CFD-Based Multi-Fidelity Machine Learning for Communication-Link-Aware Restoration of Cyber-Physical Distribution Systems under Wildfires

Alireza Nikzad, Xiaotian Zhou, and Hao Liang
Department of Electrical and Computer Engineering, University of Alberta, Edmonton, AB, Canada
Emails: anikzad@ualberta.ca, xzhou9@ualberta.ca, hao2@ualberta.ca

***Abstract*—Wildfires can simultaneously damage the physical power distribution assets and the supporting cyber infrastructures. While mobile emergency generators (MEGs) provide a flexible means to restore power for physically isolated grid sections, their deployment becomes significantly more complex when communication links also fail. In this paper, we propose a two-stage optimization framework that co-optimizes static MEG placements and dynamic restoration path selections to address this challenge. The optimization problem formulation integrates both physical and cyber constraints in a mixed-integer linear program that jointly determines static MEG placements and selects restoration paths, weighting their values by the communication success probability and enforcing a remote-switch concurrency cap. To generate the required hazard masks under wildfires, a computational fluid dynamics (CFD)-based multi-fidelity machine learning (MFML) model is developed. Based on this model, the optimization problem is solved monolithically, and a power flow audit is performed to ensure the feasibility of solutions. A case study using real-world digital elevation model (DEM) and fuel map data from Southern California indicates that the proposed framework improves service to critical loads in communication-constrained systems while maintaining low computational complexity.**

***Index Terms*—Computational fluid dynamics, cyber-physical distribution systems, mobile emergency generators, multi-fidelity machine learning, wildfire.**

I. Introduction

Future cyber-physical distribution systems will rely heavily on tightly coupled power and communication networks. In recent years, wildfire severity and critical infrastructure resilience are becoming increasingly intertwined as wildfires occur more frequently, particularly in regions like western North America, threatening not only physical assets but also the cyber domain that underpins power delivery. During the 2023 Northwest Territories wildfire emergency, more than 20,000 residents were evacuated as reliance on a single fiber-optic cable hampered emergency operations and communications [1]. This vulnerability resurfaced in May 2024, when wildfires affected major fiber-optic routes, isolating communities and disabling emergency calls [2]. These events underscore the need for planning tools that integrate physics-guided wildfire exposure with communication constraints when restoring distribution systems. Mobile emergency generators (MEGs) are a key emerging technology for this purpose. They offer significant benefits, such as providing flexible, local power to critical loads that are isolated from the main grid. However, realizing this potential presents new operational challenges. These include determining the optimal pre-positioning (siting) of MEGs and managing their real-time dispatch while relying on the degraded communication infrastructure that complicates remote switching and coordination [3], [4]. Fully leveraging these assets, therefore, remains an open challenge.

To address these challenges, prior studies have explored resilience enhancement via data-driven hazard modeling [5], remote-controlled switch (RCS) allocation [6], and various restoration schemes [7], [8]. However, existing models often overlook communication link reliability and typically assess restoration feasibility independently of the underlying wildfire conditions. Beyond communication constraints, accurately representing the physical wildfire hazard itself remains a major obstacle. On the other hand, wildfire modeling spans a spectrum from fast kinematic models, which offer high computational efficiency but limited accuracy, to high-fidelity computational fluid dynamics (CFD) and large-eddy simulation (LES), which capture fire–atmosphere interactions but are computationally prohibitive for operational use [9], [10]. As a result, many existing restoration schemes rely on overly simplistic or low-accuracy fire models. How to retain CFD-level fidelity in wildfire modeling for restoration decision making, while ensuring tractability for practical implementation in cyber-physical distribution systems, still requires extensive research. Several research works indicate that multi-fidelity machine learning (MFML) surrogates can fuse expensive and inexpensive models to approximate high-fidelity outcomes [11]–[14]. This is particularly appealing in wildfire applications, where CFD snapshots are costly to obtain, while kinematic features can be generated quickly. However, how to utilize such surrogates for distribution system restoration with communication link awareness is still an open issue.

