Reinforcement Learning based Proactive Control for Transmission Grid Resilience to Wildfire

S. U. Kadir, *Student Member, IEEE,* S. Majumder, *Member, IEEE,* A. Chhokra, *Member, IEEE,* A. Dubey, *Senior Member, IEEE,* H. Neema, A. Laszka, and A. Srivastava, *Senior Member, IEEE*

Abstract—Power grid operation subject to an extreme event requires decision-making by human operators under stressful condition with high cognitive load. Decision support under adverse dynamic events, specially if forecasted, can be supplemented by intelligent proactive control. Power system operation during wildfires require resiliency-driven proactive control for load shedding, line switching and resource allocation considering the dynamics of the wildfire and failure propagation. However, possible number of line- and load-switching in a large system during an event make traditional prediction-driven and stochastic approaches computationally intractable, leading operators to often use greedy algorithms. We model and solve the proactive control problem as a Markov decision process and introduce an integrated testbed for spatio-temporal wildfire propagation and proactive power-system operation. We transform the enormous wildfire-propagation observation space and utilize it as part of a heuristic for proactive de-energization of transmission assets. We integrate this heuristic with a reinforcement-learning based proactive policy for controlling the generating assets. Our approach allows this controller to provide setpoints for a part of the generation fleet, while a myopic operator can determine the setpoints for the remaining set, which results in a symbiotic action. We evaluate our approach utilizing the IEEE 24-node system mapped on a hypothetical terrain. Our results show that the proposed approach can help the operator to reduce load loss during an extreme event, reduce power flow through lines that are to be de-energized, and reduce the likelihood of infeasible power-flow solutions, which would indicate violation of short-term thermal limits of transmission lines.

Index Terms—Power System Operation, Proactive Control, Reinforcement Learning, Resiliency, Wildfire.

I. INTRODUCTION

THE frequency of extreme weather events (also classified as high-impact low-frequency, or HILF, events), such as storms, floods, and wildfires, as well as the associated impact on the power grid have soared in recent years [1]. Changing climate raises the potential for frequent wildfire events. Furthermore, increased ambient temperature due to approaching wildfires, or heatwaves in general, can result in sags, expanding the potential of faults for power transmission lines [2]. Additionally, high winds can blow nearby vegetation into transmission lines, which can create sparks and snap transmission lines with the potential of originating secondary

S. U. Kadir and A. Laszka are with the University of Houston, Houston, TX.
S. Majumder and A. Srivastava are with the Washington State University, Pullman, WA (e-mail: anurag.k.srivastava@wsu.edu).
A. Chhokra, A. Dubey and H. Neema are with the Vanderbilt University, Nashville, TN.

wildfire hazards [3]. To prevent these secondary sources, several utilities in the USA typically restrict power flow through some of their assets during emergency events (e.g., public safety power shut-off, or PSPS, events in the state of California, USA [4]). Compared to other extreme weather events, the slow progression of wildfires [5] across the large geographical span of power systems provides an opportunity for proactive control. In particular, slow-moving nature of these events provides operators with sufficient time to actively control line flows before the preemptive de-energization of power-system assets. Here, we define proactive control as any pre-event or during-event action to minimize the expected impact of an evolving extreme event.

From an operational point of view, if a heavily loaded line suffers an outage due to sags or if transmission lines exceed their short-term overload capacity, the necessary power flow redistribution may lead to a cascaded outage (as observed in the 1977 New York black-out [6]). Since the cost of power-system restoration can be enormous, it is crucial to reroute transmission flow with necessary de-energization proactively, while ensuring that consumers experience the minimum inconvenience of power shut off. While the use of traditional proactive control strategies for power-system resiliency is abundant [7], most of them revolve around controlling the operation of the distribution network. However, the bulk of power still flows through the transmission network, and while periodic vegetation management around the transmission infrastructure may help, wildfire threats exist nonetheless. Hence, the proactive control of transmission assets has to be emphasized. Eventually, this will help operators to avoid power-system operational reliability issues, while indirectly reducing the chances of secondary wildfires and event impact propagation.

Traditional power transmission system operation is governed by power system economics and the $N-1$ operational reliability criterion. However, if the emergency warning and caution (EWAC) direction is received from various stakeholders, operators typically aim to predict outage risks to take necessary control measures to minimize loss of loads, and systems will move to resiliency mode. Techniques, such as Markovian model-based proactive sequential re-dispatch of generators [8], proactive splitting the transmission grid into islands [9], [10], transmission system reconfiguration [11], the coordinated control of multiple microgrids connected via transmission system [12], [13], defensive islanding formation [14], optimal power shut-offs [15], transmission line derating [2], [16], are common approaches for resilient transmission

network operation and control. Abundant monitoring data from SCADA and PMUs can be leveraged in outage forecasting, which can be utilized for proactive resource allocation and deployment of necessary measures to serve consumers during emergency conditions. In this regard, statistical models of outage duration [17], [18] can be applied by the operators. However, most of these techniques employ classical optimization-based approaches and are combinatorial in nature, making the control task computationally challenging and resource-intensive to handle in real-time. Furthermore, recent events have shown the unsuitability of traditional forecast models in an effective determination of outage risk (e.g., majority of the wind power generators were outaged due to lack of winterization in the 2021 Texas power outage event, and this was not captured in the generation forecast [19]). Likewise, traditional model-predictive control techniques are also quite resource-intensive for real time.

Alternatively, machine learning (ML) approaches (e.g., [20]) can be suitably integrated by operators into proactive decision-making. ML models that were trained *apriori* can be leveraged to significantly reduce computation time during an event, offloading most of the decision-making onto the trained controller, which results in fast and efficient decision-making [21]. Wildfire propagation data (see ref. [22] for one such sample dataset) along with historical power-system monitoring and control data can be leveraged to this end. The controller can complement the action taken by the power system operator, resulting in an efficient operation. This also alleviates the difficulties of traditional proactive optimization and rule-based approaches while minimizing the possibility of human errors due to operating in a stressful environment. In this regard, reinforcement learning (RL) based power-system decision support has gained significant traction in recent years [23]. This, along with recent advances in RL-based control [24], indicate the plausibility of successful deployment of ML-based techniques for proactive control in the advent of a disaster.

In this paper, we propose a novel approach to aid power-system operators in effective proactive control of available resources during wildfires and develop an associated simulation testbed for training and validation. The core contributions of this paper are as follows:

- Modeled the spatio-temporal dynamics of a wildfire and the operation of a transmission systems as a Markov decision process, which captures the impact of the wildfire on the power-system assets. Our model allows the decision-support agent to control generation setpoints and de-energize lines working in tandem with existing power system operation and associated controller.
- Formulated the optimal proactive control problem for minimizing load shedding, the de-energization of power-system assets proactively, and the likelihood of infeasible power-flow solutions, considering the time horizon of the wildfire event. Proactive de-energization of line avoids outage of live line and impact propagation.
- Proposed a novel approach, which is an ensemble of a compact representation for the agent's observation of the wildfire state, a heuristic algorithm for the proactive de-energization of power lines, and a deep reinforcement

learning based approach trained using Deep Deterministic Policy Gradient (DDPG) for the proactive control of power generation.
- Developed an integrated testbed combining a wildfire and power-system simulator (available as open-source in the future). Testbed was used to train and evaluate developed agent on the IEEE transmission system test system mapped onto a topographical map. We demonstrate that our proactive approach can significantly reduce system impact compared to reactive and myopic control policies.

II. MODEL DESCRIPTION AND PROBLEM FORMULATION

Fig. 1 shows the high-level overview of the integrated power-system operation supplemented by an external controller to avoid catastrophic damage caused by the wildfire events modeled as a Markovian decision process (MDP). The integrated method comprises four major components: (i) historical data, (ii) offline integrated testbed, (iii) training and offline control environment, and (iv) online deployment of the controller in the natural environment. Here, accumulated historical data can be suitably leveraged to develop the wildfire transition model and also by the operator both during the training and real-time operation. All these historical datasets can also be leveraged for the development of a crude controller for further training. This section provides a detailed treatment of wildfire propagation and power-system operation modeling and formulates the control problem as MDP. A detailed discussion of the controller and associated training will be provided in the following section.

