problems. In particular, optimal solutions are not necessary during early PH iterations, where PH is attempting to establish reasonable approximations of the multipliers $\mathbf{w_s^v}$. Then, as PH converges, the optimality gap can be incrementally reduced.

Finally, to accelerate convergence of the PH algorithm, we employ variable fixing. Our variable fixing strategy is based on the heuristic that if a variable's value is non-anticipative for a contiguous number of PH iterations, it is unlikely to change in subsequent iterations and can therefore be fixed at its present value. Variable fixing leads to smaller scenario sub-problems, which in turn leads to faster solve times.

5.2 Computation of Lower and Upper Bounds

We now describe a methodology to compute global lower and upper bounds on the optimal system cost of our stochastic mixed-integer investment planning problem. These bounds are applicable to planning problems with stochastic right-hand-sides or stochastic cost coefficients. Although we here we only deal with the former, Jensen's lower bound is also applicable to the latter class of problems by a simple transformation.[9]

For the development of a lower bound, we first note that the operations problem defined by the expression $\sum_{s\in\Omega} P_s g_s(\mathbf{x})$ is a linear program with a stochastic right-hand-side (Constraints (21)). For a given trial investment plan $\mathbf{x}$, the standard lower bound based on Jensen's inequality can be applied [10, 44]. Therefore, given a partition $S_1, ..., S_k$ of the scenario space Ω, and a set of representative scenarios $\Psi^k = \{\psi_1, ..., \psi_k\}$ such that $\mathbf{W}(\psi_i) =$

[9] Note that the problem $min\ \sum_{s\in\Omega} c_s x_s$, subject to $x_s \geq 0$ can be re-written as $min\ \sum_{s\in\Omega} u_s$, subject to $u_s \geq c_s x_s$ and $u_s, x_s \geq 0$. Jensen's inequality is directly applicable to the latter.

$E[\mathbf{W}(s)|S_i]$, the stochastic operations problem solved using clustered parameters ($\sum_{\psi \in \Psi^k} P_\psi g_\psi(\mathbf{x})$) provides a lower bound on $\sum_{s \in \Omega} P_s g_s(\mathbf{x})$. Furthermore, if a hierarchical partitioning algorithm is used, then the bound is a monotonically increasing function of the number of clusters (i.e., $\sum_{\psi \in \Psi^k} P_\psi g_\psi(\mathbf{x}) \leq \sum_{\psi \in \Psi^{k+1}} P_\psi g_\psi(\mathbf{x})$) [10].

The lower bound can also be extended to the mixed-integer linear investment problem, as in [44]. Let the vector $\mathbf{x^*}$ denote the optimal solution to the investment planning problem $TC(\Omega)$. By applying Jensen's inequality to the operations problem, we obtain the following:

$$\mathbf{e} \cdot \mathbf{x^*} + \sum_{\psi \in \Psi^k} P_\psi g_\psi(\mathbf{x^*}) \leq \mathbf{e} \cdot \mathbf{x^*} + \sum_{s \in \Omega} P_s g_s(\mathbf{x^*}) \quad \forall k = 1, .., |\Omega|. \tag{23}$$

Let $\mathbf{x^*_k}$ denote an optimal solution to the investment planning problem $TC(\Psi^k)$. The following inequality then holds for any investment plan $\mathbf{x}$ feasible to Constraints (18) and (19):

$$\mathbf{e} \cdot \mathbf{x^*_k} + \sum_{\psi \in \Psi^k} P_\psi g_\psi(\mathbf{x^*_k}) \leq \mathbf{e} \cdot \mathbf{x} + \sum_{\psi \in \Psi^k} P_\psi g_\psi(\mathbf{x}) \quad \forall k = 1, .., |\Omega| \tag{24}$$

Combining inequalities (23) and (24), it follows that $TC(\Psi^k) \leq TC(\Omega)$ $\forall k = 1, .., |\Omega|$. Consequently, the optimal objective function value of the investment planning problem solved using clustered data ($TC(\Psi^k)$) provides a lower bound on the total system cost of the same problem considering the full distribution of stochastic parameters ($TC(\Omega)$). Convergence of $TC(\Psi^k)$ to $TC(\Omega)$ as the number of partitions is increased is discussed in [44].

Note that because PH is a heuristic in the case of mixed-integer linear problems, we cannot guarantee of tightness of the lower bound $TC(\Psi^k)$. However, we find that the linear relaxation $TC_{LP}(\Psi^k)$ empirically provides a tight lower bound for our investment planning problem. Furthermore, because a linear relaxation of our investment planning problem can be solved to optimality using a barrier or simplex algorithm, the resulting objective function value provides a valid global lower bound on $TC(\Omega)$ (i.e., $TC_{LP}(\Psi^k) < TC(\Psi^k) \leq TC_{PH}(\Psi^k)$). The bound developed in [20] can be also utilized.

The inequalities described above also allow us to develop upper bounds on $TC(\Omega)$. Let $\mathbf{x^*_k}$ denote the solution to $TC(\Psi^k)$ found using PH. From inequalities (23) and (24) it follows that $\mathbf{e} \cdot \mathbf{x^*_k} + \sum_{s \in \Omega} P_s g_s(\mathbf{x^*_k})$ is an upper bound on $TC(\Omega)$. Evaluating the expression $\mathbf{e} \cdot \mathbf{x^*_k} + \sum_{s \in \Omega} P_s g_s(\mathbf{x^*_k})$ only requires the solution of the operations problems $g_s(\mathbf{x^*_k})$ $\forall s \in \Omega$. These sub-problem are both separable and linear, and could be solved in parallel if necessary. For extremely large sample spaces ($|\Omega| >> 1$) it is also possible to approximate the expression $\sum_{s \in \Omega} P_s g_s(\mathbf{x^*_k})$ using a Monte Carlo simulation [52].

