
Hierarchical Decision Making In Electricity Grid Management

Gal Dalal
Elad Gilboa
Shie Mannor
Technion, Israel

GALD@TX.TECHNION.AC.IL
EGILBOA@TX.TECHNION.AC.IL
SHIE@EE.TECHNION.AC.IL

Abstract

The power grid is a complex and vital system that necessitates careful reliability management. Managing the grid is a difficult problem with multiple time scales of decision making and stochastic behavior due to renewable energy generations, variable demand and unplanned outages. Solving this problem in the face of uncertainty requires a new methodology with tractable algorithms. In this work, we introduce a new model for hierarchical decision making in complex systems. We apply reinforcement learning (RL) methods to learn a *proxy*, i.e., a level of abstraction, for real-time power grid reliability. We devise an algorithm that alternates between slow time-scale policy improvement, and fast time-scale value function approximation. We compare our results to prevailing heuristics, and show the strength of our method.

1. Introduction

The power grid is a complex and vital system that requires high level of reliability. Reliability is of utmost importance, as the consequences of outages can be catastrophic. System operators (SOs) achieve reliability by means of sophisticated control operations and planning, which often require solving sequential stochastic decision problems. Sequential decision making under uncertainty in energy systems is studied in different communities such as control theory, dynamic programming, stochastic programming and robust optimization (Powell & Meisel, 2015; Bertsimas et al., 2013; Bienstock, 2011; Koutsopoulos & Tassiulas, 2012; Bienstock et al., 2014).

Reliability assessment and control are highly complicated tasks in complex real-world systems such as the power grid. Complications in the power grid arise because of strict physical restrictions, such as generation must meet

consumption continuously and transmission lines can not exceed their limited thermal capacity. Further complications stem from the structure of decision making in different time-horizons. For example, long-term system expansion and development such as building a new wind farm or a high-voltage line take years, mid-term asset management decisions such as performing maintenance are decided upon months in advance, short-term generation schedules are planned daily, and real-time operational control decisions are made on the scale of minutes. In these interdependent hierarchical decision making processes decisions are taken by multiple stakeholders. Furthermore, over the last decade, wind and solar energy sources become increasingly preeminent with further significant expansion being envisaged (Talbot, 2009). These generators introduce high uncertainty to the system, making the control task significantly more difficult. The complex dependence between multiple time-horizon with growing uncertainty, the curse of dimensionality when dealing with large systems, and the non-linear dependence of reliability measures to the multiple time-horizon decisions, make this problem extremely hard to tackle.

To stress the dimensionality complexity, consider the IEEE RTS-96 power network used in our experiments (Wong et al., 1999). This network is an example for a power grid of a medium sized European country or a state in the USA. Its state-space is $O(10^{300})$, and its action space is $O(10^{100})$; see Sec. 5. Assessment of each control choice carries a computational burden as it requires solving a set of non-linear trigonometric equations named alternating current power flow (ACPF); see Sec. 2.1.

Nowadays, the common practice in industry is solving large mixed integer programs (MIP), often with a linear relaxation, in an attempt to reach a valid solution (Grainger & Stevenson, 1994; Allan et al., 2013). Although this model is extensive, its computational burden makes it hard even for deterministic predictions (taking an order of a day in real-world systems), and inappropriate in the stochastic case. This limits SOs to sample snapshots of future grid states or analyze a few sequential trajectories. The narrow view of possible outcomes is likely to miss important benefits and increase the costs of decisions, thereby offering

little in terms of dealing with uncertainty.

To handle uncertainty, work has been done in stochastic optimization and control theory. These often use restrictive simplifications such as independence between the decision processes in the different time-scales or consider myopic decisions only (Abiri-Jahromi et al., 2009; Wu et al., 2010; Abiri-Jahromi et al., 2013).

Another approach is to use approximate dynamic programming (Powell, 2007; Si, 2004). However, the natural hierarchical structure of the problem, where several stakeholders operating in different time-scales and exposed to different information are making decisions with mutual influence, does not naturally fit the standard Markov Decision Process (MDP) structure. Furthermore, the problem is heavily constrained, since physical electrical restrictions must be met at all times.

Making this problem tractable requires a level of abstraction in the form of fast proxy methods to approximate the impact of real-time decisions on longer-term reliability and costs. To our knowledge, few attempts have been made to construct such proxies using tools from machine learning. An example for such, is the work conducted in a recent European project, iTesla (Vanfretti et al., 2012). This work focuses on analyzing snapshots of system states at different time points using data-mining methods. Then, classification and clustering algorithms are used for constructing security rules for predicting reliability level, given a failure and an electrical network state (Anil, 2013). Such approaches can aid SOs in real-time control, but lack the dynamic perspective of state-action evolution needed to evaluate consequences of policies in a sequential decision making scenario.

In this work we suggest a novel approach to mitigate the intractability of the hierarchical decision making problem of the day-ahead (DA) and real-time (RT) reliability of the power grid. The contributions of our work are:

- We introduce an interleaved MDPs hierarchical structure with separate state space, action space, and reward metric.
- We devise an algorithm that alternates between high-level policy improvement and lower-level value approximation, i.e., the policy improvement in the first MDP is based on the second MDP’s value function.
- We show the efficacy of our method on a medium-sized power grid problem.
- We introduce a new real-world application to the RL community and provide a simulation environment.

The rest of paper is organized as follows. In Sec. 2 we present background on power system engineering. In

Sec. 3, we formulate the two-layer MDPs. In Sec. 4, we introduce our interleaved approximate policy improvement (IAPI) algorithm and present results on the IEEE RTS-96 network. We conclude our work in Sec. 6.