In this paper, we investigate load restoration on distribution feeders during wildfires, accounting for uncertainty in both power equipment and communication link availabilities. We develop a cyber-physical system model that defines component-level wildfire hazards using high-fidelity CFD. These hazard data inform a communication-aware, two-stage mixed-integer linear programming (MILP) formulation that integrates physical availability with communication-link success probability to guide MEG siting and path activation. Finally, to ensure computational tractability, we propose a complete solution framework leveraging a CFD-based MFML surrogate to generate hazard masks, combined with a monolithic optimization and a power flow audit to validate power flow feasibility. The main contributions of this paper are threefold: 1) A communication-link-aware cyber–physical distribution-

979-8-3315-8138-1

system model is developed to integrate CFD-based wildfire hazard characterization to support MEG-assisted restoration under dynamic fire conditions;
2) A two-stage MILP formulation is proposed, which integrates the physical availability derived based on CFD models with communication-link success probability to guide restoration;
3) A computationally tractable solution framework is developed to leverage an MFML surrogate to reduce CFD-related computational complexity, combined with a power flow audit for power-flow feasibility verification.

The remainder of this paper is organized as follows. Section II outlines the system model. Section III presents the MILP formulation. Section IV covers the preprocessing, solution, and power flow audit procedure. Section V presents the case study, and Section VI concludes this paper.

II. System Model

Consider a typical distribution feeder as shown in Fig. 1. The feeder is represented by a graph $\mathcal{D} = (\mathcal{N}, \mathcal{E})$ with a node set $\mathcal{N}$ and a line set $\mathcal{E}$. Each line $e \in \mathcal{E}$ has resistance R_e, reactance X_e, apparent power rating $\overline{S}_e$, and operates within voltage limits $V^{\min} \le V \le V^{\max}$. For restoration tasks, components such as RCSs and MEGs are commonly utilized. RCSs form a subset $\mathcal{R} \subseteq \mathcal{E}$. Critical loads are indexed by $\mathcal{C} \subseteq \mathcal{N}$; load $c \in \mathcal{C}$ has a demand of D_c and a priority weight of w_c. Candidate MEG sites $\mathcal{M} \subseteq \mathcal{N}$ host staged MEGs that energize restored paths k from sites to critical loads, subject to thermal and voltage limits. Base stations provide communication links to RCSs along each path, and wildfire exposure determines line availability and switch operability, as shown in Fig. 1. Wildfire exposure is represented by binary masks for $a^{\omega}_{e,t}$ and $h^{\omega}_{r,t}$. For each line $e \in \mathcal{E}$ and time t in scenario ω, $a^{\omega}_{e,t} \in \{0, 1\}$ equals one if e is safe to energize. For each RCS $r \in \mathcal{R}$, $h^{\omega}_{r,t} \in \{0, 1\}$ equals one if r is operable. The following subsections detail the MEG and restoration path models, and the wildfire and communication link models.

A. MEG and Restoration Path Models

In this work, the MEG is considered a means of restoring the distribution system. Site $m \in \mathcal{M}$ has staging capacity cap_m and tie-in limit $\overline{P}^{\mathrm{site}}_m$. The MEG fleet is indexed by $\mathcal{S}$, where each unit $s \in \mathcal{S}$ has a real-power rating of $\overline{P}_s$, an energy budget of $\overline{E}_s$, and a minimum power factor of $\mathrm{pf}_{\min}$. Throughout the restoration horizon, the energized topology remains radial. Restoration decisions are made over candidate paths from generator sites to critical loads. For each site $m \in \mathcal{M}$, a set $\mathcal{K}_m$ of simple paths to loads is constructed, and $\mathcal{K} := \bigcup_{m \in \mathcal{M}} \mathcal{K}_m$. Each path $k \in \mathcal{K}$ has root site $m(k)$, critical load $c(k)$, line set $\mathcal{E}_k \subseteq \mathcal{E}$, and switch set $\mathcal{R}_k \subseteq \mathcal{R}$. For each critical load $c \in \mathcal{C}$, we define $\mathcal{K}_c \subseteq \mathcal{K}$ as the set of all paths k that serve that load. During dispatch, each MEG energizes at most one path at a time, and path power is limited by thermal and voltage-drop capacities. The binary masks $a^{\omega}_{e,t}$ and $h^{\omega}_{r,t}$ defined above directly affect the operational capacity of the restoration paths. Line ratings are assumed nominal, i.e., $\overline{S}^{\omega}_{e,t} = \overline{S}_e$. The thermal bottleneck for a path $k \in \mathcal{K}$ is