Given the results presented above, we progressively improve the solution quality of our investment planning problem as follows:

1. Initialize the iteration counter ($k = 0$),and set the initial bounds ($LB_0 = -\infty$ and $UB_0 = +\infty$).
2. Set $k = k + 1$.
3. Compute a lower bound on $TC(\Omega)$ by solving the linear relaxation of the mixed-integer investment planning problem using clustered hours. Set $LB_k = TC_{LP}(\Psi^k)$.
4. Solve the mixed-integer investment planning problem with clustered hours ($TC(\Psi^k)$) using PH. Denote the resulting trial investment plan by $\mathbf{x_k}$.
5. Compute an upper bound on $TC(\Omega)$. First, solve the operations problems $g_s(\mathbf{x_k})$ $\forall s \in \Omega$ and then compute the true system cost $\mathbf{e} \cdot \mathbf{x_k} + \sum_{s\in\Omega} P_s g_s(\mathbf{x_k})$ of the trial investment plan $\mathbf{x_k}$.
6. If $UB_{k-1} > \mathbf{e} \cdot \mathbf{x_k} + \sum_{s\in\Omega} P_s g_s(\mathbf{x_k})$, set $UB_k = \mathbf{e} \cdot \mathbf{x_k} + \sum_{s\in\Omega} P_s g_s(\mathbf{x_k})$ and $\mathbf{x^*} = \mathbf{x_k}$. Otherwise, $UB_k = UB_{k-1}$ and $\mathbf{x^*} = \mathbf{x_{k-1}}$.
7. If $(UB_k - LB_k)/UB_k < \epsilon$ stop, the best investment plan found so far is $\mathbf{x^*}$. Otherwise, go to Step 2.

As discussed in [44], increasing the number of partitions (i.e., increasing k) will reduce the optimality gap with respect to the objective function value of PH ($(UB_k - TC_{PH}(\Psi^k))/UB_k$). Although convergence of $(UB_k - LB_k)/UB_k$ to zero as k is increased is not guaranteed in our case due to the duality gap, through numerical experiments we find that the optimality gap of the investment plans found using our methodology are acceptable for practical applications.

6 Test Case: WECC 240-bus System

Our numerical experiments are performed on a 240-bus representation of the Western Electricity Coordinating Council (WECC) in the US [53]. This test case was expanded in [42] to account for transmission and generation investment alternatives. The network is composed of 448 transmission elements and 157 aggregated generators. We assume two types of transmission investment alternatives: reinforcements to existing transmission backbones and new radial interconnections to renewable hubs. As in [42], we allow for a maximum of two new 500kV circuits to reinforce existing transmission corridors and up to four new transmission lines to any identified renewable hub. For illustration purposes we enforce a 33% Renewable Portfolio Standard in all the US states that belong to the WECC, similar to the scenario of renewable targets studied in [38]. An extensive description of transmission and generation investment alternatives is provided in [41], and the full instance can be obtained by contacting the authors.

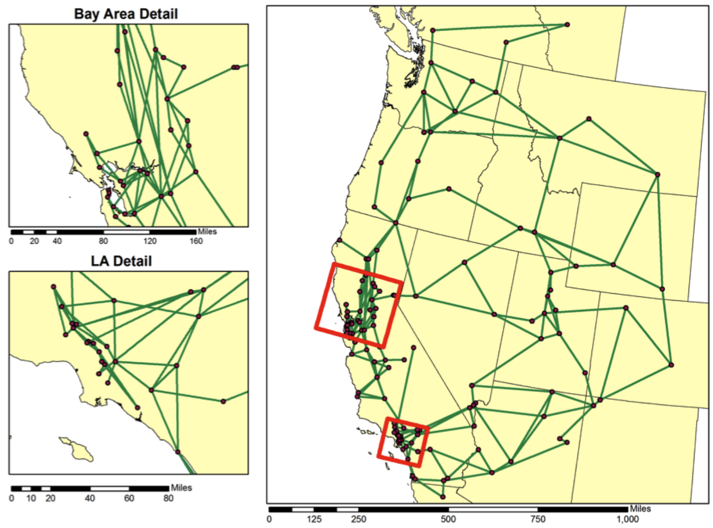

Fig. 1 Illustration of the WECC 240-bus system [42].

7 Numerical Experiments

7.1 Computational Environment

All models and algorithms were developed using the Coopr (`https://software.sandia.gov/trac/coopr`) open-source Python library for optimization, developed and maintained by Sandia National Laboratories. The core deterministic (individual scenario) optimization models are expressed in Coopr's Pyomo algebraic modeling language [25, 26]. We use Coopr's PySP library [65] to both model our stochastic investment planning problems and to solve them using PH. All data and models are available upon request. All customizations relating to $\boldsymbol{\rho}$ setting, variable fixing, cycle detection / breaking, and variable slamming are performed through a parameterization of PySP's WW extension module. All scenario solves were performed using IBM's CPLEX 12.6 mixed-integer solver.

Experimental trials were executed on a dedicated commodity 48-core Linux workstation with 512GB of RAM. The individual processors are dual quad-core, hyper-threaded Intel Xeon microprocessors, with a clock speed of 2.3GHz. Large experiments were executed on the Red Sky/Red Mesa high-performance computer, which consists of 1,920 nodes containing two quad-core 2.93 GHz Nehalem X5570 processors apiece. We used Coopr version 3.5, running Python 2.7.3.

Considering 8,736 scenarios in the planning problem results in a mixed-integer linear program with 15M variables (370 binaries for transmission) and

35M constraints. An attempt to solve the extensive form of the 100-scenario problem yielded no feasible solution after one day of wall clock time on the 48-core workstation.

7.2 Selection of the Price Ceiling of Renewable Energy Certificates (λ)

We select the value of the price ceiling of Renewable Energy Certificates (λ) to enforce the renewable target constraint by performing a sensitivity analysis on a linear relaxation of the mixed-integer planning problem. To ensure that the variability of time-dependent resources is at least partially represented, we use 50 representative scenarios obtained using the clustering methodology described in Section 4. We enforce a 33% renewable target across all regions that belong to the WECC, which is analogous to one of the regulatory scenarios analyzed in [42].