2. Background

In this section we present a brief introduction to the field of power systems engineering. This is vast a field with extensive background and theory. For more information please refer to (Grainger & Stevenson, 1994; Allan et al., 2013).

2.1. Decision Processes and Power Flow in Power Grids

To better explain the multiple time-horizon decision processes we use a toy 6-bus power grid example (Wood & Wollenberg, 1996), shown in Fig. 1. The 6-bus system is composed of 6 electrical nodes referred to as “buses”. Each bus can have loads and generators attached to it. Loads (shown in blue) are consumers (e.g., large neighborhoods or cities and factories), and generators (shown in red) are power producers such as nuclear plants, coal plants, wind turbines, and solar panels. Load values change continuously throughout the day and closely follow daily, weekly, and yearly profiles. Controllable generators are operated such that the overall power generation meets the overall load at all times (up to transmission losses). The edges connecting the buses represent transmission lines which, due to thermal restrictions, can only transfer a limited amount of power before risking tripping.

Given a snapshot of loads and generation values, and the power grid topology (buses and transmission lines), it is possible to solve the complete alternating current power flow (ACPF) equations. The ACPF is a set of non-convex trigonometric equations that model the physical electrical characteristics of the power grid, i.e., voltage magnitude and angles of each node (Cain et al., 2012). The ACPF solution includes the amount of power passing through each transmission line (shown in green in Fig. 1).

In general, reliability of a power system is measured based on the avoidance of full or partial blackouts (both planned and unplanned) and their negative effect on social welfare. A blackout is an event where demand cannot be met. This can occur predominantly because of *contingencies* (i.e., asset malfunctions) which lead to unsafe operation and may require the SOs to disconnect loads in order to avoid catastrophes. Contingencies can stem from multiple causes, such as a tree falling, lightning strike, poor maintenance or exceeding the thermal limits of a transmission line. To maintain a high reliability level at all times, the current practice of SOs is to immunize the system against a pre-determined contingency list. A common choice for this list is all *single* asset contingencies, resulting in the so-called

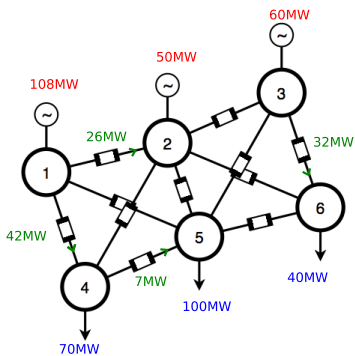


Figure 1. Wood & Wollenberg 6-bus system, with generation values in red; load values in green; and transmission line flow values in blue, obtained from an AC power-flow solution.

$N - 1$ reliability criterion.

However, contingency probabilities are difficult to obtain and their impact is hard to assess. Furthermore, the high penetration of stochastic and often uncontrollable renewable generators, makes the planning tasks significantly harder for several reasons. First, generation must equal demand at all times. Second, multiple decision making processes are taking place simultaneously on multiple time-scales. Third, each decision process involves high dimensional decision variables, and complex non-linear (Powell & Meisel, 2015), often intractable mathematical formulations.

For example, in the 6-bus system in Fig. 1, a system developer might plan to expand the system by building a new transmission line between buses 3 and 4. Expanding the grid is a long term process and a decision must be taken years in advance. However, this decision affects the future maintenance decisions, which will affect future daily planning that in turn affects the future real-time control room operations. Ideally, the system developer should consider all possible future realizations of the environment, grid, and the decision processes in all other time horizons.

2.2. Related work

Several works in the literature of power systems, operational research and more recently machine learning offer approaches for solving sequential stochastic problems using dynamic programming. Of which, the majority of these works focus on energy storage (Lai et al., 2010; Xi et al., 2014; Jiang et al., 2014; Scott & Powell, 2012), unit commitment (Padhy, 2004; Dalal & Mannor, 2015; Ernst et al., 2007), and energy market bidding strategies (Song & Wang, 2003; Urieli & Stone, 2014; Jiang & Powell, 2015). To our knowledge, no work has been done to use MDPs for assessing the reliability in power grids.

For our proxy abstraction we devise a hierarchical model.

Hierarchical models, offer several benefits over flat models when appropriate. They can improve exploration, enable learning from fewer trials, and allow faster learning for new problems by reusing subtasks learned on previous problems (Dietterich, 1998). Standard approaches for hierarchical models include: planning with options (often referred to as skills) (Sutton et al., 1999), task hierarchy (Barto & Mahadevan, 2003) and hierarchy of abstract machines (Parr & Russell, 1998). These models include levels of decision making that share the same state-space and a termination condition to switch between controllers. This structure does not fit our problem well where two separate decision makers run on *different* state-spaces and *temporal* resolutions.

3. Problem Formulation

Here we present a formulation for the two sequential decision processes occurring in the day ahead (DA) and real-time (RT) in terms of a hierarchical two MDP model. DA decisions are taken in order to maximize the system’s next day reliability. However, the next day reliability can only be assessed in RT, and is dependent on the system operator decision taken in RT. This results in a complex dependence between DA and RT actions and system reliability. We therefore formulate the problem using two layers of interleaved MDPs: a RT-MDP, describing the state of the system, reliability, and decisions on an hourly basis, and a DA-MDP describing the DA action of choosing a daily subset of active generators based on the upcoming day predictions. In our terminology, the former serves as a *proxy* for assessing decisions taken in the latter, see Fig. 2.