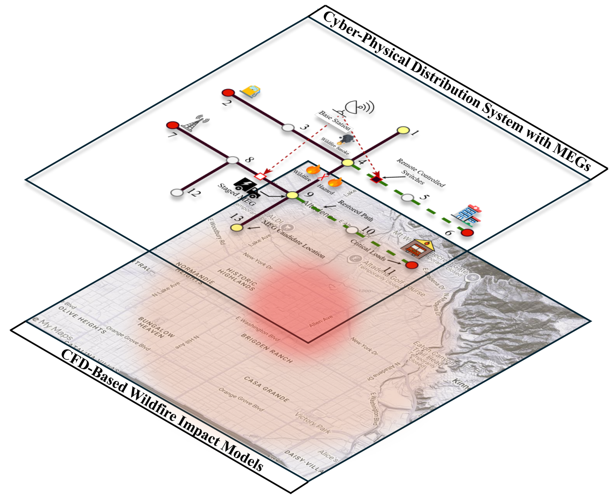

Fig. 1. An illustration of the communication-link-aware restoration of cyber-physical distribution systems under wildfires.

$\overline{S}^{\omega}_{k,t} = \min_{e \in \mathcal{E}_k} a^{\omega}_{e,t} \overline{S}^{\omega}_{e,t}$, and the corresponding real-power cap is $\overline{P}^{\omega}_{k,t} = \mathrm{pf}_{\min} \overline{S}^{\omega}_{k,t}$.

B. Wildfire and Communication Link Models

CFD is used to model the wildfire hazard as a physics-based approach that solves the Navier-Stokes equations with combustion and radiation, producing wind, temperature, and heat flux fields that inform line and switch safety. The Navier-Stokes equations express a conservation of Mass:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = \dot{m}_{\mathrm{pyro}}, \tag{1}$$

Momentum:

$$\rho \frac{\mathrm{D}\mathbf{u}}{\mathrm{D}t} = -\nabla p + \nabla \cdot (\mu(\nabla \mathbf{u} + \nabla \mathbf{u}^{\top})) + \rho \mathbf{g} + \dot{\mathbf{m}}, \tag{2}$$

and Energy:

$$\rho c_p \frac{\mathrm{D}T}{\mathrm{D}t} = \nabla \cdot (k_{\mathrm{th}} \nabla T) + \dot{Q}_{\mathrm{comb}} + \dot{Q}_{\mathrm{rad}}, \tag{3}$$

where ρ, $\mathbf{u}$, p, μ, c_p, T, and k_{th} represent density, velocity, pressure, viscosity, specific heat, temperature, thermal conductivity, respectively. Also, $\dot{Q}_{\mathrm{comb}}$ and $\dot{Q}_{\mathrm{rad}}$ are the volumetric heat release rates from combustion and radiation, respectively. Here, $\dot{m}_{\mathrm{pyro}}$ is the mass source from pyrolysis, $\mathbf{g}$ is gravity, and $\dot{\mathbf{m}}$ is the momentum source. In this framework, $\dot{Q}_{\mathrm{comb}}$ is governed by an Arrhenius-based pyrolysis model, and $\dot{Q}_{\mathrm{rad}}$ is solved using a vertical two-sweep discrete ordinates model (DOM) combined with a direct view-factor calculation. The binary mask values are determined by thresholding the high-fidelity CFD outputs (e.g., temperature, heat flux) against equipment safety limits.

Just as the physical wildfire determines distribution system equipment's availability, the cyber layer determines their remote operation. The restoration process requires functional communication links to operate RCSs. The cyber layer is modeled via a set of base stations $\mathcal{B}$ that govern both switching concurrency and command success. A binary map $\nu_{k,b}$ indicates that path k draws commands from base station b ($b \in \mathcal{B}$). The per-base-station concurrency capacity is defined as $\mathrm{cap}^{\mathrm{comm}}_{b,t}$, which indicates the maximum number of new

979-8-3315-8138-1

switching operations that base station b can issue at time t. Let $\pi^{\text{base}} \in [0,1]$ denote baseline base station uptime. The probability for all commands along path k to succeed is modeled as $\widehat{\psi}^{\omega}_{k,t} = \prod_{r\in\mathcal{R}_k}\big(\pi^{\text{base}}\cdot \hat{p}_h(r,t)\big)$, where $\hat{p}_h(r,t)$ is the probability of RCS r being operational at time t. If a path involves switches served by multiple base stations, we set $\nu_{k,b}=1$ for each base station so that a single new activation along the path consumes one unit of capacity.