Table 1 Sensitivity analysis of renewable supply as a function of λ.

λ [$/MWh]	0	10	20	30	40	50	60	70	80	90	100
Renewables [%]	7.5	10.3	14.8	20.1	25.1	28.7	31.7	33.7	37.5	37.9	38.2

As shown in Table 1, not enforcing a renewable target ($\lambda = 0$) would only result in a 7.5% of demand being supplied from renewables. Such investments occur for purely economic reasons and do not require additional revenue streams from the production of RECs to be cost-effective. Raising the price ceiling λ naturally yields a monotonic increase in the penetration of renewables. We find that a price ceiling of $70 per MWh results in a 33.7% penetration of renewables (Table 1), which meets the target. All of the remaining numerical experiments are run assuming $\lambda = 70$ [$/MWh]. The true amount of renewable supply per year is computed in the next section using a 52-week economic dispatch.

7.3 PH Configuration

For illustration purposes we limit the number of threads assigned to each PH subproblem to one. For larger problems, increasing the number of threads assigned to each subproblem could potentially reduce solution times or improve solution quality for mixed-integer linear programs that require solution of several nodes of the MILP branch and cut tree. In our case, the relaxation induced neighborhood search (RINS) heuristic[10] provided by CPLEX yields solutions within 1-2% of optimality at the root node. Therefore, the efficiency gains due to parallelization of the MILP branch and cut process are rather modest.

[10] We invoke the RINS heuristic every 100 nodes in the MILP branch and cut tree.

Additional options passed to the CPLEX solver include aggressive scaling and numerical emphasis.

Within PH we set the MILP gap tolerance for sub-problem solves at iterations 0 and 1 to 1%. This MILP gap is then progressively tightened to 0.001% as the first stage variables (investments) converge. PH is terminated once the convergence metric value is equal or below 0.0001.

As mentioned in Section 5, the selection of the $\boldsymbol{\rho}$ parameter, the delay used to fix variables that have converged, and the MILP optimality gap are all features that influence the convergence of PH. Here, we perform a sensitivity analysis with respect to subsets of these parameters based on our previous experience solving stochastic optimization problems using PH. As in Section 7.2, we use a 50-cluster investment planning problem to test different configurations of the PH algorithm. The optimal objective function value of the LP relaxation is \$551.7B and provides a lower bound that we use to measure the quality of the investment plan from the mixed-integer linear problem found using PH.

Table 2 shows the results of a sensitivity analysis with respect to different values of the $\boldsymbol{\rho}$ parameter. In the table, the columns "F.S. Cost" and "Sol. Time" respectively denote the total investment cost and the wall clock time required for PH to converge. We find that a common default value of the parameter ($\boldsymbol{\rho} = 1$ for all variables) forces a too-rapid convergence of the first stage variables, which yields an investment plan with a 23.8% optimality gap with respect to the LP relaxation. Defining variable-specific values of ρ in proportion to their objective function cost coefficients immediately results in an improved solution. However, the cost-proportional scaling parameter μ_G on generation investment variables ultimately makes the largest difference in solution quality.

In our formulation, transmission and generation infrastructure are mutually substitutable. Lack of generation capacity to meet demand in one region can be compensated by additional transmission infrastructure that allows for imports from a neighboring region and vice versa. In the WECC test case the capital costs per MW for transmission range from \$400.8 (an extremely short line) to \$1.8M (mean=\$0.2M and standard deviation=\$0.3M) and from \$0.7M to \$5.6M for generation (mean=\$2.9M and standard deviation=\$1.7M).[11] Ignoring geographical constraints, the addition of a new MW of generation costs between 1,000 to 10,000 times more than adding a MW of transmission capacity (the cheapest transmission alternative with respect to cheapest and average generation cost, respectively). Consequently, using cost proportional values of $\boldsymbol{\rho}$ directly in PH imposes a disproportionately large pressure on the convergence of generation investment variables. As shown in Table 2, scaling the $\boldsymbol{\rho}$ parameters associated with transmission improves the solution quality significantly and reduces first stage investment costs. We find that the values of μ_G that yield the lowest total costs (0.001 and 0.0001) are inversely propor-

[11] Histograms of generation and transmission capital costs per MW are shown in Figures 3 and 4 in the Appendix.

Table 2 Sensitivity analysis of PH with respect to different configurations of $\boldsymbol{\rho}$. We impose a time limit for each PH subproblem solve equal to 60 seconds. Binary and continuous variables are fixed if values are non-anticipative after 6 or 10 contiguous iterations, respectively.

Rho configuration	Total Cost [$B]	Gap LP [%]	F.S. Cost [$B]	Sol. Time [s]	N. Iter.
$\boldsymbol{\rho = 1}$	723.8	23.8	496.0	337.5	151
Cost proportional $\boldsymbol{\rho}$, $\mu_T = 1$, and $\mu_G = 1$	705.6	21.8	487.2	2674.8	60
Cost proportional $\boldsymbol{\rho}$, $\mu_T = 1$, and $\mu_G = 0.1$	699.5	21.1	484.6	2232.1	58
Cost proportional $\boldsymbol{\rho}$, $\mu_T = 1$, and $\mu_G = 0.01$	651.2	15.3	450.5	2480.7	63
Cost proportional $\boldsymbol{\rho}$, $\mu_T = 1$, and $\mu_G = 0.001$	562.2	1.9	360.1	2041.0	122
Cost proportional $\boldsymbol{\rho}$, $\mu_T = 1$, and $\mu_G = 0.0001$	560.8	1.6	351.2	2661.8	152
Cost proportional $\boldsymbol{\rho}$, $\mu_T = 1$, and $\mu_G = 0.00001$	565.5	2.4	346.1	2231.8	174

tional to the ratios between the cheapest and average generation alternatives and the cheapest transmission investment alternative per MW. Decreasing μ_G below those values yields even lower investment costs, but results in higher total cost. Based on the results of this sensitivity analysis, all of our remaining experiments are conducted using cost proportional values of $\boldsymbol{\rho}$ with scaling factors $\mu_T = 1$ and $\mu_G = 0.0001$.