3.1. Day-Ahead MDP

The DA-MDP is a tuple $(\mathcal{S}^{\text{DA}}, \mathcal{A}^{\text{DA}}, \mathcal{P}^{\text{DA}}, \mathcal{R}^{\text{DA}})$. Time index is t_d , denoting days. Day-ahead state $s_{t_d}^{\text{DA}} \in \mathcal{S}^{\text{DA}}$ consists of a day ahead prediction of hourly demand on each bus, and wind generation of each wind generator.¹ Therefore, $\mathcal{S}^{\text{DA}} = \mathbb{R}^{T_D \cdot (n_b + n_g)}$, where T_D is the number of intraday time steps (24 in our case), and n_b , n_g are the number of buses and wind generators. For the day ahead action $a_{t_d}^{\text{DA}} \in \mathcal{A}^{\text{DA}}$ we use a simplified model which considers a binary vector indicating which generators participate in the next day’s generation process. The sets of generators contained in \mathcal{A}^{DA} represent common settings an SO can choose from. This set can be constructed by experts or inferred from data. An action $a_{t_d}^{\text{DA}}$ is chosen according to a policy $a_{t_d}^{\text{DA}} = \pi^{\text{DA}}(s_{t_d}^{\text{DA}})$. The next day state is chosen according to \mathcal{P}^{DA} , and is purely exogenous, i.e., $P(s_{t_d+1}^{\text{DA}} | s_{t_d}^{\text{DA}}, a_{t_d}^{\text{DA}}) = P(s_{t_d+1}^{\text{DA}} | s_{t_d}^{\text{DA}})$. The reward function \mathcal{R}^{DA} is a complicated function of the reliability in RT. The

¹In this work we consider only wind generation as a renewable source for simplicity.




Figure 2. Day-ahead and Real-Time hierarchical MDPs. The real-time process serves as a proxy for assessing decisions taken in day-ahead process.

complication arises since the DA decision affects not only the following day's RT reliability level, but also all subsequent days' RT reliability level. In addition, the subsequent days' RT reliability level is affected by future DA decisions. Since we cannot obtain the day ahead reward directly, we revert to use the RT reward as a surrogate for comparing DA policies. Notice that we cannot directly use the sum of RT rewards between consecutive days as a replacement for the DA reward since the model will no longer be Markovian.

3.2. Real-Time MDP

The RT-MDP is a tuple $(\mathcal{S}^{\text{RT}}, \mathcal{A}^{\text{RT}}, \mathcal{P}^{\text{RT}}, \mathcal{R}^{\text{RT}})$. It represents the real time reliability control process. Time index is t , denoting intra-day time steps (e.g., hours). In RT power network operation, an operator may choose preventive actions at each time step, trying to immunize the system against potential malfunctions by attempting to avoid unreliable states. We model this decision making process using post-decision states (Powell, 2007), where at the beginning of each time interval, the agent observes the current state $s_t^{\text{RT}} \in \mathcal{S}^{\text{RT}}$, i.e., the realized demand and wind values for this interval and chooses an action $a_t^{\text{RT}} \in \mathcal{A}^{\text{RT}}$. Following the agent's action, the system is now in a post-decision state $s_t^{\text{RT},a}$, which is the new state, after performing action a_t^{RT} from state s_t^{RT} . Next, exogenous random information W_t is obtained, informing whether equipment malfunction (contingency) occurred during time interval t . Given $s_t^{\text{RT},a}$ and W_t , the real time reward $r_t^{\text{RT}}(s_t^{\text{RT},a}, W_t)$ which represents the system's reliability, can be calculated, and a transition to s_{t+1}^{RT} occurs, governed by $P(s_{t+1}^{\text{RT}} | s_t^{\text{RT},a}, W_t)$. The history of this RT process can be written as $h_t^{\text{RT}} = (s_0^{\text{RT}}, a_0^{\text{RT}}, s_0^{\text{RT},a}, W_0, r_0^{\text{RT}}, s_1^{\text{RT}}, \dots, s_t^{\text{RT}})$.

3.2.1. REAL-TIME STATE SPACE

We define a RT state s_t^{RT} to be the tuple (d, w, g, e) , where:

d is a vector of stochastic nodal demand.

w is a vector of stochastic nodal wind generation.

g is a vector of controllable generation values. The DA action $a_{t_d}^{\text{DA}}$ determines which generators will have positive values, and which will be set to 0 throughout the

day. Each generator has minimal and maximal generation limits while in operation.

e is the topology of the grid. Includes information of current state of each edge (transmission line). $e \in \{0, 1, \dots, E\}$, where 0 is 'operational', and the others are countdown states until the line is fixed.

3.2.2. REAL-TIME ACTION SPACE

A RT action a_t^{RT} is a preventive action, that attempts to achieve better reliability of the system by immunizing against potential contingencies. The action involves redispatch Δg , i.e., change the generation values of the working controllable generators (chosen in DA):

$$s_t^{\text{RT}} \xrightarrow{a_t^{\text{RT}}} s_t^{\text{RT},a} = (d, w, g + \Delta g, e).$$

Any action is allowed as long as it is within the minimal and maximal generator limits. Notice that $\Delta g_i \neq 0$ for working generators only ($a_{t_d,i}^{\text{DA}} = 1$).