III. Formulation of the Cyber-Physical Distribution System Restoration Problem

To address the challenges of cyber-physical distribution system restoration under wildfires, we need to co-optimize discrete strategic decisions (e.g., MEG siting, path switching) with continuous power dispatch under uncertainty. Accordingly, we formulate this problem as a two-stage MILP, where the first stage determines the static MEG assignments, while the second stage performs the dynamic path activation and power dispatch for each time period and scenario. The formulation integrates both physical feasibility based on binary masks derived from CFD models and cyber-layer operational constraints defined by communication concurrency caps.

A. Objective Function

The objective is to maximize the expected, discounted, and priority-weighted energy delivered to critical loads while penalizing new remote switching and MEG staging, given by

$$\max_{\mathcal{X}} \sum_{\omega\in\Omega}\pi_\omega \sum_{t\in\mathcal{T}}\beta^t\Big[\sum_{s\in\mathcal{S}}\sum_{k\in\mathcal{K}} w_{c(k)}\,\widehat{\psi}^{\omega}_{k,t}\,p^{\omega}_{ks,t}\,\Delta t - \alpha\sum_{k\in\mathcal{K}} N_k\,\Delta y^{\omega}_{k,t}\Big] - \rho_{\text{site}}\sum_{s\in\mathcal{S}}\sum_{m\in\mathcal{M}}\kappa_{sm}\,x_{sm}, \tag{4}$$

where $\mathcal{X}$ is the set of decision variables. These include binary variables for MEG staging (x_{sm}), path assignment to MEGs ($z^{\omega}_{ks,t}$), path energization status ($y^{\omega}_{k,t}$), new path activations ($\Delta y^{\omega}_{k,t}$), and radiality ($u^{\omega}_{em,t}$, v^{ω}_{nmt}), as well as continuous variables for power dispatch ($p^{\omega}_{ks,t}$) and load served ($r^{\omega}_{c,t}$). Here, π_ω is scenario probability, β is the discount factor, and Δt is the time-step duration. The term $w_{c(k)}$ is the priority weight for the critical load $c(k)$ served by path k, and $\widehat{\psi}^{\omega}_{k,t}$ is the communication success probability defined in the previous section. Penalties include α for new activations (scaled by N_k, the number of RCSs on path k), ρ_{site} for siting, and κ_{sm} for the cost of staging MEG s at site m.

B. Constraints

The model enforces staging, activation, topology, operational limits, and load exclusivity in a single set of constraints. Staging assigns each MEG to a unique site and respects per-site capacity:

$$\sum_{m\in\mathcal{M}} x_{sm}=1, \forall s\in\mathcal{S};\quad \sum_{s\in\mathcal{S}} x_{sm}\le \text{cap}_m, \forall m\in\mathcal{M}. \tag{5}$$

Activation and topology are determined (gated) by siting, wildfire, and communication link feasibility. For path k with root $m(k)$, we define line gate $g^{\omega}_{k,t}=\prod_{e\in\mathcal{E}_k} a^{\omega}_{e,t}$ and switch gate $\eta^{\omega}_{k,t}=\min_{r\in\mathcal{R}_k} h^{\omega}_{r,t}$, which satisfy

$$z^{\omega}_{ks,t}\le \min(x_{s\,m(k)}, g^{\omega}_{k,t}, \eta^{\omega}_{k,t}),\quad \forall k,s,t,\omega. \tag{6}$$

$$\sum_{k\in\mathcal{K}} z^{\omega}_{ks,t}\le 1,\ \forall s,t,\omega;\quad z^{\omega}_{ks,t}\le y^{\omega}_{k,t}\le \sum_{s\in\mathcal{S}} z^{\omega}_{ks,t},\ \forall k,t,\omega. \tag{7}$$

Each MEG selects at most one path, and an energized path reflects at least one selection, constrained as follows:

$$\Delta y^{\omega}_{k,t}\ge y^{\omega}_{k,t}-y^{\omega}_{k,t-1},\quad \Delta y^{\omega}_{k,t}\le y^{\omega}_{k,t},\quad \Delta y^{\omega}_{k,t}\le 1-y^{\omega}_{k,t-1}, \tag{8}$$

with initial $y^{\omega}_{k,0}$ given by the pre-dispatch configuration. Also, at most one path may serve a given critical load at any time, so that $\sum_{s\in\mathcal{S}}\sum_{k\in\mathcal{K}_c} z^{\omega}_{ks,t}\le 1$ for all $c\in\mathcal{C}$.