A. Wildfire Model

1) Wildfire Propagation Dynamics in Topographical Space: We utilize a stochastic model as given in ref. [25] for wildfire propagation dynamics. We divided the geographical region into multiple cells and defined it by X as a grid cell set. The temporal horizon is also divided into uniform-length contiguous steps identified by the variable k. The state $s_{x,k}^f$ of fire in each cell $x \in X$ at the beginning of k^{th} interval (i.e., at the k^{th} time step) is captured by a Boolean variable d_k^x, representing the status of fire, and an integer variable $h_k^x (\geq 0)$, representing available wildfire fuel within the cell (determining the duration of fire).

The status of fire d_k^x in cell x can be either one of two states: non-ignited $d_k^x = 0$ and ignited $d_k^x = 1$. Once a cell is ignited, it consumes fuel at a constant rate of $\mathbf{C}^x$ until it exhausts (burns) all fuel in the cell. Precise fuel dynamics are given by (1). Once all the fuel is burnt, the cell returns to the non-ignited state.

$$h_k^x = \begin{cases} h_{k-1}^x & \text{if } \neg d_{k-1}^x \vee h_{k-1}^x \leq 0 \\ h_{k-1}^x - \mathbf{C}^x & \text{otherwise.} \end{cases} \quad (1)$$

The evolution of d_k^x is stochastic and driven by the transition probability ρ_k^x. Specifically, the probability of cell x being ignited at k^{th} time step (i.e., $d_k^x = 1$) is given by (2).

$$\rho_k^x = \begin{cases} 0 & \text{if } \neg d_{k-1}^x \wedge |\mathcal{H}_k^x| = 0 \\ 1 - \prod_{y \in \mathcal{H}_k^x} (1 - \mathbf{P}_{x,k}^y) & \text{if } \neg d_{k-1}^x \wedge |\mathcal{H}_k^x| > 0 \\ 1 & \text{otherwise.} \end{cases} \quad (2)$$
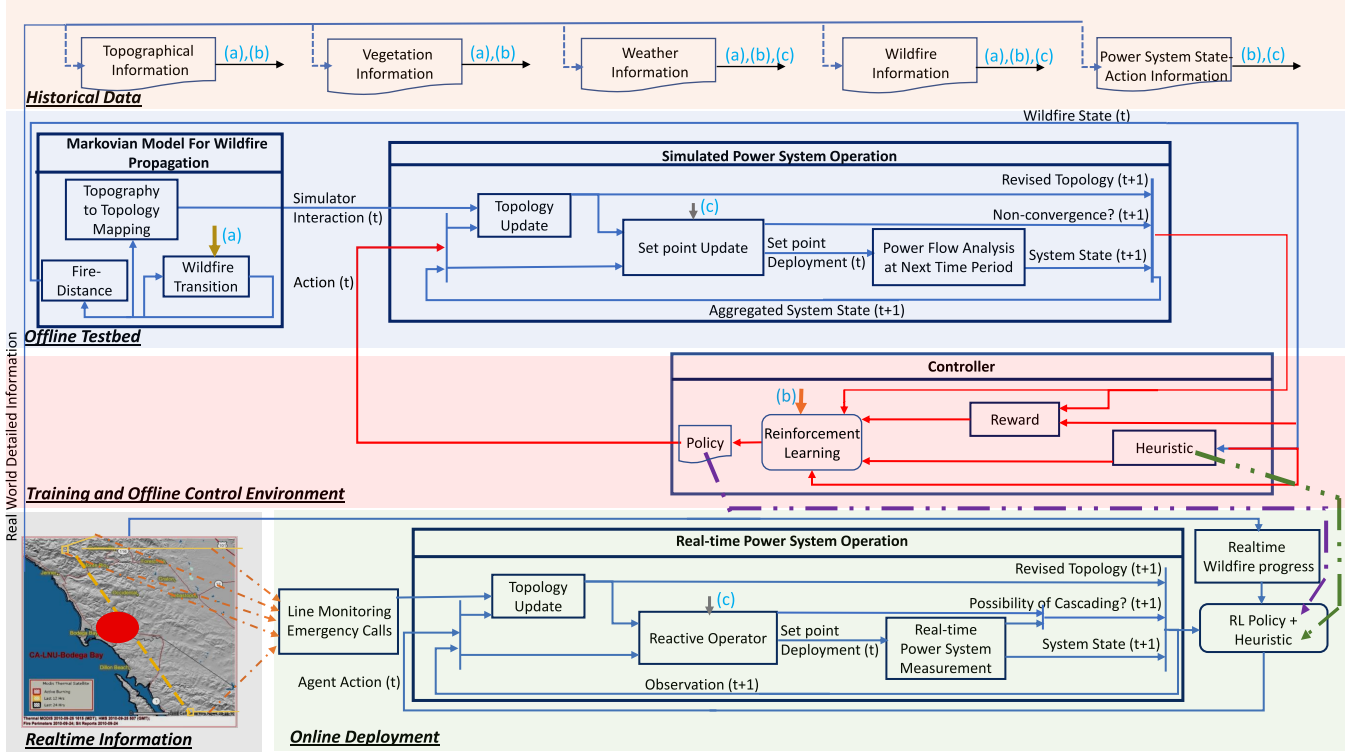
Fig. 1: An integrated wildfire propagation and reactive power system operation — testbed, training, and deployment.

Here, $\mathrm{P}^y_{x,k}$ denotes the probability of fire spreading from cell y to x. For a given cell x, $\mathcal{H}^x_k$ is the set of neighboring cells that can contribute to the spreading of fire to cell x in k^{th} time step, that is, $\mathcal{H}^x_k$ is the set of neighboring cells y for which $d^y_k = 1$. Note that geographical and environmental factors are implicitly captured by the cell-to-cell wildfire spread probability $\mathrm{P}^y_{x,k}$. Real world topographical information can be stored as historical data and utilized to determine or update these probabilities. However, the calculation of these probabilities is beyond the scope of this paper, and we assume them to be given.

As shown in (3), the state of the whole topographical grid s^f_k can be obtained by composing the state of every cell.

$$s^f_k = \langle (s_{1,k}), \ldots, (s_{M,k}) \rangle$$
$$\mathcal{L}(s^f_k) = \langle d^1_k, \ldots, d^M_k \rangle \tag{3}$$

where $(1, ..., M)$ are M cells in the grid.

Note that the selection of a larger spatiotemporal grid would speed up the simulation of the fire-propagation model at the expense of a significant reduction in model accuracy [26]. Therefore, it is assumed that the fire propagation model operates on a finer timescale, and the result of the simulation is appropriately down-sampled for determining the outcome of wildfire-power system interactions. The selection of optimal grid size is beyond the scope of this paper.

2) Topography to Topology Mapping: Given the wildfire dynamics in Section II-A1, here we model the impact of propagating wildfire on the power system operation. The power system environment is geotagged with cell information, and the state of one or more cells—in terms of $\mathcal{L}(s^x_k)$—correspond

to a given power system asset or equipment (transmission lines or substations). Topologically, the power system can be represented as a collection of nodes, N, representing generation and transmission substations, connected through a set of transmission lines, T. At a given time step, for each node $i \in N$ and branch $t \in T$, given cell-level fire propagation status obtained earlier, we assign binary variables $z^f_{i,k}$ and $z^f_{t,k}$ to indicate the operational status of substations and transmission lines respectively ($= 1$ denotes availability, and $= 0$ denotes assets encroached by fire) and is captured using (4).

$$z^f_{(\cdot),k} = \begin{cases} 0 & \text{if } \exists x \in G_{(\cdot)} \text{ s.t. } \mathcal{L}(s^x_k) \neq \text{non-ignited} \\ 1 & \text{otherwise} \end{cases} \tag{4}$$

where set $G_{(\cdot)}$ symbolizes the cells corresponding to a given power system asset. Although the wildfire propagation model operates on a finer time scale compared to power system operation, we reuse the variable k to determine power system operational time steps without the loss of generality. In the real-world deployment, this input mimics the emergency call from the generating and transmission substations, or thermal overload condition of the transmission lines [27] based on temperature measurement, indicating wildfire encroachment.