7.4 Data clustering and computation of lower bounds

We construct a sample space of load and wind, solar, and hydro output using 151 profiles and 8,736 observations from historical data [42]. Clustering of this data was performed using the Python SciPy package. Because loads and renewables data are supplied in different units, we normalize all profiles by subtracting the parameter mean and dividing by the standard deviation, yielding uniform weights for all parameters when executing k-means. User preferences can be supplied, in order to modify weights in situations where errors in some parameters (e.g., loads) are more important than in others (e.g., solar). However, implementation of a weighted k-means algorithm is beyond the scope of our research.

Table 3 shows the fraction of variance present in the full dataset captured using cluster counts ranging from 1 to 1000. The table also shows the number of variables and constraints, total cost, and solution time of the LP relaxation of the planning problem formulated using the clustered data. Note that only 100 representative clusters are needed to capture 49.2% of the variance of all the stochastic parameters. Refining the number of partitions to 1000 increases the fraction of variance captured to 80.6%. As shown in Figure 2, the asymptotic convergence of the fraction of variance captured to 100% (depicted by the blue line) – a standard behavior of clustering algorithms – is also mirrored in the improvement of the lower bound (depicted by the red line) as the cluster count is increased. Thus, even though the size of the problem and solution times scale linearly as a function of the number of clusters (see Table 3), the

Table 3 Fraction of variance captured relative to the full dataset, number of variables and constraints, total cost, and solution times of the LP relaxation as a function of the number of clusters.

Number of Clusters	Fraction of Variance [%]	Number of Variables [10^3]	Number of Constraints [10^3]	Total Cost [$B]	Solution Time [s]
1	0.0	2.6	4.5	530.1	0.3
10	32.2	18.8	40.1	548.2	5.7
20	36.8	36.9	79.6	549.1	9.9
30	39.6	54.9	119.2	550.3	16.2
40	41.7	72.9	158.7	550.4	17.9
50	43.2	90.9	198.3	551.7	19.9
100	49.2	181.1	396.0	555.4	38.2
200	56.2	361.4	791.5	557.8	75.3
300	61.7	541.7	1187.0	559.7	208.4
400	66.0	722.0	1582.5	561.3	192.2
500	69.7	902.3	1978.0	560.8	314.5
1000	80.6	1803.8	3955.5	562.3	731.9

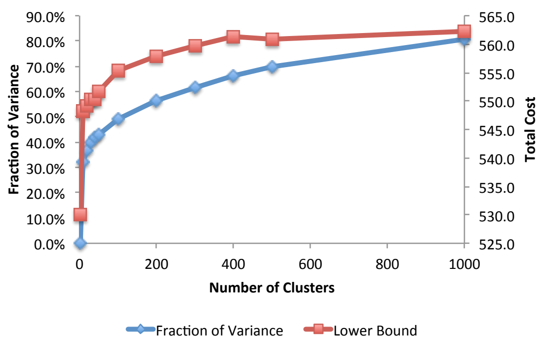

Fig. 2 Fraction of variance captured relative to the full dataset and objective function value of the LP relaxation of the investment-planning problem as a function of the number of clusters.

computational effort required to improve the lower bound increases exponentially (i.e., improvement in lower bound or fraction of variance captured versus solution time).

7.5 Solution of the Stochastic MILP and Computation of Upper Bounds

For the computation of (primal) investment plans and upper bounds on total system cost we follow the procedure outlined in Section 5.2. Table 4 summarizes results for experiments considering between 1 and 500 representative scenarios created using our clustering-based scenario reduction framework. We compute the GAP_{LB} (8th column) as the percentage difference between the objective function value of the MILP investment problem solved using PH

(2nd column) and the linear relaxation of the 1000-scenario problem.[12] This gap provides a representative metric of the relative quality of the solution found using PH with respect to the true optimal solution of the MILP problem (which cannot be computed exactly in tractable run times).

The GAP_{UB} quantity measures the relative error associated with using a reduced number of clustered hours in the MILP problem with respect to the upper bound (UB). We compute the upper bound (7th column) by adding the first stage cost from the MILP problem (investment costs, 3rd column) and the operating costs that result from testing the performance of the trial investment plan against a 52-week economic dispatch simulation (6th column). The GAP quantity is then the optimality gap of the upper bound (UB) with respect to the objective function value of the linear relaxation of the 1000-scenario problem.

Although we cannot guarantee convergence of PH to the optimal solution of the full MILP investment planning problem considering all scenarios, we find that the total system costs increase monotonically as we increase the number of scenarios. Improving the resolution of the operations problem in the MILP planning problem yields lower investment costs (first stage costs) and reductions in operating costs. Both effects combined are reflected in a reduction of the total optimality gap (GAP). The single-cluster experiment ($k = 1$) corresponds to the expected-value problem and its optimality gap provides an upper bound on the value of the stochastic solution (VSS) [10,44]. For our set of scenarios, an upper bound on VSS is 20.03% of the upper bound (\$662.9B) or \$132.8B. This quantity reflects the potential cost savings that would result from refining the partitioning of the space of loads and wind, solar, and hydro output observations.

Increasing the number of clusters from 1 to 50 reduces the optimality gap from 20.03% to 5.36%, but at the expense of an increase in solution time of the MILP problem from 15.1 seconds to 596.0 seconds. Further refinements of the number of partitions yield even tighter optimality gaps, but improvements beyond 50 scenarios (clusters) yields low marginal improvements. Increasing the number of scenarios to 500 reduces GAP_{UB} to 0.02%, which indicates that the difference between the predicted operating costs using clustered hours and the 52-week economic dispatch simulation is negligible. Note that the solution time for the 500-scenario problem is only 1.9 hours, which represents a significant reduction in solution time relative to our attempt to solve the extensive form of the smaller 100-scenario problem in 24 hours. As an additional experiment we also solved the larger problems ($k \geq 100$) on our 48-core workstation. We find that solution times modestly increase, to 0.5 hours for the 100-scenario experiment and 8.7 hours for the 500-scenario instance. Thus, our solution framework can be used to obtain high-quality solutions to large-scale planning problems on either high-performance computers or a shared-memory commodity workstations.