The radial requirement is enforced by binary ownership indicators for lines ($u^{\omega}_{em,t}$) and nodes (v^{ω}_{nmt}) that partition energized components by site. These constraints (not shown for brevity), together with the prefix-closed and post-divergence edge-disjoint path construction, produce a radial forest per site.

Operational limits related to MEG ratings, path bottlenecks and voltage caps, line thermals, site tie-in limits, communication concurrency, and MEG energy budgets are as follows:

$$0\le p^{\omega}_{ks,t}\le \overline{P}_s\, z^{\omega}_{ks,t},\qquad \forall k,s,t,\omega, \tag{9}$$

$$\sum_{s\in\mathcal{S}} p^{\omega}_{ks,t}\le \min(\overline{P}^{\omega}_{k,t}, \overline{P}^{\text{volt}}_k)\, y^{\omega}_{k,t},\quad \forall k,t,\omega, \tag{10}$$

$$r^{\omega}_{c,t}=\sum_s\sum_{k\in\mathcal{K}_c} p^{\omega}_{ks,t},\quad 0\le r^{\omega}_{c,t}\le D_c,\ \forall c,t,\omega, \tag{11}$$

$$\sum_s\sum_{k:\,e\in\mathcal{E}_k} p^{\omega}_{ks,t}\le \text{pf}_{\min}\, a^{\omega}_{e,t}\,\overline{S}^{\omega}_{e,t},\quad \forall e,t,\omega, \tag{12}$$

$$\sum_s\sum_{k\in\mathcal{K}_m} p^{\omega}_{ks,t}\le \overline{P}^{\text{site}}_m,\quad \forall m,t,\omega, \tag{13}$$

$$\sum_{k\in\mathcal{K}} \nu_{k,b}\,\Delta y^{\omega}_{k,t}\le \text{cap}^{\text{comm}}_{b,t},\quad \forall b,t,\omega, \tag{14}$$

$$\sum_t\sum_k p^{\omega}_{ks,t}\,\Delta t\le \overline{E}_s,\quad \forall s,\omega, \tag{15}$$

where $\overline{P}^{\text{volt}}_k$ is the voltage-drop cap computed once per path from its total resistance and reactance.

IV. Problem Solution with CFD-based MFML

In this section, we present the solution methodology of the cyber-physical distribution system restoration problem. Solving the two-stage optimization problem is challenging due to CFD's computational complexity and the optimization problem's high dimensionality. We address this challenge with a two-step framework: 1) A data preprocessing step using a CFD-based MFML surrogate to efficiently generate hazard inputs; and 2) A monolithic MILP solver with power flow post-validation. The MILP yields an initial schedule that an iterative power flow audit subsequently de-rates to enforce power-flow feasibility, achieving a practical balance between modeling accuracy and computational tractability.

A. Data Preprocessing and MFML Surrogate

We first preprocess the grid topology and wildfire data to create the necessary optimization inputs. Candidate path sets $\mathcal{K}_m$ are constructed for each MEG site by a breadth-first enumeration of simple feeder paths to critical loads. To

979-8-3315-8138-1

manage the problem's complexity, the path sets are prefix-closed: if a path is included, then all its prefixes are also included. After the first divergence of any pair of paths, the remainder of the two paths must be edge-disjoint, ensuring that once the lines split, there are no re-convergences. To further bound the problem size, we retain only a user-defined number of the shortest $m \to c$ paths per load.

Then, an MFML surrogate is trained offline to efficiently map inexpensive and low-fidelity features such as perimeter distance, wind exposure, and a CFD optical-depth proxy to the safety masks $(a^{\omega}_{e,t}, h^{\omega}_{r,t})$. This provides the physics-based inputs without online CFD cost. The surrogate maps features to high-fidelity classifiers f_A (line availability) and f_H (switch operability), anchored by CFD samples, which are generated by solving the governing equations presented in Subsection II-B. We use a strict train-on-A, test-on-B holdout. Thresholds τ_a, τ_h are calibrated on CFD anchors with caps on False Negative Rate (FNR) and False Positive Rate (FPR), or via a weighted cost $J = w_{\mathrm{FN}}\mathrm{FNR} + w_{\mathrm{FP}}\mathrm{FPR}$. Physics guidance acts as a safety control: Hard standoff rules revert to simple thermal logic when inputs fall outside the validated feature-space, preventing unsafe extrapolation. At runtime, low-fidelity features $X^{\mathrm{LFM}}_{e,t}, X^{\mathrm{LFM}}_{r,t}$ are computed, feeding the surrogates to produce probabilities $\hat{p}_a(e,t), \hat{p}_h(r,t)$, which are thresholded to yield binary gates $a^{\omega}_{e,t}, h^{\omega}_{r,t}$. From these, we can derive path-level gates $g^{\omega}_{k,t}, \eta^{\omega}_{k,t}$, communication maps $\nu_{k,b}$, success weights, and initial states $y^{\omega}_{k,0}$.