B. Power System Operation Model

Contrary to an existing energy management system (EMS) providing decision-support to the operator in solving traditional multi-period scenario-based deterministic or stochastic operational optimization problem for the decision making, as

shown in Fig. 1, we treat our system operator as a myopic entity, actively seeking the recommendation of the controller in the decision-making in the wake of the disaster, while monitoring the state of the system along with the system-wide emergency condition. Since the applicability of the proposed controller is limited to the wildfire time horizon, and the operator is expected to return to economic mode following expiration emergency conditions, the operating horizon of the controller is finite. Operator actions can be divided into two stages:

1) Topology Revision: We consider that the action taken by the operator is also discrete in time, and is based on the observation made at the beginning of k^{th} interval, or at k^{th} time step. Here, the operator receives the emergency responses from the wildfire simulation block, given by $z_{i,k}^{f}$ and $z_{t,k}^{f}$ (see (4)), and substation and transmission line shut-off actions from the controller, given as $z_{i,k}^{e}$ and $z_{t,k}^{e}$ respectively, and first deploys them faithfully[1]. Thus, transmission asset switching actions are not automatic, rather, those actions will be supervised by the operator. Consequently, the revised operational status of a substation and transmission line for the $k+1^{th}$ time step becomes $z_{i,k+1}^{o} = z_{i,k}^{e} z_{i,k}^{f} z_{i,k}^{o}$ and $z_{t,k+1}^{o} = z_{t,k}^{e} z_{t,k}^{f} z_{t,k}^{o}$ respectively (where $= 1$ denotes availability, and $= 0$ denotes shut-off condition of the transmission assets).

2) Reactive Operator Actions: Given the updated system topology ($z_{i,k+1}^{o}$ and $z_{t,k+1}^{o}$), existing system state (s_{k}^{p}), and partial generator control setpoints from the controller at k^{th} time step, the operator also utilizes the demand forecast to derive the complete list of setpoints of the generators and load shedding schedule while aiming to minimize the value of the lost load at $k+1^{th}$ time step. As indicated earlier, contrary to traditional economics-driven power system operation, we consider that the operator's objective should be to minimize the value of the shredded load at $k+1^{th}$ time step[2], and such an objective will remain in effect until the emergency condition is lifted. The use of such an objective is often notable in the post-blackout power distribution system restoration problem [28]. The consequent problem to be solved is given as:

$$\min_{\Phi_{k+1}} \sum_{i \in N} w_i^{c\_l} \Delta P_{i,k+1}^{c\_l} + w_i^{nc\_l} \Delta P_{i,k+1}^{nc\_l} + \epsilon \left(P_{i,k+1}^{g} - P_{i,k}^{g} \right)^2 \tag{5}$$

where Φ_{k+1} consists of the set of decision variables of the operator, that includes power generation outputs $P_{i,k+1}^{g}$ (without any loss of generality, all the generators at a given bus are aggregated), operational state of generators group $z_{i,k+1}^{g}$, critical, non-critical loads shedded ($\Delta P_{i,k+1}^{c\_l}$, $\Delta P_{i,k+1}^{nc\_l}$), nodal angles $\theta_{i,k+1}$, of node $i \in N$ along with line flows, $P_{t,k+1}^{flow}$, through branch $t \in T$. Also, $w_i^{c\_l}$, $w_i^{nc\_l}$ (> 0) are the values of critical and non-critical loads respectively. The criticality of the loads (hospitals, fire stations, police stations, etc. are generally treated as critical loads) imposes the condition of $w_i^{c\_l} \geq w_i^{nc\_l} > 0$ ($\forall i \in N$). These values can be differ-

ent for different nodes. The nullity of the decision variable space Φ_{k+1} indicates non-existence of the solution, implying, possible overloading and outage of transmission resources along with secondary wildfire origin, and is captured through variable q_{k+1}. Overall state of the power system is captured by $s_{k+1}^{p} = \{\Phi_{k+1} \cup q_{k+1} \cup z_{t,k+1}^{o} \cup z_{i,k+1}^{o}\}$.

Sole utilization of the value of load loss expression as an objective, $\sum_{i \in N} w_i^{c\_l} \Delta P_{i,k+1}^{c\_l} + w_i^{nc\_l} \Delta P_{i,k+1}^{nc\_l}$, may lead to the existence of multiple feasible solutions with undesirable ramping of the generators. Addition of an expression $\sum_{i \in N} \left(P_{i,k+1}^{g} - P_{i,k}^{g} \right)^2$ in the objective with minuscule positive bias of ϵ inhibits such possibility, with little to no impact on the original objective. Furthermore, $\min w_i^{nc\_l} \gg \epsilon > 0$. Determination of these weights are beyond the scope of this paper, and is assumed to be known *apriori*.

The discussed objective function must be subject to the following power system operation and safety constraints:

a) Generation Constraints: Finite generation capacity (lower and upper limits of generators given by P_i^{min}, P_i^{max}, respectively) and ramping capabilities (given by R_i^{max}) limit the operability of the generators. Since the network topology has already been revised, we use $z_{i,k+1}^{o}$, and $z_{t,k+1}^{o}$ to determine nodal and line availability, respectively. The controller provides the list of 'to be controlled' generators through $v_{i,k}^{e}$ ($= 1$ symbolizes to be controlled, $= 0$ otherwise), and their incremental setpoint adjustment, $\Delta P_{i,k}^{e}$. The operator, in return, checks the feasibility in terms of upper limits of the received set-point based on the current operating state of the generators $z_{i,k}^{g}$ using (6) ($= 1$ symbolizes operational, $= 0$ otherwise), failing which the operator reports nullity of the decision space q_{k+1}.

$$0 \leq v_{i,k}^{e} \left(P_{i,k}^{g} + \Delta P_{i,k}^{e} \right) \leq v_{i,k}^{e} z_{i,k+1}^{o} z_{i,k}^{g} \mathrm{P}_i^{max} \tag{6}$$

If the lower limit is violated (as indicated by $z_{i,k}^{e} = 0$ in (7)), the limits remain unchanged; otherwise, minimum and maximum operating limits of the generators ($\mathrm{P}_i^{min,*}, \mathrm{P}_i^{max,*}$) are updated using (8)-(9).

$$z_{i,k}^{e} = \begin{cases} 0 & \text{if } v_{i,k}^{e} \left(P_{i,k}^{g} + \Delta P_{i,k}^{e} \right) \leq v_{i,k}^{e} z_{i,k+1}^{o} z_{i,k}^{g} \mathrm{P}_i^{min} \\ 1 & \text{otherwise} \end{cases} \tag{7}$$

$$\mathrm{P}_i^{min,*} = z_{i,k}^{e} \left(P_{i,k}^{g} + \Delta P_{i,k}^{e} \right) + (1 - z_{i,k}^{e}) \mathrm{P}_i^{min} \tag{8}$$

$$\mathrm{P}_i^{max,*} = z_{i,k}^{e} \left(P_{i,k}^{g} + \Delta P_{i,k}^{e} \right) + (1 - z_{i,k}^{e}) \mathrm{P}_i^{max} \tag{9}$$

In a continuously changing environment, generators can suffer from forced outages. Here, we assume that the generators are equipped with load rejection capabilities [29], and when generators face forced outages, their ramping rate and operating limit can be allowed to contravene (see (11)), subject to the outaged generators will not be brought online without a thorough safety check. For this paper, these generators will remain outaged indefinitely (see (10)). To model such a condition, we seek the help of a large positive real constant, Γ^0, as shown in (12). These conditions, along with revised generating and ramping capability, are given in the following

---

[1]The operator may choose to not strictly follow the decision support provided by the controller. However, such an operator needs to be appropriately modeled to ensure controller robustness.