[12] Note that the lower bound obtained using the linear relaxation is global and it does not depend on the number of scenarios used to find a trial investment plan. Therefore, the more scenarios of clustered hours considered in the linear relaxation, the tighter the lower bound.

Table 4 Summary of PH runs and computation of upper bounds using the 52-week economic dispatch. All experiments for $k \geq 100$ were run on the Red Mesa high-performance computer. TC and OC denote Total and Operating Costs, respectively. ED stands for economic dispatch.

	MILP Investment Problem (PH)				52-week ED				
k	TC [\$B]	FS Cost [\$B]	Sol. Time [s]	Its.	OC [\$B]	UB [\$B]	GAP_{LB} [%]	GAP_{UB} [%]	GAP [%]
1	542.9	386.2	15.1	1	276.7	662.9	2.36	18.10	20.03
5	554.8	357.9	177.1	89	237.8	595.7	2.11	6.86	8.82
10	557.7	356.5	241.0	100	233.9	590.4	1.70	5.54	7.15
20	558.0	353.4	366.3	134	232.3	585.7	1.59	4.73	6.25
30	559.6	351.9	467.2	136	233.4	585.3	1.67	4.39	5.98
40	560.0	352.4	524.8	148	231.7	584.1	1.72	4.12	5.77
50	561.0	350.5	596.0	136	232.4	582.9	1.66	3.76	5.36
100	565.7	351.7	940.3	186	231.0	582.7	1.82	2.92	4.69
200	568.3	349.5	1500.4	242	231.1	580.6	1.85	2.12	3.93
300	572.0	347.3	2360.8	326	231.5	578.8	2.16	1.18	3.31
400	573.7	349.2	3721.2	414	230.0	579.2	2.16	0.95	3.09
500	575.4	345.1	6933.6	700	230.4	575.5	2.53	0.02	2.55

Next, we examine the impact of our dualization scheme for the RPS constraint. We find that our relaxation of the renewable target yields renewable penetration levels that range from 31.0% to 31.9%, with an average of 31.3% for all experiments listed in Table 4. This implies that on average, only 1.7% of the 33.0% target is met using Renewable Energy Certificates capped at \$70 per MWh. Achieving a smaller gap between the actual and target would require minor increases to this price ceiling and re-execution of our experiments; however, such analysis is beyond the present scope.

Finally, we analyze the duality gap for 100 different single-scenario investment planning problems, solved using CPLEX. We impose a 0.5% optimality gap and set a time limit of 1 hour. We find that the average gap with respect to the best incumbent solution from the MILP problems (upper bound) is 2.17% and with respect to the best lower bound is 1.39%. These results lead us to hypothesize that the quality of the solution found using PH is in practice far better than what is implied via the GAP_{LB} quantities displayed in Table 4. For instance, if these results held for multi-scenario problems, and the linear relaxation gap was 1.5% for the 500-scenario problem, the total optimality gap would be 0.81% instead of 2.55%. Therefore, the development of a tighter lower bound could directly improve our methodology. Nevertheless, we believe that a solution with an optimality gap of 2.55% provides invaluable insights to decision makers that rely on these tools for long-term investment planning.

7.6 Analysis of Primal Solutions

Finally, we study the nature of changes in the primal investment solutions as we increase the number of scenarios (clusters) in the MILP planning problem. In Table 5, we show investments as a function of the number of clusters for

Table 5 Sample of aggregate transmission and generation capacity investments per state as a function of the number of scenarios. All percentages are calculated with respect to the 500-scenario results. Total investments account for all areas in the WECC. Results in the last row are total capacity investments for the 500-scenario problem in GW.

	Transmission					Generation				
k	CA	MT	WA	WY	Total	CA CCGT	MT CT	WA CT	WY Wind	Total
	[%]	[%]	[%]	[%]	[%]	[%]	[%]	[%]	[%]	[%]
1	66.1	49.1	42.3	498.6	95.3	78.3	285.5	0.0	323.4	106.9
5	66.1	67.5	59.4	221.2	91.2	85.9	187.9	66.7	128.5	102.5
10	66.1	67.5	74.3	121.2	91.8	83.2	78.6	109.4	112.7	101.9
20	66.1	67.5	93.9	148.6	98.6	86.8	96.7	103.5	111.2	101.8
30	100.0	67.5	74.3	148.6	96.8	87.1	124.0	117.4	109.8	101.4
40	66.1	67.5	93.9	148.6	98.6	95.4	120.8	107.7	108.4	101.5
50	66.1	67.5	93.9	148.6	98.6	93.5	129.3	118.0	104.5	101.2
100	100.0	85.1	93.9	100.0	98.1	94.3	111.2	97.4	104.4	101.5
200	100.0	85.1	93.9	100.0	98.1	98.2	101.3	98.0	102.4	101.0
300	100.0	85.1	93.9	100.0	98.1	98.0	112.9	100.9	99.6	100.5
400	100.0	100.0	108.9	100.0	102.8	99.7	107.0	99.8	100.4	100.7
500	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
[GW]	9.5	5.4	13.3	0.7	84.2	13.9	7.0	9.8	4.7	195.8

a sample of four states in the WECC. We find that the expected-value problem ($k = 1$) underestimates transmission capacity by 4.7% and overestimates the need for generation by 6.9% with respect to the 500-scenario solution. Although the discrepancies are not large in magnitude at an aggregate, system-wide level, using a small number of representative scenarios can significantly bias the distribution of investments at the state level.