3.2.3. REAL-TIME TRANSITION KERNEL

The RT transition kernel can be factorized to exogenous transitions of demand, wind generation, and contingencies. It is conditioned on the last RT state and action (encoded in the RT post-state), and on the corresponding last DA decision taken to determine participating generators:

$$s_{t+1}^{\text{RT}} = f(s_t^{\text{RT},a}, W_t, a_{t_d}^{\text{DA}}).$$

The dependence between RT and DA states is expressed using two sets of equations. The first is RT demand process, based on DA demand prediction:

$$d_t^{\text{RT}} = d_t^{\text{DA}} + \delta_t \quad (1)$$

$$\delta_{t+1} = \delta_t + \epsilon_t, \quad (2)$$

where d_t^{RT} is the RT demand vector at time t , d_t^{DA} is the DA prediction vector for time t of the day, and ϵ_t is Gaussian noise. The dynamics in Eqs. (1)-(2) also hold for the wind generation process. For this work we chose a Brownian bias process for simplicity, however more complicated methods, such as in (Box et al., 2015; Papavasiliou & Oren, 2013; Taylor & Buizza, 2002), can be considered. The second equation coupling DA and RT determines the generators participating in current day generation process:

$$g_{t+1,i} = \begin{cases} g_{t,i} + \Delta g_{t,i} & \text{if } i \in \mathcal{I}(a_{t_d,i}^{\text{DA}}) \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

where $\mathcal{I}(a_{t_d,i}^{\text{DA}})$ is the index set of generators chosen by DA action $a_{t_d,i}^{\text{DA}}$.

Lastly, random exogenous information W_t specifies whether a contingency happened in the system, causing

transmission line i to fail, changing the network topology to e_{t+1} . The probability of line i to fail at each time-step is p_i if at the last time-step e_i was 0, and 0 otherwise.

3.2.4. REAL-TIME REWARD

We choose the RT reward to be the reliability level of the power system at the current time. To assess the level of reliability, we employ the common criterion used in the industry, termed $N - 1$, which assesses the system ability to withstand any contingency of a single asset.

To calculate the reliability of the system, it is examined using a sequence of tests (contingency list), where each test is an attempt to take out a single line (contingency) and check if the system retains safe operation. Hence, the reward $r_t^{\text{RT}}(s_t^{\text{RT},a}, W_t)$ is a number in $[0, 1]$, expressing the portion of tests passed out of the predetermined contingency list, which includes all single contingencies $c \in \mathcal{N}_{-1}$. The reliability is calculated for a given state of the grid, and is dependent of current topology (e_t) and the changes to the topology due to possible new contingencies (W_t). In practice, preserving the system in safe operation means being able to obtain a feasible solution to the power flow equations (see Sec. 2) of the network circuit. We define $\mathbb{I}_{[\text{PF}(c, s_t^{\text{RT},a}, W_t)]}$ to be 1 if a power flow solution exists, and 0 otherwise. As a result, the RT reward is:

$$r_t^{\text{RT}}(s_t^{\text{RT},a}, W_t) = \frac{1}{|\mathcal{N}_{-1}|} \sum_{c \in \mathcal{N}_{-1}} \mathbb{I}_{[\text{PF}(c, s_t^{\text{RT},a}, W_t)]}.$$

4. Interleaved Approximate Policy Improvement

In this section we present our algorithm, called *Interleaved Approximate Policy Improvement* (IAP), presented in Alg. 1, for jointly learning the RT reliability value function while searching for an optimal DA policy. We use the term *interleaved* since the policy improvement in one MDP is based on the second MDP's value function. We use simulation based value learning to assess the RT reliability of the system and the cross entropy method (De Boer et al., 2005; Szita & Lőrincz, 2006) for improving the DA policy. Our method scales to large systems since it uses simple models with carefully engineered features and designed to run on distributed computing. Since the algorithm is massively parallelizable, the more cores available, the faster the convergence will be.

Our goal is to find an optimal DA policy π^{DA} , under the assumption that the RT policy π^{RT} is known. Henceforth, we will use π to symbolize π^{DA} . As explained in Sec. 3, reliability is not explicitly defined on the DA level and we instead use the RT value function v^π as a surrogate for comparing between different DA policies. Differently than the common notation, v^π denotes the RT value function, under

the fixed RT policy π^{RT} , and a DA policy π .

Algorithm 1 IAPI Algorithm

Input: initial distribution $P_\psi^{(0)}$ for DA policy parameters

Output: optimal DA policy $\pi(\psi^*)$

- 1: initialize $\mathcal{S}_{test}^{\text{RT}} = \emptyset$, $k = 0$
 - 2: **repeat**
 - 3: **for** $i \leq N$ **do**
 - 4: draw $\psi_i \sim P_\psi^{(k)}$
 - 5: sample N_{episodes} trajectories using $\pi_i = \pi(\psi_i)$
 - 6: approximate v^{π_i} using TD(0)
 - 7: add TD(0) trajectories to $\mathcal{S}_{test}^{\text{RT}}$
 - 8: **end for**
 - 9: set $\hat{v}_i^{(k)} = \frac{1}{|\mathcal{S}_{test}^{\text{RT}}|} \sum_{s \in \mathcal{S}_{test}^{\text{RT}}} v^{\pi_i}(s)$, $\forall i \in [N]$
 - 10: rank policies π_i according to \hat{v}_i
 - 11: use ψ_i of the top percentile π_i to update $P_\psi^{(k)}$
 - 12: $k = k + 1$
 - 13: **until** convergence
-

Our method includes the following components:

Day Ahead Policy Approximation We define a parametric DA policy as $\pi(s^{\text{DA}}; \psi) = \arg \max_{a^{\text{DA}} \in \mathcal{A}^{\text{DA}}} \psi^\top \Phi(s^{\text{DA}}, a^{\text{DA}})$, where a^{DA} is the day ahead action dictating which generators will be active during the day, $\Phi(s^{\text{DA}}, a^{\text{DA}})$ are features of DA state s^{DA} and action a^{DA} .