[2]The operator may also use a multi-period greedy algorithm to derive and deploy set points for next time step.

equations. Here, Δk represents the power system operating interval.

$$0 \leq z^g_{i,k+1} \leq z^g_{i,k} \tag{10}$$

$$z^g_{i,k+1} \mathbf{P}^{min,*}_i \leq P^g_{i,k} \leq z^g_{i,k+1} \mathbf{P}^{max,*}_i \tag{11}$$

$$-\Gamma^0 \left(1 - z^g_{i,k+1}\right) - \Delta k R^{max}_i \leq P^g_{i,k+1} - P^g_{i,k}$$
$$\leq \Delta k R^{max}_i + \Gamma^0 \left(1 - z^g_{i,k+1}\right) \tag{12}$$

b) Load Demand Constraints: The necessity of the deployment of control action for the generators at k^{th} time step to ensure load-generation balance at $k + 1^{th}$ time step requires prediction of load demand. Available historical data, as shown in Fig. 1, can facilitate such computation. However, the development of load prediction models is also beyond the scope of this paper and is assumed to be given.

Limited availability of generation during the prevailing contingencies necessitates demand curtailment. Suppose, $\mathbf{P}^l_{i,k+1}$ is the operator predicted load demand, then, following removal of associated substation, updated load demand will be $z^o_{i,k+1} \mathbf{P}^l_{i,k+1}$. Additionally, α_i is the parameter representing the critical load fraction (positive real number) served. Availability of a large number of switchable loads within the distribution network connected at the transmission substation enables treating the sheddable loads as a continuous variable [30]. Consequently, critical and non-critical sheddable loads $(\Delta P^{c\_l}_{i,k+1}, \Delta P^{nc\_l}_{i,k+1})$ is bounded as follows:

$$0 \leq \Delta P^{c\_l}_{i,k+1} \leq \alpha_i z^o_{i,k+1} \mathbf{P}^l_{i,k+1} \tag{13}$$

$$0 \leq \Delta P^{nc\_l}_{i,k+1} \leq (1 - \alpha_i) z^o_{i,k+1} \mathbf{P}^l_{i,k+1} \tag{14}$$

c) Load Flow Constraints: Since the typical operating voltage within the power system remains close to 1.00 *pu*, and the difference in the voltage angle of the adjacent buses is tiny, we consider a DC power flow model [31]. An associated mathematical expression is given in (16). Here, $\mathbf{B}_t$ is the element corresponding to the t^{th} branch in the imaginary part of the nodal admittance matrix. Equation (15) represents the nodal flow balance equation. Also, the set $T^i \subseteq T$ consists of all the branches that are connected to node i. Here, θ^{min} and θ^{max} are upper and lower bound of nodal angle, respectively. The power flow constraint is described in (18).

$$P^g_{i,k+1} - \mathbf{P}^l_{i,k+1} z^o_{i,k+1} + \Delta P^{c\_l}_{i,k+1} + \Delta P^{nc\_l}_{i,k+1}$$
$$- \sum_{t \in T^i} P^{flow}_{t,k+1} = 0 \tag{15}$$

$$P^{flow}_{t,k+1} - z^o_{t,k+1} \mathbf{B}_t \left(\theta_{i,k+1} - \theta_{j,k+1}\right) = 0 \tag{16}$$

$$\theta^{min} \leq \theta_{i,k+1} \leq \theta^{max} \tag{17}$$

$$-z^o_{t,k+1} \mathbf{P}^{maxflow}_t \leq P^{flow}_{t,k+1} \leq z^o_{t,k+1} \mathbf{P}^{maxflow}_t \tag{18}$$

3) Power Flow Analysis: At each time step, the deenergization actions, along with calculated setpoints determined earlier, are deployed, and the operator waits for the duration of Δk. It is notable that the delay in the measurement of the state and deployment of the control action is minuscule enough to account for. It is also assumed that the availability of sufficient load-following reserve can help mitigate intra-period slower fluctuation [32]. Since the forecasted system operating

conditions (load demands[3]) may not materialize, we need to solve an actual set of power flow equations[4] to determine the system state. As shown in the online deployment part of Fig. 1, the correct system state can be directly obtained from the real environment. This revised state is fed to the controller to determine requisite operator action for the next time step.

C. Reward Function and MDP Formulation

In our proposed model, the power system operator is myopic. As a result, the controller needs to consider future trajectory and provide an appropriate control signal to the operator, facilitating prevention from running into reliability related issues while maximizing the value of load served. Given the probabilistic nature of wildfire propagation and power system loads, we formulate the control problem solved by an agent as a Markov decision process (MDP) problem. An MDP is a tuple $\mathcal{D} = \langle S, A, \mathcal{P}, \mathcal{R} \rangle$, where S is a finite state space, A is a finite action space, $\mathcal{P}$ is the transition probability function and $\mathcal{R}$ is the reward function. The agent chooses an action from the possible action space to lead the system from one state to another. For the given problem, these elements are defined as follows:

States: At any time step k, the state $s^{\mathcal{E}}_k$ of $\mathcal{D}$ is defined by a tuple $s^{\mathcal{E}}_k = \langle s^f_k, s^p_k \rangle$, including the states of wildfire and power system respectively. The complete state space S is defined in (19).

$$S \subseteq \underbrace{\{0,1\}^M \times \mathbb{R}^M}_{S^f \text{ (wildfire)}} \times \underbrace{\{0,1\}^{|T|+|N|+1} \times \mathbb{R}^{|T|+3|N|}}_{S^p \text{ (power system)}} \tag{19}$$

As discussed in Section II-A1, M is the number of cells for depicting the topographical map for wildfire propagation, $|N|$ is the number of substations, and $|T|$ is the numbers of transmission lines. We assume that the state of the environment is entirely observable to the controller (i.e., to the agent).

Actions: The agent can provide emergency shut off commands to the substations and transmission lines ($z^e_{i,k}$, $z^e_{t,k}$, respectively), and change power injections for the generators v^e_i, by an amount ΔP^e_i. The action space A is defined as $A \subseteq \{0,1\}^{2|N|+|T|} \times \mathbb{R}^{|N|}$.

Transitions: An MDP evolves as a result of the set of actions taken. The transition probability function, $\mathcal{P}(s_{k+1}|s_k, a_k)$ indicates that the action a_k at time step k in state s_k will lead to the next state s_{k+1}. Here, the stochasticity arises from the wildfire propagation model and variation in the load demand (along with generation uncertainty with renewable energy resources) within the power system.

Reward Function: The reward function, $r_k = \mathcal{R}(s_k, a_k, s_{k+1})$, indicates that the reward is received for taking action a_k in state s_k to reach next state s_{k+1}. The reward function is formally defined by (20), where c_1, c_2, c_3 and c_4 are positive, constant weights of each types of penalty. Determining the optimal set of these weights is beyond the scope of this paper but is expected to have a major impact on the training speed of the agent.

---

[3]Renewable generation uncertainties can also be treated similarly.
[4]AC power flow equations can be invoked here.

$$\mathcal{R}(s_k, a_k, s_{k+1}) = -1 \times (c_1 r_1 + c_2 r_2 + c_3 r_3 + c_4 r_4) \quad (20)$$

$$r_1 = q_k \sum_{i \in N} \left((1 - z_{i,k}^o) \mathrm{P}_{i,k}^l + \Delta P_{i,k}^{c\_l} + \Delta P_{i,k}^{nc\_l} \right) \quad (21)$$

$$r_2 = q_k \left(\sum_{i \in N} \psi_{i,k}^e |P_{i,k}^g| + \sum_{t \in T} \psi_{t,k}^e |P_{t,k}^{flow}| \right) \quad (22)$$

$$r_3 = q_k \left(\sum_{i \in N} \psi_{i,k+1}^f + \sum_{t \in T} \psi_{t,k+1}^f \right) \sum_{i \in N} \mathrm{P}_{i,k}^l \quad (23)$$

$$r_4 = (|N \cup T|)(1 - q_k) \sum_{i \in N} \mathrm{P}_{i,k}^l \quad (24)$$

We defined the *load loss* penalty r_1 in (21), where the expression $(1 - z_{i,k}^o)\mathrm{P}_{i,k}^l$ signifies the shedded load demand following isolation of a substation, and $\Delta P_{i,k}^{c\_l} + \Delta P_{i,k}^{nc\_l}$ denotes aggregated critical and non-critical load curtailment of substation i.