For instance, the expected-value problem overestimates the need for transmission infrastructure by 398.6% and wind capacity by 223.4% in Wyoming. It also underestimates the need for new transmission by 57.7% and CT generation by 100% in Washington. Increasing the number of scenarios in the operations problem embedded in the MILP planning problem reduces the error with respect to the 500-scenario problem. However, we find that at least 400 scenarios are needed to obtain investment plans with an error below 10%.[13] Although these results are not general, they highlight the importance of considering large samples of hourly observations of loads and wind, solar, and hydro output to capture the true economic value of these resources.

8 Conclusions

In this article we propose a practical and scalable solution framework for stochastic transmission and generation investment planning problems. We develop a simple stochastic investment planning problem assuming that uncer-

[13] We are implicitly assuming that the 500-scenario solutions are extremely close to the true optimal. We justify this assumption by the small optimality gap attained with that experiment.

tainty is restricted to load and wind, solar, and hydro power. To reduce the computational complexity of the problem, we utilize the k-means clustering algorithm as a scenario selection method. This algorithm automatically groups observations into clusters with similar characteristics (e.g., peak-load and low-wind or low-load and high-wind conditions) and returns one representative scenario from each cluster, the centroid. We solve the resulting stochastic mixed-integer linear program using Progressive Hedging, which decomposes the problem on a scenario basis.

Because all stochastic parameters of the planning model are right-hand-side components of constraints, we apply Jensen's inequality to compute a lower bound on the optimal system cost. Application of this lower bound requires dualizing the renewable target constraint that links variables across scenarios. This approximation yields policies that enforce renewable targets set price ceilings on the value of violating the renewable target constraint. Therefore, mismatches on the quantities of renewable power supplied per year with respect to the target have an economic interpretation and can be priced in the objective function. We approximate the value of the Lagrangian multiplier needed to attain the renewable target by performing a sensitivity analysis on the linear relaxation of the planning model.

For each trial investment plan found using PH, we compute an upper bound by testing its performance against a 52-week economic dispatch problem that can be solved in parallel. Although we cannot guarantee convergence of PH to the optimal solution of the mixed-integer investment problem, our numerical experiments suggest that careful configuration of the algorithm can yield high-quality investment plans. We find that only 50 scenarios are needed to attain optimality gaps of 5.36% and refining the partitioning of the probability space to include up to 500 representative scenarios can reduce the gap to 2.55%. Solution times for the 500 scenario instance range from 1.9 hours on the Red Mesa high-performance computer to 8.7 hours on a 48-core commodity workstation. In contrast, an attempt to solve the extensive form of the 100-scenario problem resulted in no feasible solution after more than a day of wall clock time.

As mentioned in Section 4, here we only consider uncertainty associated with time-dependent parameters. A potential extension of our methodology is to account for long-term market and regulatory uncertainties, as in [42]. Such modeling would involve consideration of stochastic parameters that are not just on the right-hand-side of constraints, but also in the objective function. Through a simple transformation described in Section 4, one could still apply Jensen's inequality to compute a lower bound on the optimal system cost, but the bound would only be valid for continuous investment variables. Another important future research direction is the application of our methodology to solve multi-stage investment planning problems. All our results based on Jensen's inequality and PH are also applicable to multi-stage formulations.

Acknowledgements The research in this article was supported by the Department of Energy's Office of Advanced Scientific Computing Research. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

References

1. PSR NETPLAN Software (2012). URL `http://www.psr-inc.com.br/portal/psr/servicos/modelos_de_apoio_a_decisao/studio_plan/netplan/`
2. Energy Exemplar PLEXOS Software (2014). URL `http://energyexemplar.com/software/plexos-desktop-edition/`
3. VENTYX PROMOD IV Software (2014). URL `http://www.ventyx.com/en/enterprise/business-operations/business-products/promod-iv`
4. Akbari, T., Rahimikian, A., Kazemi, A.: A Multi-Stage Stochastic Transmission Expansion Planning Method. Energy Conversion and Management **52**(8-9), 2844–2853 (2011)
5. Alguacil, N., Motto, A., Conejo, A.: Transmission expansion planning: a mixed-integer LP approach. IEEE Transactions on Power Systems **18**(3), 1070–1077 (2003)
6. Bahiense, L., Oliveira, G., Pereira, M., Granville, S.: A Mixed Integer Disjunctive Model for Transmission Network Expansion. IEEE Transactions on Power Systems **16**(3), 560–565 (2001)
7. Benders, J.: Partitioning Procedures for Solving Mixed-Variables Programming Problems. Numerische Mathematik **4**(1), 238–252 (1962)
8. Binato, S., De Oliveira, G.C., De Araújo, J.L.: A greedy randomized adaptive search procedure for transmission expansion planning. IEEE Transactions on Power Systems **16**(2), 247–253 (2001)
9. Binato, S., Pereira, M., Granville, S.: A New Benders Decomposition Approach to Solve Power Transmission Network Design Problems. IEEE Transactions on Power Systems **16**(2), 235–240 (2001)
10. Birge, J., Louveaux, F.: Introduction to Stochastic Programming. New York, NY:Springer (1997)
11. Bloom, J.A.: Solving an electricity generating capacity expansion planning problem by generalized benders' decomposition. Operations Research **31**(1), 84–100 (1983)
12. Bloom, J.A., Caramanis, M., Charny, L.: Long-range generation planning using generalized benders' decomposition: Implementation and experience. Operations Research **32**(2), 290–313 (1984)
13. Castillo, E., Minguez, R., Conejo, A., Garcia-Bertrand, R.: Decomposition techniques in mathematical programming. Springer (2006)
14. Chupka, M.W., Earle, R., Fox-Penner, P., Hledik, R.: Transforming Americas Power Industry. The Brattle Group (2008)
15. Dantzig, G., Glynn, P.W., Avriel, M., Stone, J., Entriken, R., Nakayama, M.: Decomposition Techniques for Multi-area Generation and Transmission Planning Under Uncertainty: Final Report. Electric Power Research Institute (1989)
16. Dupačová, J., Gröwe-Kuska, N., Römisch, W.: Scenario reduction in stochastic programming. Mathematical programming **95**(3), 493–511 (2003)
17. EIA: Energy Information Administration, Electricity Power Annual. Retrieved August 30, 2014, from http://www.eia.gov/electricity/annual/html/epa_01_01.html (2014)
18. Feng, Y., Ryan, S.M.: Scenario construction and reduction applied to stochastic power generation expansion planning. Computers & Operations Research **40**(1), 9 – 23 (2013)
19. Feng, Y., Ryan, S.M.: Solution sensitivity-based scenario reduction for stochastic unit commitment. Computational Management Science pp. 1–34 (2014)
20. Gade, D., Hackebeil, G., Ryan, S., Watson, J., Wets, R., Woodruff, D.: Obtaining lower bounds from the progressive hedging algorithm for stochastic mixed-integer programs. Under Review (2014)
21. Garille, S., Gass, S.: Stigler's diet problem revisited. Operations Research **49** (2001)
22. Garver, L.L.: Transmission network estimation using linear programming. IEEE Transactions on Power Apparatus and Systems (7), 1688–1697 (1970)