A plausible choice for mapping DA state s^{DA} to an action a^{DA} is using multi-class classifiers. However, for large number of classes (20 in our experiments) these methods require a significant number of simulations for training (Bishop, 2006). Furthermore, approaches for classification-based policy learning often require obtaining multiple rollouts for all the actions from a state during the training procedure (Gabillon et al., 2011), which in our case will result in a full value evaluation per each action and might prove overly encumbering. To mitigate these complexities, our policy chooses the action that maximizes the inner product with ψ .

Real Time Value Function Approximation For a fixed DA policy π we approximate the RT value function using the TD(0) algorithm (Sutton & Barto, 1998); see Fig 3. The RT value function is parametrized as $v^\pi(s^{\text{RT}}; \theta_\pi) = \theta_\pi^\top \phi(s^{\text{RT}})$, where the parameter vector θ_π depends on π , and $\phi(s^{\text{RT}})$ are the features of RT state s^{RT} .

Day Ahead Policy Comparison A comparison between different DA policies π_i is done by calculating the empirical mean of RT value function $\mathbb{E}v^{\pi_i} \approx \frac{1}{|\mathcal{S}_{test}^{\text{RT}}|} \sum_{s \in \mathcal{S}_{test}^{\text{RT}}} v^{\pi_i}(s)$, using a set of representative RT initial states $\mathcal{S}_{test}^{\text{RT}}$. This set is composed of the full history of all




Figure 3. Day-ahead policy comparison using TD-learning of real-time value function.

RT states visited during the current IAPI iteration, enabling expected value estimation using many probable states with only linear computational complexity in $|\mathcal{S}_{\text{test}}^{\text{RT}}|$.

Day Ahead Policy Improvement using Cross Entropy

DA policy improvement is achieved using the cross entropy method (De Boer et al., 2005; Szita & Lőrincz, 2006). In this method, initial policies are sampled from a distribution $P_{\psi}^{(0)}$. Following which, in each iteration k policy parameters ψ are drawn from $P_{\psi}^{(k)}$, and their top percentile, according to the RT value, is used to update $P_{\psi}^{(k+1)}$ (De Boer et al., 2005; Szita & Lőrincz, 2006). In our experiments we set $P_{\psi}^{(0)}$ to be a distribution, whose support is the set of ψ 's that make $\psi^{\top} \Phi(s_0^{\text{DA}}, a_i^{\text{DA}})$ equal for all the different actions a_i^{DA} . The distribution $P_{\psi}^{(k)}$ is a Gaussian mixture with means set to $\psi^{(k-1)}$ that belong to the top percentile. The convergence criterion we use in our experiments is $\frac{1}{N_{\text{top}}} \sum_{i=1}^{N_{\text{top}}} (\hat{v}_i^{(k)} - \hat{v}_i^{(k-1)})^2 < \epsilon$, where N_{top} is the number of the top-percentile values. By using the cross entropy method, we avoid using gradient-based optimization which may be difficult to compute in our case due to the discrete, non-linear nature of ACPF solutions and their dependence of generation (Cain et al., 2012), which dictate the level of reliability.

The criterion for comparing policies is a parametric RT value function, $v^{\pi}(s^{\text{RT}}; \theta_{\pi})$, as oppose to using rollouts for policy evaluation (Gabillon et al., 2011). The reason for this choice is three-fold. First, since a rollout only explores a small part of the space, assuming a structure allows us to better generalize to unvisited states. This assumption is supported by our experiments; see Fig. 5. Second, this functional representation allows us to fairly compare different DA policies using a common set of representative RT initial states $\mathcal{S}_{\text{test}}^{\text{RT}}$. Third, our end-goal is to use the value function learned by this algorithm as a proxy for system reliability in RT.

5. Experiments

In this section we show results of IAPI algorithm on the IEEE RTS-96 test system, that is considered a standard test-




Figure 4. Diagram of the IEEE-RTS96 network we use for our experiments.

case in the power systems literature (Wong et al., 1999); see Fig. 4. This test-case is an example for a power grid of a medium-sized country, containing 73 buses, 99 generators, and 120 transmission lines. We updated the test-case to include 9 additional wind generators to better represent current power grids. We use daily demand and wind profiles based on real historical records as published in (Pandzic et al., 2015). As stated in Sec. 1, this is a complicated, high dimensional system, which cannot be solved using brute-force methods. The state space of this system can have $O(2^{120})$ line configurations, with $O(D^{73} \cdot G_r^9)$ demand values (D) and wind generation values (G_r) at each time, which are of a stochastic nature. This is without accounting for the day-ahead prediction, which will be the power of 24 of this number (for each hour of the day). Controlling which controllable generators are on/off makes $O(2^{99})$ integer decisions, and $O(G_c^{99-24})$ generation levels for G_c possible values per each generator.

To compose the DA action set \mathcal{A}^{DA} , we define 20 subsets of active generators chosen at random, and fix it for the rest of the simulation. These subsets contain varying numbers of generators with different capacities, to enable meeting demand for the different possible daily profiles. For the DA we use a $K + 4$ feature vector

$$\Phi(s^{\text{DA}}, a^{\text{DA}}) = (1, U_v, L_v, P, \mathbb{I}_{[a^{\text{DA}}=a_1]}, \dots, \mathbb{I}_{[a^{\text{DA}}=a_K]}),$$

where

K is the number of actions ($K = 20$ in our experiments).

U_v indicates if generation can meet maximal predicted daily demand.

L_v indicates if generation can meet minimal predicted daily demand.

P is a barrier penalty function that penalizes if the average demand is close to the upper or lower generation bounds achieved by a^{DA} .

$\mathbb{I}_{[a^{\text{DA}}=a_i]}$ is an indicator function over the selected DA action.