The *Proactive Isolation of grid Assets with expected Wildfire (PIAW)* penalty r_2 is proportional to the amount of power flowing through the transmission lines $P_{t,k}^{flow}$ or the substations $P_{i,k}^g$ during the isolation of live assets as defined in (22), where $\psi_{j,k}^e = (z_{j,k}^o \oplus z_{j,k}^e)$ denotes the change in the operational status of the equipment $j \in \{N \cup T\}$ as a result of controller action. Here, $\psi_{j,k}^e$ evaluates to true if asset j status changes from k^{th} time step to $k+1^{th}$ time step due to the agent action at time step k. This penalty stems from the fact that isolation of heavily loaded assets has comparably higher chance of creating power system reliability related problems.

The *Asset Damage and Isolation due to encroached Wildfire (ADIW)* penalty r_3 is defined in (23), where the agent is penalized in proportion to aggregated load demand $\sum_{i \in N} \mathrm{P}_{i,k+1}^l$. Here, $\psi_{j,k}^f = (z_{j,k}^o \oplus z_{j,k+1}^f)$ denotes the change in the operational status of the equipment $j \in \{N \cup T\}$ as a result of safety related action following encroached wildfire. This penalty is to prohibit safety related actions after assets are encroached by wildfire.

The *non-convergence* penalty r_4, defined in (24), is equal to the aggregated network-wide load $\sum_{i \in N} \mathrm{P}_{i,k+1}^l$ times the number of assets $|N \cup T|$ in the system. It is notable that while non-convergence may not result into cascading outage, we use this term as a negative feedback to the controller to inhibit thermal overloading. Detailed model of slow cascading will be considered in our future research.

Policy: A policy $\mu(s)$ is a function that specifies the action to be taken in each state s. At every time step k, the agent selects an action, $a_k = \mu(s_k)$, based on the deployed policy. This experiment aims to find an optimal policy that can maximize the cumulative sum of expected rewards in each episode. The expected rewards in the model are defined in (25), where n is the finite number of steps in each episode. We consider the power system operator has no inherent temporal preference as the episodes' length is short (system operates in resiliency mode only for a relatively short period of time).

$$\mu^* = \arg\max_\mu \mathbb{E} \left[\sum_{k=0}^{n} \mathcal{R}(s_k, a_k, s_{k+1}) \right] \quad (25)$$

III. DEEP REINFORCEMENT LEARNING BASED PROACTIVE CONTROL

We apply deep reinforcement learning to train the agent. We consider a standard reinforcement learning approach, where the agent interacts with its environment in discrete time steps k, as illustrated in Fig. 1. The agent learns by iteratively updating its policy μ^*, defined in (26), where $\gamma \in (0, 1)$ is a temporal discount factor for infinite-horizon future rewards:

$$\mu^* = \arg\max_\mu \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k \cdot \mathcal{R}(s_k, a_k, s_{k+1}) \right] \quad (26)$$

Our control problem, modeled as an MDP, is fully observable, so the state S and observation spaces O are equivalent. However, due to the enormous observation O and action spaces A defined in Section II-C, training our agent can be difficult and time consuming. To improve the agent's learning pace, we propose switching to a low-dimensional, "compressed" observation space $\hat{O}$. Since the removal of an important feature to compress the observation space may also impair the agent's performance, we need to choose the compressed space $\hat{O}$ carefully. To reduce the action space, we propose a hybrid approach, where control actions are partially chosen by a heuristic, inspired by [15], and by a neural network.

A. Reducing the Observation Space

As described in Section II-C, the complete observation space is $O \subseteq \{0, 1\}^{M+|T|+|N|+1} \times \mathbb{R}^{M+|T|+3|N|}$. In practice, the number of cells M in the wildfire propagation model is far greater than the number of substations N and transmission lines T (i.e., $M \gg |N \cup T|$). So, we can reduce the size of the observation space O by reducing only the dimension of wildfire observation $S^f \subseteq \{0, 1\}^M \times \mathbb{R}^M$ without significantly affecting the agent's accuracy.

To this end, we propose to replace cell states in the observation with a "fire-distance metric" for each asset: for each power system asset in $N \cup T$ (substation or transmission line), calculate and observe the geographical distance from the nearest ignited cell. This transformation reduces the wildfire observation space from S^f to $\hat{S}^f \subseteq \mathbb{R}^{|T|+|N|}$. The reduced observation space $\hat{O}$ of the agent is given in (27):

$$\hat{O} \subseteq \underbrace{\mathbb{R}^{|T|+|N|}}_{\hat{S}^f \, (\text{wildfire})} \times \underbrace{\{0, 1\}^{|T|+|N|+1} \times \mathbb{R}^{|T|+3|N|}}_{S^p \, (\text{power system})} \quad (27)$$

B. Reducing the Action Space

The impact of large action space, $\{0, 1\}^{2|N|+|T|} \times \mathbb{R}^{|N|}$, can be significantly reduced through direct utilization of fire-distance metric in the decision making through a simple heuristic. This heuristic can control the energization status of the component of the power system. At the same time, the neural network-trained RL setting can determine the generator's setpoints. Such a hybrid approach can alleviate the monolithic agent's bottleneck while emphasizing learning a relatively good power generation control policy. However, the action space for the RL agent ($\hat{A} \subseteq \{0, 1\}^{|N|} \times \mathbb{R}^{|N|}$) is still large and we use two strategies to further scale down the action space.

6

- The first strategy, referred to as *full-control*, forces the agent to control all generator outputs. Here, the external control input $v_{i,k}^e$ is defined per (28), and the resulting action space reduces to $\hat{A}_1 \subseteq \mathbb{R}^{|N^{gen}|}$, where $N^{gen} \subseteq N$ is the set of nodes with generation capabilities.

$$v_{i,k}^e = \begin{cases} 1 & \text{if} \quad i \in N^{gen} \\ 0 & \text{otherwise.} \end{cases} \quad (28)$$

- The second strategy, referred to as *partial-control*, allows the agent to control exactly g number of generators. Due to the nature of the power transmission network, mostly the generators that are in the vicinity of the wildfire need to be actively controlled. Therefore, we can select g generators by dividing the set of generating nodes N^{gen} into g number of subsets, and select one generator from each subset.

The heuristic and DRL training algorithm are described in following subsections.

1) Heuristics: The heuristic rule de-energizes power system assets proactively based on the fire-distance metric of the reduced observation space $\hat{S}^f$. If the distance between the fire and an asset $j \in \{N \cup T\}$ is less than or equal to a pre-defined threshold β, then asset j is de-energized, as defined in (29).

$$z_{j,k}^e = \begin{cases} 0 & \text{if} \quad \hat{s}_k^j \leq \beta \\ 1 & \text{otherwise} \end{cases} \quad (29)$$

Ideally, the value of β depends upon the thermal characteristics of the equipment and the rate of temperature increase in the vicinity of the equipment's physical location. However, in this work, since the stochastic wildfire propagation model does not explicitly capture temperature change, we use a constant value $\beta = 2$ (i.e., if the fire is 2 cells away from an asset j, then the corresponding external control input $z_{j,k}^e$ is set to 0).

2) DRL Training Algorithm: To solve the problem described in Section II-C, we closely follow the Deep Deterministic Policy Gradient (DDPG) [33] algorithm for the agent. The algorithm is a model-free, off-policy actor-critic algorithm using deep function approximators that can learn policies in high-dimensional, continuous action spaces.

The agent takes an action a_k at every time step k based on its trained policy μ, $a_k = \mu(s_k|\theta^\mu)$, where θ^μ is the weights of the agent (actor) for policy μ. During training, the agent adds noise with the action it takes to explore the network. The actor updates its weights θ^μ based on the critic value $Q(s, a|\theta^Q)$. The agent also uses target actor μ' and critic Q' networks for stability. The target networks are initialized with the same random weights that $\theta^{Q'} \leftarrow \theta^Q$ and $\theta^{\mu'} \leftarrow \theta^\mu$.