B. Monolithic MILP Solver and Power Flow Post-Validation

To reduce the computational complexity related to the coupled discrete switchings with non-linear power flows, we decouple the MILP from full power flow physics. First, the two-stage model is solved monolithically using linearized caps $(\overline{P}^{\mathrm{volt}}_k)$ to propose a schedule. Then, this schedule is audited with a single-phase power flow on each energized subgraph. If violations are detected, a scaling factor $\xi_k \in (0,1]$ is computed for the offending path k. The MILP is re-solved with a tightened cap $\overline{P}^{\omega}_{k,t} \leftarrow \xi_k \, \overline{P}^{\omega}_{k,t}$. This audit-and-derating loop repeats until the dispatch is feasible.

V. Case Study

The proposed framework is evaluated on a modified IEEE 13-Node Test Feeder [15], subjected to a 60-minute wildfire scenario. The simulation is based on real-world digital elevation model (DEM) and fuel map data from the Eaton Canyon area in Southern California, sourced from [16] and [17], respectively. The system includes four critical loads, three MEGs at four candidate sites, and two RCSs, as shown in Fig. 1. The MILP models are solved using Gurobi 12.0.3 [18]. The following subsections benchmark the wildfire models, validate the MFML surrogate, and test the power flow audit. To illustrate the core physical mismatch between modeling approaches central to this case study, a conceptual diagram based on the Eaton Canyon topography is shown in Fig. 2. Three distinct wildfire models are considered:

(1) The baseline Kinematic model;
(2) High-fidelity CFD benchmark;
(3) The proposed MFML surrogate.

The Kinematic model, widely used in literature for its speed, applies a simple 2D geometric standoff rule (25 m) to the 600 K ground-level flame perimeter, while the CFD and MFML models use physics-based inputs to determine line availability.

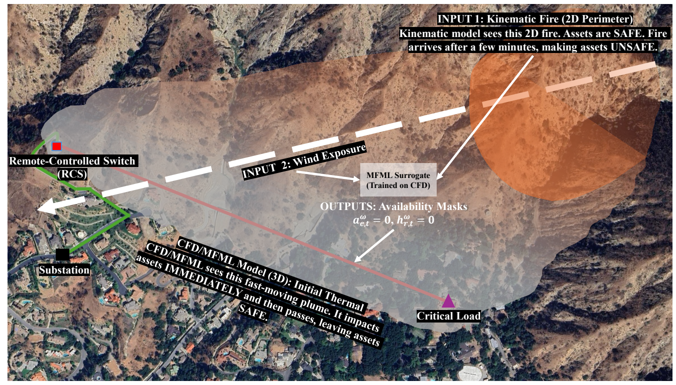

Fig. 2. A conceptual illustration of the wildfire and the case study's hazard.

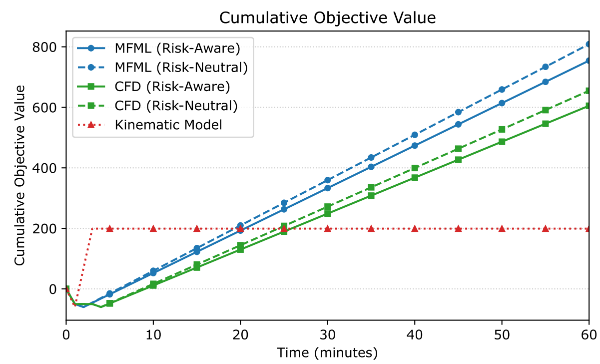

Fig. 3. Comparison of cumulative objective value by model.