For the RT policy π^{RT} we employ a simple heuristic, of shifting the hourly generation values to meet the realized effective demand. We consider effective demand to be demand values minus wind generation values. This is a natural approach as wind generation is not under the decision maker’s control and therefore is not considered a part of regular controllable generation. The RT feature vector $\phi(s^{\text{RT}})$ contains polynomial features of (D, e_d, e_g) , where

D is the total RT effective demand,
 e_d is the demand entropy across the different buses, and
 e_g is the generation entropy across the different buses,

resulting in a 10 dimensional vector. We use the entropy feature since it compactly maps the spread of generation and demand across the network. The spread is important as concentrations of generation and demand are directly linked to reliability issues, see Fig. 5.

For parameters for the dynamics described in Eqs. (1)-(2) we use (Lu et al., 2013) and choose $\delta_0^w \sim \mathcal{N}(0, 0.05 \cdot w_0^{\text{RT}})$ for the wind forecast error, and $\delta_0^d \sim \mathcal{N}(0, 0.01 \cdot d_0^{\text{RT}})$ for the demand forecast error. The real time variation is chosen to be $\epsilon_t \sim \mathcal{N}(0, 0.05 \cdot \delta_0)$. Line failure probability p_i is set to $5 \cdot 10^{-4}$ for each line, and its time-fill-fix $E = 5$.

In our simulation we use $N_{\text{episodes}} = 50$ episodes, each with a 3 day horizon. Each episode starts from a random DA state s_0^{DA} , drawn from several representing demand and wind profiles, to which we add normally distributed noise. The next day transition corresponds to adding a normally distributed bias to the previous day profile. In each cross-entropy iteration we evaluate 200 DA policies ($N = 200$) and choose the top 20-th percentile for updating P_ψ . The DA policies are evaluated in parallel, on a 200 cores cluster. For the TD(0) algorithm we use discounting with $\gamma = 0.95$.

In Fig. 5 we show the learned RT value $v^\pi(s_t^{\text{RT}}; \theta_\pi)$, as a function of the deviation of the overall effective demand (demand minus wind) from the DA prediction, and generation entropy across the different buses. The RT value shown is marginalized over the rest of the features, time, and daily profiles. As shown in the figure, as the real-time demand deviates from the predicted demand, reliability suffers in a quadratic dependence. This is because the generators chosen in the DA will reach their upper or lower thresholds, causing generation to not meet the demand. The monotonic dependence in generation entropy implies that the higher the entropy the more reliable the system. This can be understood since high entropy corresponds to a more distributed generation throughout the network, mitigating the consequence of line outages. The reason this mitigation occurs is that when less emphasis is put on specific areas of the network, the system has more flexibility to find alternative routes from generation to demand. This, however, incurs a price in real life since generation cannot be concentrated only on cheap generators.

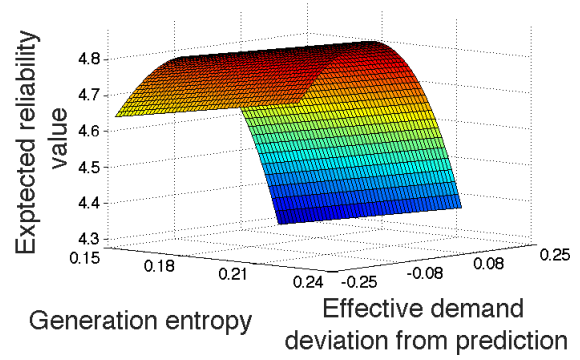


Figure 5. Learned RT value $v^\pi(s^{\text{RT}}; \theta_\pi)$ as a function of effective demand and generation entropy across the network.




Figure 6. Convergence of the IAPI algorithm. We show the top 20-th percentile, which is used in the algorithm to update the distribution P_ψ .

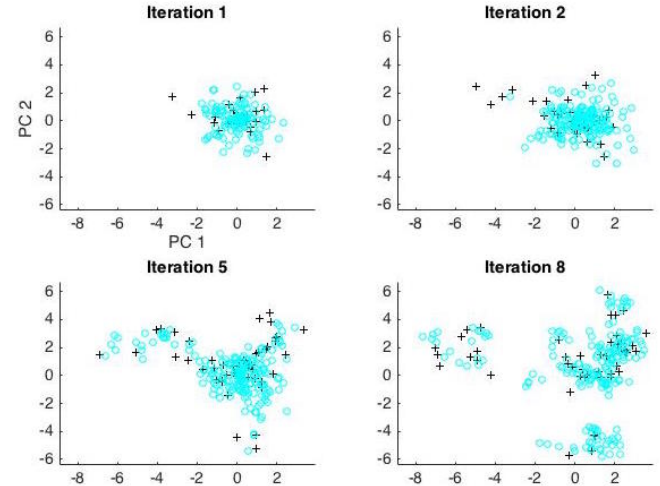


Figure 7. Projection of the top two principal components of the DA policy parameters ψ . The figure shows the scattering of the drawn policies parameters ψ_i in each iteration, where the dark dots mark the ψ s corresponding to top percentile \hat{v}_i .




Figure 8. Daily effective demand profiles, colored according to the chosen DA action using the policy learned by IAPI algorithm.

Next, in Fig. 6 we show the top 20-th percentile convergence of the IAPI algorithm. As can be seen, the average value is increasing and converging after 8 iterations, while the variance of the top percentile solutions is decreasing. In Fig. 7, we visualize the convergence of the IAPI algorithm by projecting on the top two principal components (PC) of the DA policy parameters ψ . We use the same PCs for all the plots. The figure shows the scattering of the drawn ψ_i in each iteration. As described in Alg. 1, each ψ_i defines a policy $\pi(\psi_i)$ for which we calculate the estimated expected value \hat{v}_i . The dark '+' mark the ψ s corresponding to the top percentile of \hat{v}_i . As can be seen, the IAPI algorithm explores the policy space until converging to local minima.