The agent first calculates the expected return, $y_k = r_k + \gamma Q'(s_{k+1}, \mu'(s_{k+1}|\theta^{\mu'})|\theta^{Q'})$, based on the target networks, where r_k is the reward value returned from the environment calculated based on (20) and γ is the discount factor for future rewards. Here, y_k is called the moving target that the critic model tries to achieve. The agent updates the critic network θ^Q by minimizing the loss, $L = \frac{1}{N} \sum_k (y_k - Q(s_k, a_k|\theta^Q))^2$.

The actor network θ^μ is updated based on the critic value. To calculate the actor loss, the agent first selects the action a_k based on the actor policy μ and then applies the action to the critic network. The agent updates the actor policy by applying the sampled policy gradient, $\nabla_{\theta_\mu} J \approx \frac{1}{N} \sum_k \nabla_a Q(s, a|\theta^Q)|_{s=s_k, a=\mu(s_k)} \nabla_{\theta_\mu} \mu(s|\theta^\mu)|_{s_k}$.

Finally, the agent updates the target actor $\theta^{\mu'} = \tau\theta^\mu + (1 - \tau)\theta^{\mu'}$ and target critic $\theta^{Q'} = \tau\theta^Q + (1-\tau)\theta^{Q'}$ networks, where $\tau \in (0, 1)$ is the target-network update constant. To stabilize the learning process, constant τ is chosen to be a very small value, which prevents divergence due to fast updates.

IV. SIMULATION RESULTS

A. Simulation Setup

A standard IEEE 24-bus reliability test system (RTS), superimposed on a hypothetical geospatial terrain, divided into 350×350 grid, has been considered here for analysis. Power system operational parameters can be obtained from [34]. As discussed earlier, to reduce the size of the solution space, the outputs of these generators are aggregated. Entire power system control horizon is divided into several time steps with a 15-minute interval. We consider the time step for the fire propagation is six times faster than power system operation. Here, every episode consists of 300 time steps (spanning the entire power system control horizon) or until we reach the non-convergence condition.

For simplicity, we also consider that the load demand within the power network is deterministic in nature with constant magnitude. As discussed, network-wide loads are comprised of critical and non-critical fractions. This fraction also remains constant throughout the network for simplicity. Our proposed controller tracks spatio-temporal wildfire propagation and provides set points for a partial set of generators. Successively, the power system operator would solve the optimization problem to calculate and deploy the complete set of requisite control actions based on load forecasts. To determine the setpoints for the operator, we utilize the SCIP solver in *General Algebraic Modeling System (GAMS)* due to its versatility. As discussed, non-convergence of power system operational problem is treated as a *simulation-ending* condition with a severe penalty. To calculate the total penalty (see (20)), constants c_1, c_2, c_3, and c_4 (corresponding to value of lost load, PIAW, ADIW, and non-convergence conditions) are 1, 1, 2, and 10, respectively.

We conduct our initial experiments considering a given wildfire origin, which remains the same during training and evaluation, and we refer to this as a *fixed source*. Successively, we expand our experiments by considering a random wildfire origin within the geographical area, which is chosen at random for each episode, and we refer to this as a *random source*. We set a square boundary of 250×250 cells inside the 350×350 grid for the *random source* of wildfire to ensure that the wildfire impacts the power system assets.

We use a deep neural network with two hidden layers, consisting of 450 and 300 neurons respectively, for both the actor and the critic network of the agent. In the actor network, we use *tanh* activation function in the hidden layers and *sigmoid* activation in the output layer. In the critic network, we use *rectified linear* activation in the hidden layers and *linear* activation in the output layer. The actor and critic learning rates are 0.001 and 0.002, respectively. We set 0.01 for the constant

τ to update the target network and 0.99 as the discount factor γ to calculate the expected return described in Section III-B2.

B. Control Approaches

We consider three different control approaches in this experiment, which are defined as follows:

1) Reactive control: In this approach, the system operator determines appropriate asset de-energization sequence, load shedding, and generator setpoints to be deployed based on its current observation of the interaction of wildfire with the power system (monitored transmission line status and emergency calls from substations) and predicted system-wide load demand. Here, the operator is blind to the wildfire propagation in the topographical space. External decision support is unavailable here.

2) Myopic control: The external decision-making agent, based on its observation of wildfire propagation over the topographical space, provides the operator with proactive power system asset de-energization control input. The operator also observes the impact of wildfire on the power system, estimates system wide-load demand, and considers de-energization decision support from the controller to determine the control input for the next time step.

3) Proactive control: In proactive control, the controller provides the operator with (i) power system asset de-energization control input based on the heuristic defined in Section III-B1, and (ii) set points of the generators utilizing an RL-based controller. The RL-based controller can provide setpoints for a partial set or, entire generator fleet to the operator, and corresponding controllers are identified as full-control and partial-control, as described in Section III-B. The operator is also expected to observe the impact of wildfire on the power system, estimate system wide-load demand, and account for external control input, to calculate and deploy the requisite control actions.

C. Training the Proactive Control Agent

Training of DRL-based controllers can often be time-consuming, and it may take several days for an agent to learn a satisfactory policy. The pace of learning can also vary based on the random seed and selected hyper-parameter values.

During training, the agent explores the solution space by adding random noise to the actions chosen by the actor. Therefore, penalties received during training cannot indicate the actual performance of the agent. Further, overfitting may diminish the agent's performance. In light of this, we evaluated the agent for 4 episodes (without exploration) after every 20 episodes of training to find the best version of the agent. After training has finished, we chose a set of trained agents based on these initial evaluations. Finally, we evaluated them thoroughly over 100 episodes. We trained agents in this manner for both the fixed- and the random-source environments.

D. Results

1) Fixed fire source: As shown in Fig. 2, the operator encounters no simulation-ending condition with the proactive control approach, while the reactive and myopic control
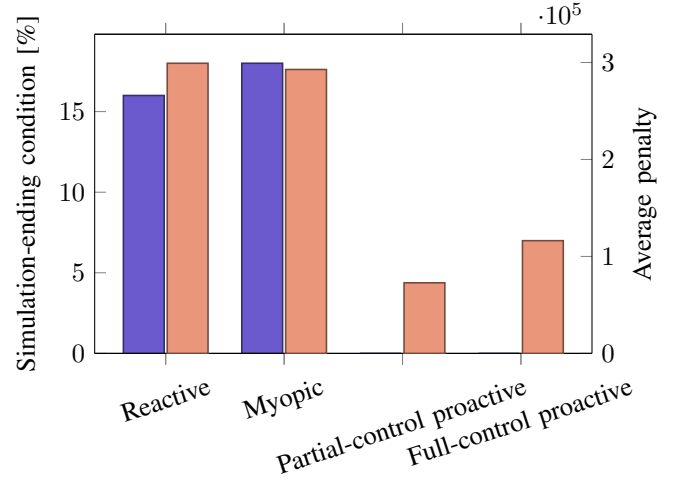

Fig. 2: Ratio of episodes that encountered a simulation-ending condition (■) and average total penalty (■) for various control approaches (reactive, myopic, partial-control proactive, and full-control proactive) with fixed fire source.
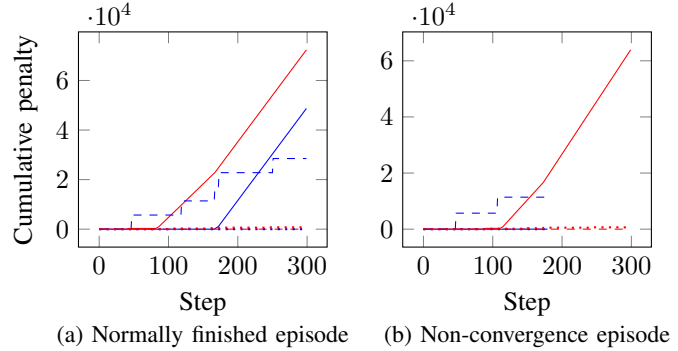

(a) Normally finished episode   (b) Non-convergence episode

Fig. 3: Cumulative penalty components for reactive (load loss (——), ADIW (– – –) and PIAW (⋯⋯)) and proactive (load loss (——), ADIW (– – –) and PIAW (⋯⋯)) control approaches over each step of two example episodes.

approaches encounter simulation-ending states in 16% and 18% of episodes, respectively. The reason for encountering simulation-ending states is that the operator does not consider the entire time horizon when switching transmission assets and providing set points to the generators with these two approaches. This problem is addressed by the proactive control approach, as the controller is inherently trained to consider possible future states and penalties. In Fig. 2, we can also see that the average penalty is much lower for the proactive control agent (measured over 100 episodes).