A. Benchmarking Wildfire Models for Grid Restoration

The impact of physical modeling accuracy, which was conceptually illustrated in Fig. 2, is quantified in Fig. 3. The Kinematic model's 2D-only view provides a misleading hazard assessment. It fails to detect the initial 3D thermal plume and then becomes overly conservative, leading to a low objective of 199.17 (red dotted line). In contrast, the CFD-aware (605.19, green solid line) and MFML-aware (754.42, blue solid line) models leverage their superior physical accuracy to achieve 3.0x to 3.7x higher objective values. This gap is attributable to the physics-aware models safely identifying 57-59 minutes of available operation versus the 3 minutes identified by the Kinematic model.

To demonstrate the importance of cyber-physical coupling, we also evaluate a "Risk-Neutral" model that ignores the communication success probability ($\widehat{\psi}^{\omega}_{k,t} = 1$ in (4)). The results in Fig. 3 show that this model is overly optimistic (blue dashed line) due to its assumption of perfect communication links. By ignoring the probability of communication failure, these models select a different, higher-risk set of restoration paths, believing them to be more valuable than they actually are. This highlights that both high-fidelity physical hazard assessment and probabilistic cyber-risk modeling are essential.

979-8-3315-8138-1

TABLE I
MFML Surrogate Model Holdout Performance

Metric	Availability (f_A)	RCS Operability (f_H)
Model Type	Bagged Tree	Bagged Tree
Holdout AUC	0.941	0.985
Holdout Accuracy	89.3%	92.3%
Holdout FNR	11.7%	13.0%
Holdout FPR	10.6%	5.4%

B. MFML Surrogate Performance Validation

The statistical accuracy and computational efficiency of the MFML surrogate, as a replacement for the full CFD simulation, are shown in Table I. The surrogate, which is a Bagged Tree model, provides hazard masks in under 5 minutes, a significant reduction from the 39.4 hours required by the high-fidelity CFD model. As described in Section IV, the model was trained offline on computationally inexpensive features (such as perimeter distance, wind exposure, and a CFD optical-depth proxy) anchored by high-fidelity CFD samples. The framework demonstrates high predictive effectiveness, achieving a holdout Area Under Curve (AUC) of 0.941 for line availability (f_A) and 0.985 for RCS operability (f_H). These metrics prove the model's ability to reliably distinguish between safe and unsafe conditions. The error rates also demonstrate a safety-conscious tradeoff. For RCS operability, the model achieved a low 5.4% FPR, demonstrating high safety, with an acceptable 13.0% FNR. For line availability, the MFML surrogate predicted 59 available minutes, slightly more optimistic than the CFD's 57 available minutes. This 2-minute discrepancy leads to the 10.6% FPR.

C. Validation of the Power Flow Audit Workflow

To validate the power flow audit workflow, a congestion test was designed to create an intentional mismatch between the model's assumptions and realistic load flows. We selected one path to the 843.0 kW load at bus 675, synthetically increased its series impedance by a factor of 10 (lowering its true capacity to ≈34.8 kW), and set its linearized cap $\overline{P}_k^{\text{volt}}$ to an optimistic 900.0 kW. All other paths to this load were disabled, guaranteeing the initial MILP dispatch would be infeasible. Then, the power flow audit loop is engaged as designed. The iterative convergence, as shown in Fig. 4, demonstrates the workflow's corrective action. In Iteration 1, the MILP dispatches 843.0 kW (bound by load demand). The power flow audit failed, detecting a voltage collapse to 0.668 p.u. The audit applied a 50% de-rating factor, setting the new cap to 450.0 kW. In Iteration 2, the MILP is solved again with this new cap, which still resulted in a voltage limit violation (0.902 p.u.). The process repeated, cutting the cap to 225.0 kW for Iteration 3, which also narrowly failed (0.941 p.u.). Finally, in Iteration 4, the cap was reduced to 112.5 kW. The MILP dispatched 112.5 kW, which the power flow audit confirmed as feasible, raising the minimum voltage to 0.975 p.u. and successfully terminating the loop. This test case confirms that the linearized $\overline{P}_k^{\text{volt}}$ cap can be optimistic and validates that the proposed power flow audit workflow is essential for ensuring a feasible final restoration schedule.

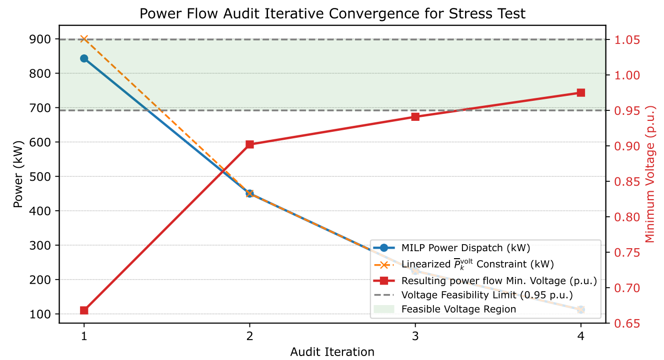

Fig. 4. Iterative convergence of the power flow audit workflow for stress test.