23. Geoffrion, A.M.: Generalized benders decomposition. Journal of optimization theory and applications **10**(4), 237–260 (1972)
24. Granville, S., Pereira, M.V.F.: Analysis of the Linearized Power Flow Model in Benders Decomposition. System Optimization Lab., Dept. of Operations Research, Stanford University (1985). Tech. Rep. SOL 85-04
25. Hart, W., Watson, J., Woodruff, D.: Pyomo: Modeling and solving mathematical programs in python. Mathematical Programming Compuation **3**(3) (2011)
26. Hart, W.E., Laird, C., Watson, J.P., Woodruff, D.L.: Pyomo: Optimization Modeling in Python. Springer (2012)
27. Hedman, K., Gao, F., Sheble, G.: Overview of transmission expansion planning using real options analysis. In: Proceedings of the 37th Annual North American Power Symposium, 2005., pp. 497–502. IEEE (2005)
28. Hobbs, B.F.: Optimization Methods for Electric Utility Resource Planning. European Journal of Operational Research **83**(1), 1–20 (1995)
29. Hobbs, B.F., Ji, Y.D.: Stochastic Programming-Based Bounding of Expected Production Costs for Multiarea Electric Power Systems. Operations Research **47**(6), 836–848 (1999)
30. Høyland, K., Kaut, M., Wallace, S.W.: A heuristic for moment-matching scenario generation. Computational optimization and applications **24**(2-3), 169–185 (2003)
31. Infanger, G.: Monte carlo (importance) sampling within a benders decomposition algorithm for stochastic linear programs. Annals of Operations Research **39**(1), 69–95 (1992)
32. Jin, S., Botterud, A., Ryan, S.M.: Temporal versus stochastic granularity in thermal generation capacity planning with wind power. IEEE Transactions on Power Systems **PP**(99), 1–9 (2014)
33. Jin, S., Ryan, S.M., Watson, J.P., Woodruff, D.L.: Modeling and solving a large-scale generation expansion planning problem under uncertainty. Energy Systems **2**(3-4), 209–242 (2011)
34. Kahn, E.: Wind Integration Studies: Optimization vs. Simulation. The Electricity Journal **23**(9), 51–64 (2010)
35. Latorre, G., Cruz, R.D., Areiza, J.M., Villegas, A.: Classification of Publications and Models on Transmission Expansion Planning. IEEE Transactions on Power Systems **18**(2), 938–946 (2003)
36. Latorre-Bayona, G., Perez-Arriaga, I.J.: Chopin, a heuristic model for long term transmission expansion planning. IEEE Transactions o nPower Systems **9**(4), 1886–1894 (1994)
37. Løkketangen, A., Woodruff, D.L.: Progressive hedging and tabu search applied to mixed integer (0,1) multistage stochastic programming. Journal of Heuristics **2**, 111–128 (1996)
38. Mills, A., Phadke, A., Wiser, R.: Exploration of Resource and Transmission Expansion Decisions in the Western Renewable Energy Zone initiative. Energy Policy **39**(3), 1732 – 1745 (2011)
39. Mills, A., Wiser, R.: Changes in the economic value of variable generation at high penetration levels: A pilot case study of california. Tech. rep., Lawrence Berkeley National Laboratory (2012)
40. Mills, A., Wiser, R.: An evaluation of solar valuation methods used in utility planning and procurement processes. Tech. rep., Lawrence Berkeley National Laboratory (2012)
41. Munoz, F.D.: Engineering-Economic Methods for Power Transmission Planning under Uncertainty and Renewable Resource Policies. Ph.D. thesis, The Johns Hopkins University, Baltimore (2014)
42. Munoz, F.D., Hobbs, B.F., Ho, J., Kasina, S.: An Engineering-Economic Approach to Transmission Planning Under Market and Regulatory Uncertainties: WECC Case Study. IEEE Transactions on Power Systems **29**(1), 307–317 (2014)
43. Munoz, F.D., Hobbs, B.F., Kasina, S.: Efficient Proactive Planning to Accommodate Renewables. In: IEEE Power and Energy Society General Meeting, pp. 1–7 (2012)
44. Munoz, F.D., Hobbs, B.F., Watson, J..P.: New Bounding and Decomposition Methods for MILP Investment Problems: Multi-Area Transmission and Generation Planning Under Policy Constraints. JHU Working Paper (2014)