Figs. 3a and 3b show temporal variation of various penalty component (without *non-convergence*) for given wildfire origin and propagation paths. Both Figs. 3a and 3b shows that although the load loss penalty in the proactive control approach is higher compared to the reactive control approach, this is compensated by other components of the reward function, indicating reduced flow through the 'to be' de-energized lines or alleviation of non-convergence. It is important to note that load loss penalty in the proactive control approach need not
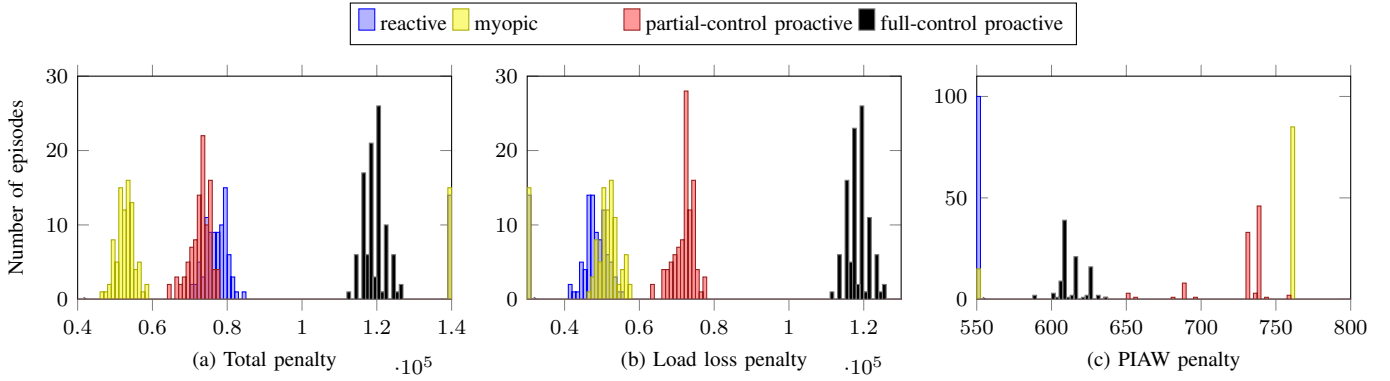
8

Fig. 4: Distribution of total penalty over 100 episodes for different types of controls (reactive, myopic, partial-control proactive, and full-control proactive) with a fixed fire source.
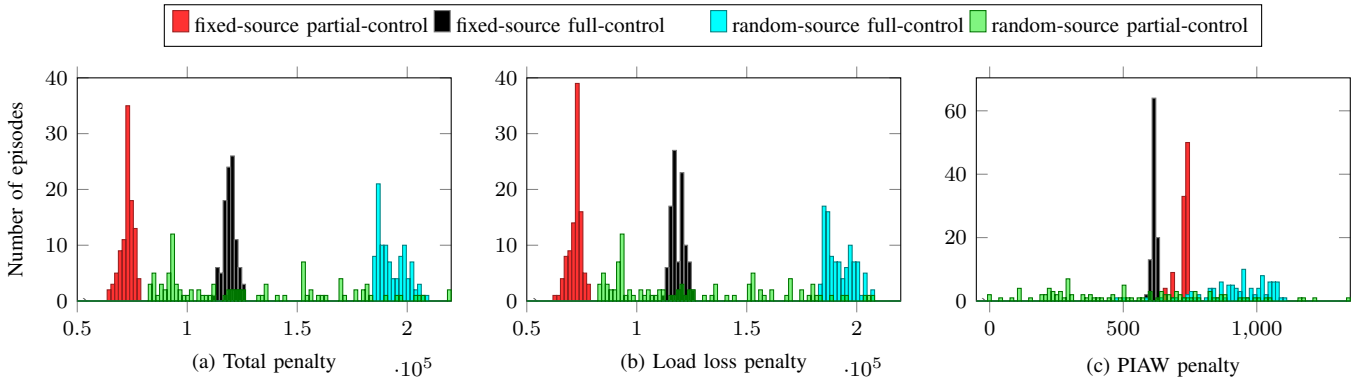

Fig. 5: Distribution of penalty over 100 episodes for various types of proactive control with fixed and random fire sources.

always be higher compared to the reactive control approach. Fig. 3(b) shows that the ability to estimate the future and accounting for it in the decision making, even with the myopic operator, ensures that the system operation does not end prematurely due to non-convergence.

In Figure 4a, we see the distribution of episodic *total penalty* for different types of control approaches. Since the total load within the system remains constant, the non-convergence penalty, if encountered, will also be constant. Therefore, we did not include the distribution of the non-convergence penalty, and Fig. 2 can be referred to in this regard. It can be seen from the figures that the median of the total penalty for the proactive partial-control agent falls in between the reactive control approach and the myopic control agent if non-convergence penalties are ignored. The ability to avoid non-convergence with a proactive control agent indicates the superiority of this approach. It can also be seen that partial control works better than controlling the entire generating fleet.

Figs. 4b and 4c show the load loss and the PIAW penalty, respectively. From the earlier discussion, it is evident that the proactive control agent tries to find equilibrium by reducing the power flow through the assets, minimizing ADIW and PIAW and avoiding non-convergence, while also minimizing load loss penalty. Note that the reactive control approach does not have a PIAW penalty as it does not remove lines proactively. Instead, it incurs an ADIW penalty for not taking any action.

2) Proactive control agents based on random fire sources: We also trained and evaluated agents considering random fire sources, as shown in Fig. 5a. Here, both proactive control agents can again learn to avoid simulation-ending conditions. Although the random-source partial-control agent seems better than the random-source full-control agent, the agent encountered simulation-ending conditions in 2% of the cases.

Note that in some of the figures, we see distribution columns at the edge of the figures. These columns represent all the values that are outside of the plotted range. In Figure 4a, we see a yellow column (myopic control agent) at the right edge. Those are the total penalties that include a non-convergence penalty for reaching simulation-ending conditions. Similar phenomenon can also be observed in Fig. 5a.

V. Conclusions

In this work, we developed a deep reinforcement learning (DRL) based proactive control to supplement decision support for operators given a wildfire event. Testbed has been developed integrating a wildfire-propagation model with a power-system operation model to train and validate a controller that can supplement traditional computationally-intensive, forecast-driven power-system operations during a wildfire. We formulated the control problem as a Markov Decision Process (MDP). To address the challenges posed by the large observation space, we introduced a compact representation of fire-state observations. To tackle the challenges of the large action

space, we proposed a hybrid agent, consisting of a heuristic and a deep neural network, which facilitate the training of the controller. In particular, we utilized a model-free, off-policy, actor-critic algorithm using deep function approximators as a part of the Deep Deterministic Policy Gradient (DDPG) algorithm to train a deep RL-based policy. We also aggregated generator outputs to reduce the size of the action space.

Numerical results indicate that the DRL-based proactive control agent can avoid thermal overloading of the transmission network when one or more transmission assets are outaged due to safety concerns with encroaching wildfire, while reactive and myopic approaches struggle to do so. Since our agent needs to be trained before the wildfire event, a crucial question is whether it can perform well if the origin of the fire is unknown at training time: our results demonstrate that while it is easier to handle fixed-source fires, our agent performs well even with random fire sources. We also observed that our proactive control agent performs significantly better when it provides setpoints for only a subset of generators and lets the operator derive setpoints for the remaining ones (compared to providing setpoints to all generators).