In Fig. 8 we present different daily effective demand profiles, colored according to the DA action chosen by the DA policy $\pi(s^{\text{DA}}; \psi)$, that was learned by the IAPI algorithm. A clear clustering can be observed between different daily demand profiles and the resulting action taken by the DA policy. The policy distinguishes between different consumption patterns and maps them to a corresponding set of active generators for reliable operation of the day to come.

To test our algorithm we compare the learned DA policy to three common heuristics. Taking the daily state as an input, these heuristics choose an eligible generator subset that can satisfy the maximal effective demand according to that day's DA prediction. The difference between them is the way they choose among the eligible subsets of each day. 'Random' chooses one at random, 'Cost' chooses the cheapest combination of generators, and 'Elastic' chooses the subset with the most flexible generators, having the largest ratio between upper and lower generation limits. We evaluate the performance of the different policies using rollouts of 2000 episodes per each policy. Fig. 9 presents the box-plots of the results. As can be seen, the value varies greatly between the different methods. In the 'Random' policy, there is an almost flat spread, demonstrating a lack of preference for a single subset when encountering a new day. The 'Cost' and 'Elastic' policies produce a more concentrated spread, corresponding to their preference of subset choices. The policy learned using IAPI obtains higher




Figure 9. Box-plot summary of the three heuristic policies and the policy learned using the IAPI algorithm. Higher is better.

reward than the heuristics. This result shows the IAPI algorithm's ability to learn a diverse DA policy.

6. Discussion

In this work we present an interleaved two-MDP model, inspired by the hierarchical decision making problem of managing power grid reliability. The IAPI algorithm alternates between improving the DA policy and learning the RT reliability value. The IEEE RTS-96 network in our experiments is a large enough network to capture computational complexities that arise in real-world networks. Here we focus on the power grid, however our model can be adapted to other important applications with an hierarchical decision making structure in different time-scales where high level of reliability and sustainability is required. Examples are sewer systems, smart cities and traffic control.

The coarse model presented in this work was crafted jointly as an initial step with several SOs. This work is the tip of the iceberg and many enhancements can be considered. For example, an important aspect that is not covered by it is budget consideration. Following the practice in the power system industry, reliability and money are often treated as different "currencies". Considering a budget will impose limitations on action selection and will complicate this problem even more. Another addition that can be made to extend the IAPI algorithm to interleave in reverse, i.e., alternating the DA improvement with improving the RT policy. Suspected drawbacks in this case are convergence problems, and the need for even more intense simulation.

Managing high reliability in stochastic complex systems, with interleaved decision making in different time horizons, is inherently difficult and results in intractable formulations. To mitigate this, there is a growing interest in the power system community to utilize proxies that will enable quick assessment of reliability for different states of the grid. In this work we introduce new models and formulations, along with a simulation environment. Our hope is that this will provide a platform for other researchers in the community to develop and explore their own innovative methods, and will help to bring these two fields closer. The code for the simulation environment is available at https://github.com/gald1/icml16_iapi.

Acknowledgments

We wish to thank the reviewers for their valuable comments. The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement No 608540, project acronym GARPUR. The scientific responsibility rests with the authors.