45. Munoz, F.D., Sauma, E.E., Hobbs, B.F.: Approximations in Power Transmission Planning: Implications for the Cost and Performance of Renewable Portfolio Standards. Journal of Regulatory Economics **43**(3), 305–338 (2013)
46. Oliveira, G., Costa, A., Binato, S.: Large scale transmission network planning using optimization and heuristic techniques. IEEE Transactions on Power Systems **10**(4), 1828–1834 (1995)
47. O'Neill, R.P., Krall, E.A., Hedman, K.W., Oren, S.S.: A Model and Approach to the Challenge Posed by Optimal Power Systems Planning. Mathematical Programming **140**(2), 239–266 (2013)
48. ONeill, R.P., Dautel, T., Krall, E.: Recent ISO software enhancements and future software and modeling plans. Retrieved August 30, 2014, from http://www.ferc.gov/industries/electric/indus-act/rto/rto-iso-soft-2011.pdf (2011)
49. Papavasiliou, A., Oren, S.S.: Multiarea stochastic unit commitment for high wind penetration in a transmission constrained network. Operations Research **61**(3), 578–592 (2013)
50. Park, H., Baldick, R.: Transmission Planning Under Uncertainties of Wind and Load: Sequential Approximation Approach. IEEE Transactions on Power Systems **PP**(99), 1–8 (2013)
51. Pereira, M.V.F., Pinto, L.M.V.G., Cunha, S.H.F., Oliveira, G.C.: A Decomposition Approach to Automated Generation Transmission Expansion Planning. IEEE Transactions on Power Apparatus and Systems **104**(11), 3074–3083 (1985)
52. Pierre-Louis, P., Bayraksan, G., Morton, D.P.: A Combined Deterministic and Sampling-Based Sequential Bounding Method for Stochastic Programming. Proceedings of the 2011 Winter Simulation Conference (Wsc) pp. 4167–4178 (2011)
53. Price, J.E., Goodin, J.: Reduced Network Modeling of WECC as a Market Design Prototype. In: IEEE Power and Energy Society General Meeting, pp. 1–6 (2011)
54. Reis, F.S., Carvalho, P., Ferreira, L.: Reinforcement scheduling convergence in power systems transmission planning. IEEE Transactions on Power Systems, **20**(2), 1151–1157 (2005)
55. Rockafellar, R.T., Wets, R.J.B.: Scenarios and Policy Aggregation in Optimization under Uncertainty. Mathematics of Operations Research **16**(1), 119–147 (1991)
56. Romero, R., Monticelli, A.: A hierarchical decomposition approach for transmission network expansion planning. IEEE Transactions on Power Systems, **9**(1), 373–380 (1994)
57. Ryan, S.M., Wets, R.J.B., Woodruff, D.L., Silva-Monroy, C., Watson, J.P.: Toward scalable, parallel progressive hedging for stochastic unit commitment. In: 2013 IEEE Power and Energy Society General Meeting (PES),, pp. 1–5. IEEE (2013)
58. Samarakoon, H., Shrestha, R., Fujiwara, O.: A Mixed Integer Linear Programming Model for Transmission Expansion Planning with Generation Location Selection. International Journal of Electrical Power & Energy Systems **23**(4), 285–293 (2001)
59. da Silva, E.L., Gil, H.A., Areiza, J.M.: Transmission network expansion planning under an improved genetic algorithm. In: Proceedings of the 21st 1999 IEEE International Conference on Power Industry Computer Applications, 1999., pp. 315–321. IEEE (1999)
60. Silva, I., Rider, M., Romero, R., Murari, C.: Transmission Network Expansion Planning Considering Uncertainty in Demand. IEEE Transactions on Power Systems **21**(4), 1565–1573 (2006)
61. Stein, M.: Large sample properties of simulations using latin hypercube sampling. Technometrics **29**(2), 143–151 (1987)
62. Tibshirani, R., Walther, G., Hastie, T.: Estimating the Number of Clusters in a Data Set Via the Gap Statistic. Journal of the Royal Statistical Society Series B-Statistical Methodology **63**, 411–423 (2001)
63. Villasana, R., Garver, L., Salon, S.: Transmission network planning using linear programming. IEEE Transactions on Power Apparatus and Systems (2), 349–356 (1985)
64. Watson, J., Woodruff, D.: Progressive hedging innovations for a class of stochastic mixed-integer resource allocation problems. Computational Management Science **8**(4) (2011)
65. Watson, J.P., Woodruff, D., Hart, W.: Pysp: Modeling and solving stochastic programs in python. Mathematical Programming Computation **4**(2) (2012)

66. van der Weijde, A.H., Hobbs, B.F.: The Economics of Planning Electricity Transmission to Accommodate Renewables: Using Two-Stage Optimisation to Evaluate Flexibility and the Cost of Disregarding Uncertainty. Energy Economics **34**(6), 2089–2101 (2012)
67. Wogrin, S., Duenas, P., Delgadillo, A., Reneses, J.: A new approach to model load levels in electric power systems with high renewable penetration. IEEE Transactions on Power Systems **29**(5), 2210–2218 (2014)

9 Appendix

9.1 Cost data

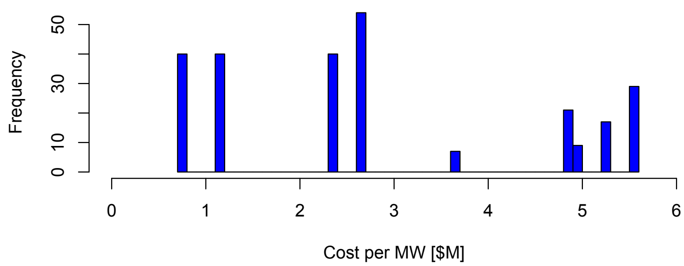

Fig. 3 Histogram of generation capital costs per MW.

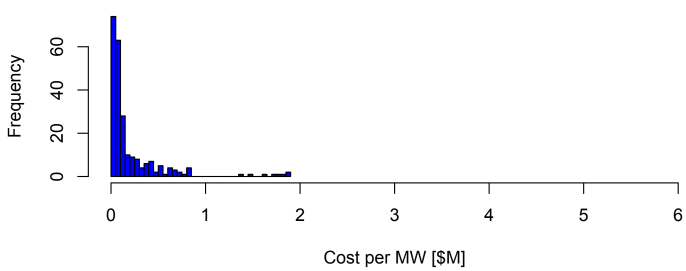

Fig. 4 Histogram of transmission capital costs per MW.