In this work, we reduced the size of the observation space by utilizing heuristics for the sake of computational tractability, which enables the application of our approach to large geographical areas. However, actions chosen based on such a reduced space might be sub-optimal. In future work, we will explore approaches for handling large observation spaces, such as (graph) convolutional neural networks.

REFERENCES

[1] A. B. Smith and R. W. Katz, "US billion-dollar weather and climate disasters: data sources, trends, accuracy and biases," *Natural Hazards*, vol. 67, no. 2, pp. 387–410, 2013.

[2] M. Choobineh, B. Ansari, and S. Mohagheghi, "Vulnerability assessment of the power grid against progressing wildfires," *Fire Safety Journal*, vol. 73, pp. 20 – 28, 2015.

[3] Western Area Power Lines, "Trees and power lines." [Online]. Available: https://www.wapa.gov/newsroom/FactSheets/Pages/trees-powerlines.aspx

[4] J. T. Abatzoglou, C. M. Smith, D. L. Swain, T. Ptak, and C. A. Kolden, "Population exposure to pre-emptive de-energization aimed at averting wildfires in Northern California," *Environmental Research Letters*, vol. 15, no. 9, p. 094046, 2020.

[5] Y. Wang, C. Chen, J. Wang, and R. Baldick, "Research on resilience of power systems under natural disasters–A review," *IEEE Trans. Power Syst.*, vol. 31, no. 2, pp. 1604–1613, 2016.

[6] J. Latson, "Why the 1977 blackout was one of New York's darkest hours," *Time*, 2015.

[7] M. A. Mohamed, T. Chen, W. Su, and T. Jin, "Proactive resilience of power systems against natural disasters: A literature review," *IEEE Access*, vol. 7, pp. 163 778–163 795, 2019.

[8] C. Wang, Y. Hou, F. Qiu, S. Lei, and K. Liu, "Resilience enhancement with sequentially proactive operation strategies," *IEEE Trans. Power Syst.*, vol. 32, no. 4, pp. 2847–2857, 2016.

[9] S. Biswas, E. Bernabeu, and D. Picarelli, "Proactive islanding of the power grid to mitigate high-impact low-frequency events," in *2020 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*. IEEE, 2020, pp. 1–5.

[10] J. Q. Tortós and V. Terzija, "Controlled islanding strategy considering power system restoration constraints," in *2012 IEEE Power and Energy Society General Meeting*, 2012, pp. 1–8.

[11] P. Dehghanian, S. Aslan, and P. Dehghanian, "Quantifying power system resiliency improvement using network reconfiguration," in *2017 IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS)*, 2017, pp. 1364–1367.

[12] K. P. Schneider, F. K. Tuffner, M. A. Elizondo, C. Liu, Y. Xu, S. Backhaus, and D. Ton, "Enabling resiliency operations across multiple microgrids with grid friendly appliance controllers," *IEEE Trans. Smart Grid*, vol. 9, no. 5, pp. 4755–4764, 2018.

[13] C. Chen, J. Wang, F. Qiu, and D. Zhao, "Resilient distribution system by microgrids formation after natural disasters," *IEEE Trans. Smart Grid*, vol. 7, no. 2, pp. 958–966, 2016.

[14] M. Panteli, D. N. Trakas, P. Mancarella, and N. D. Hatziargyriou, "Boosting the power grid resilience to extreme weather events using defensive islanding," *IEEE Transactions on Smart Grid*, vol. 7, no. 6, pp. 2913–2922, 2016.

[15] N. Rhodes, L. Ntaimo, and L. Roald, "Balancing wildfire risk and power outages through optimized power shut-offs," *IEEE Transactions on Power Systems*, 2020.

[16] T. Wang, Z. Tang, X. Wang, J. Yang, X. Wu, K. Zhao, C. Zuo, and W. Chen, "Whether the high-voltage transmission lines have enough load capacity after wildfire," in *5th International Conference on Electrical Engineering and Automatic Control*, 2016, pp. 1243–1251.

[17] H. Liu, R. A. Davidson, D. V. Rosowsky, and J. R. Stedinger, "Negative binomial regression of electric power outages in hurricanes," *Journal of Infrastructure Systems*, vol. 11, no. 4, pp. 258–267, 2005.

[18] S. D. Guikema, R. Nateghi, S. M. Quiring, A. Staid, A. C. Reilly, and M. Gao, "Predicting hurricane power outages to support storm response planning," *IEEE Access*, vol. 2, pp. 1364–1373, 2014.

[19] B. Storrow, "Why the deep freeze caused Texas to lose power," *Scientific American*, February 2021. [Online]. Available: https://www.scientificamerican.com/article/why-the-deep-freeze-caused-texas-to-lose-power/

[20] R. Nateghi, S. D. Guikema, and S. M. Quiring, "Forecasting hurricane-induced power outage durations," *Natural Hazards*, vol. 74, no. 3, pp. 1795–1811, 2014.

[21] C. Ji, Y. Wei, and H. V. Poor, "Resilience of energy infrastructure and services: Modeling, data analytics, and metrics," *Proceedings of the IEEE*, vol. 105, no. 7, pp. 1354–1366, 2017.

[22] S. Singla, T. Diao, A. Mukhopadhyay, A. Eldawy, R. Shachter, and M. Kochenderfer, "WildfireDB: A spatio-temporal dataset combining wildfire occurrence with relevant covariates," in *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*, 2020.

[23] X. Chen, G. Qu, Y. Tang, S. Low, and N. Li, "Reinforcement learning for decision-making and control in power systems: Tutorial, review, and vision," *arXiv preprint arXiv:2102.01168*, 2021.

[24] P. P. Khargonekar and M. A. Dahleh, "Advancing systems and control research in the era of ML and AI," *Annual Reviews in Control*, vol. 45, pp. 1 – 4, 2018.

[25] D. Bertsimas, J. D. Griffith, V. Gupta, M. J. Kochenderfer, and V. V. Mišić, "A comparison of Monte Carlo tree search and rolling horizon optimization for large-scale dynamic resource allocation problems," *European Journal of Operational Research*, vol. 263, no. 2, pp. 664–678, 12 2017.

[26] D. Boychuk, W. J. Braun, R. J. Kulperger, Z. L. Krougly, and D. A. Stanford, "A stochastic forest fire growth model," *Environmental and Ecological Statistics*, vol. 16, no. 2, pp. 133–151, 2009.

[27] *Line Thermal Monitoring —- dynamic rating of transmission lines*, ABB Switzerland Ltd, 2005, rev. 10.05. [Online]. Available: https://library.e.abb.com/public/eace83bd6a60d884c12570d0002f99b4/1002_LTM_PSGuard_Datasheet.pdf

[28] S. Poudel and A. Dubey, "Critical load restoration using distributed energy resources for resilient power distribution system," *IEEE Trans. Power Syst.*, vol. 34, no. 1, pp. 52–63, 2019.

[29] W. P. Gorzegno and P. V. Guido, "Load rejection capability for large steam generators," *IEEE Trans. Power Apparatus and Systems*, vol. PAS-102, no. 3, pp. 548–557, 1983.

[30] M. Esfahani, N. Amjady, B. Bagheri, and N. D. Hatziargyriou, "Robust resiliency-oriented operation of active distribution networks considering windstorms," *IEEE Transactions on Power Systems*, vol. 35, no. 5, pp. 3481–3493, 2020.

[31] W. D. Stevenson, "Element of power system analysis," *McGraw-Hill*, 1975.

[32] D. Kirschen and G. Strbac, *System Security and Ancillary Services*. John Wiley & Sons, Ltd, 2004, ch. 5, pp. 105–139.

[33] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.

[34] Probability Methods Subcommittee, "IEEE reliability test system," *IEEE Trans. Power Apparatus and Systems*, vol. PAS-98, no. 6, pp. 2047–2054, 1979.