References

- Abiri-Jahromi, A, Fotuhi-Firuzabad, M, and Abbasi, E. An efficient mixed-integer linear formulation for long-term overhead lines maintenance scheduling in power distribution systems. *Power Delivery, IEEE Transactions on*, 24(4):2043–2053, 2009.
- Abiri-Jahromi, Amir, Parvania, Masood, Bouffard, Francois, and Fotuhi-Firuzabad, Mahmud. A two-stage framework for power transformer asset maintenance management – Part I: Models and formulations. *Power Systems, IEEE Transactions on*, 28(2):1395–1403, 2013.
- Allan, RN et al. *Reliability evaluation of power systems*. Springer Science & Business Media, 2013.
- Anil, Can. Benchmarking of data mining techniques as applied to power system analysis. 2013.
- Barto, Andrew G and Mahadevan, Sridhar. Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems*, 13(4):341–379, 2003.
- Bertsimas, Dimitris, Litvinov, Eugene, Sun, Xu Andy, Zhao, Jinye, and Zheng, Tongxin. Adaptive robust optimization for the security constrained unit commitment problem. *Power Systems, IEEE Transactions on*, 28(1): 52–63, 2013.
- Bienstock, Daniel. Optimal control of cascading power grid failures. In *Decision and control and European control conference (CDC-ECC), 2011 50th IEEE conference on*, pp. 2166–2173. IEEE, 2011.
- Bienstock, Daniel, Chertkov, Michael, and Harnett, Sean. Chance-constrained optimal power flow: Risk-aware network control under uncertainty. *SIAM Review*, 56(3): 461–495, 2014.
- Bishop, Christopher M. Pattern recognition. *Machine Learning*, 2006.
- Box, George EP, Jenkins, Gwilym M, Reinsel, Gregory C, and Ljung, Greta M. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.
- Cain, Mary B, Oneill, Richard P, and Castillo, Anya. History of optimal power flow and formulations. *Federal Energy Regulatory Commission*, 2012.
- Dalal, Gal and Mannor, Shie. Reinforcement learning for the unit commitment problem. In *PowerTech, 2015 IEEE Eindhoven*, pp. 1–6. IEEE, 2015.
- De Boer, Pieter-Tjerk, Kroese, Dirk P, Mannor, Shie, and Rubinstein, Reuven Y. A tutorial on the cross-entropy method. *Annals of operations research*, 134(1):19–67, 2005.
- Dietterich, Thomas G. The MAXQ method for hierarchical reinforcement learning. In *ICML*, pp. 118–126. Citeseer, 1998.
- Ernst, Damien, Glavic, Mevludin, Stan, Guy-Bart, Mannor, Shie, and Wehenkel, Louis. The cross-entropy method for power system combinatorial optimization problems. In *2007 Power Tech*, 2007.
- Gabillon, Victor, Lazaric, Alessandro, Ghavamzadeh, Mohammad, and Scherrer, Bruno. Classification-based policy iteration with a critic. 2011.
- Grainger, John J and Stevenson, William D. *Power system analysis*. McGraw-Hill, 1994.
- Jiang, Daniel R and Powell, Warren B. Optimal hour-ahead bidding in the real-time electricity market with battery storage using approximate dynamic programming. *INFORMS Journal on Computing*, 27(3):525–543, 2015.
- Jiang, Daniel R, Pham, Thuy V, Powell, Warren B, Salas, Daniel F, and Scott, Waymond R. A comparison of approximate dynamic programming techniques on benchmark energy storage problems: Does anything work? In *Adaptive Dynamic Programming and Reinforcement Learning (ADPRL), 2014 IEEE Symposium on*, pp. 1–8. IEEE, 2014.
- Koutsopoulos, Iordanis and Tassioulas, Leandros. Optimal control policies for power demand scheduling in the smart grid. *Selected Areas in Communications, IEEE Journal on*, 30(6):1049–1060, 2012.
- Lai, Guoming, Margot, François, and Secomandi, Nicola. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations research*, 58(3):564–582, 2010.
- Lu, Ning, Diao, Ruisheng, Hafen, Ryan P, Samaan, Nancy, and Makarov, Yuri V. A comparison of forecast error generators for modeling wind and load uncertainty. In *Power and Energy Society General Meeting (PES), 2013 IEEE*, pp. 1–5. IEEE, 2013.
- Padhy, Narayana Prasad. Unit commitment-a bibliographical survey. *Power Systems, IEEE Transactions on*, 19(2):1196–1205, 2004.

- Pandzic, Hrvoje, Wang, Yannan, Qiu, Ting, Dvorkin, Yury, and Kirschen, Daniel S. Near-optimal method for siting and sizing of distributed storage in a transmission network. 2015.
- Papavasiliou, Anthony and Oren, Shmuel S. Multiarea stochastic unit commitment for high wind penetration in a transmission constrained network. *Operations Research*, 61(3):578–592, 2013.
- Parr, Ronald and Russell, Stuart. Reinforcement learning with hierarchies of machines. *Advances in neural information processing systems*, pp. 1043–1049, 1998.
- Powell, Warren B. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons, 2007.
- Powell, Warren B and Meisel, Stephan. Tutorial on stochastic optimization in energy – Part I: Modeling and Policies. 2015.
- Scott, W and Powell, Warren B. Approximate dynamic programming for energy storage with new results on instrumental variables and projected bellman errors. *Submitted to Operations Research (Under Review)*, 2012.
- Si, Jennie. *Handbook of learning and approximate dynamic programming*, volume 2. John Wiley & Sons, 2004.
- Song, Yong-Hua and Wang, Xi-Fan. *Operation of market-oriented power systems*. Springer Science & Business Media, 2003.
- Sutton, Richard S and Barto, Andrew G. *Reinforcement learning: An introduction*. MIT press, 1998.
- Sutton, Richard S, Precup, Doina, and Singh, Satinder. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1):181–211, 1999.
- Szita, István and Lőrincz, András. Learning Tetris using the noisy cross-entropy method. *Neural computation*, 18(12):2936–2941, 2006.
- Talbot, David. Lifeline for renewable power. *Technol Rev*, 112:40–47, 2009.
- Taylor, James W and Buizza, Roberto. Neural network load forecasting with weather ensemble predictions. *Power Systems, IEEE Transactions on*, 17(3):626–632, 2002.
- Urieli, Daniel and Stone, Peter. Tactex'13: a champion adaptive power trading agent. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pp. 1447–1448. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- Vanfretti, Heyberger, ELIA, NGC, Statnett, IPTO, KTH, DTU, Tractebel, and RTE. Innovative tools for electrical system security within large areas. <http://www.itesla-project.eu/>, 2012. Accessed: 2016-02-03.
- Wong, Paul, Albrecht, P, Allan, R, Billinton, Roy, Chen, Qian, Fong, C, Haddad, Sandro, Li, Wenyuan, Mukerji, R, Patton, Diane, et al. The IEEE reliability test system-1996. a report prepared by the reliability test system task force of the application of probability methods subcommittee. *Power Systems, IEEE Transactions on*, 14(3): 1010–1020, 1999.
- Wood, Allen J and Wollenberg, B. Power generation operation and control 2nd edition. In *Fuel and Energy Abstracts*, volume 3, pp. 195, 1996.
- Wu, Lei, Shahidepour, Mohammad, and Fu, Yong. Security-constrained generation and transmission outage scheduling with uncertainties. *Power Systems, IEEE Transactions on*, 25(3):1674–1685, 2010.
- Xi, Xiaomin, Sioshansi, Ramteen, and Marano, Vincenzo. A stochastic dynamic programming model for co-optimization of distributed energy storage. *Energy Systems*, 5(3):475–505, 2014.