VI. Conclusion

This paper proposed a two-stage optimization for wildfire-aware MEG siting and dispatch, integrating a CFD-based MFML surrogate with an MILP that uses communication success probability to guide its decisions. Case study results validated the proposed approach, showing the MFML model can enable superior physics-aware outcomes compared to simpler kinematic models. Furthermore, the power flow audit workflow was proven necessary to ensure the final dispatch schedules were physically valid. By integrating multi-fidelity physics with cyber-layer constraints, this framework provides a tractable tool for utilities to dispatch MEG assets efficiently. Future work will focus on the generalization of the surrogates for the integration of stochastic repair crew dispatch.

References

[1] Government of the Northwest Territories, "2023 Wildfire Emergency Response After-Action Review." Yellowknife, NT: Department of Municipal and Community Affairs, 2024. [Online]. Available: https://www.maca.gov.nt.ca/sites/maca/files/resources/nwt_after-action_review_-_bilingual.pdf. Accessed: Nov. 5, 2025.

[2] L. Cecco, "No internet, no phone: Canada wildfires expose fragility of rural infrastructure," *The Guardian*, May 15, 2024. [Online]. Available: https://www.theguardian.com/world/article/2024/may/15/internet-access-rural-canada-wildfires. Accessed: Nov. 5, 2025.

[3] S. Lei, J. Wang, C. Chen, and Y. Hou, "Mobile emergency generator pre-positioning and real-time allocation for resilient response to natural disasters," *IEEE Trans. Smart Grid*, vol. 9, no. 3, pp. 2030–2041, May 2018.

[4] S. Lei, C. Chen, Y. Li, and Y. Hou, "Resilient disaster recovery logistics of distribution systems: co-optimize service restoration with repair crew and mobile power source dispatch," *IEEE Trans. Smart Grid*, vol. 10, no. 6, pp. 6187–6202, Nov. 2019.

[5] W. Shi, H. Liang, and M. Bittner, "Data-driven resilience enhancement for power distribution systems against multishocks of earthquakes," *IEEE Trans. Ind. Informat.*, vol. 20, no. 5, pp. 7357–7369, May 2024.

[6] S. Lei, J. Wang, Y. Hou, Y. Sun, M. Shahidehpour, and C. Chen, "Remote-controlled switch allocation enabling prompt restoration of distribution systems," *IEEE Trans. Power Syst.*, vol. 33, no. 3, pp. 3129–3142, May 2018.

[7] A. Arif, B. Cui, and Z. Wang, "Switching device-cognizant sequential distribution system restoration," *IEEE Trans. Power Syst.*, vol. 37, no. 1, pp. 317–329, Jan. 2022.

[8] W. Shi, H. Liang, and M. Bittner, "Stochastic sequential restoration for resilient cyber-physical power distribution systems," *IEEE Trans. Ind. Informat.*, vol. 21, no. 2, pp. 1200–1209, Feb. 2025.

[9] J. L. Coen, W. Schroeder, S. Conway, and L. Tarnay, "Computational modeling of extreme wildland fire events: a synthesis of scientific understanding with applications to forecasting, land management, and firefighter safety," *J. Comput. Sci.*, vol. 45, art. 101152, Sep. 2020.

[10] G. Accary and D. Morvan, "How can CFD contribute to the understanding of wildfire behaviour?" *Comput. Fluids*, vol. 279, art. 106322, July 2024.

[11] B. Peherstorfer, K. Willcox, and M. Gunzburger, "Survey of multifidelity methods in uncertainty propagation, inference and optimization," *SIAM Rev.*, vol. 60, no. 3, pp. 550–591, Aug. 2018.

[12] L. Brunel, M. Balesdent, L. Brevault, R. Le Riche, and B. Sudret, "A survey on multi-fidelity surrogates for simulators with functional outputs: unified framework and benchmark," *Comput. Methods Appl. Mech. Eng.*, vol. 435, art. 117577, Feb. 2025.

[13] M. C. Kennedy and A. O'Hagan, "Bayesian calibration of computer models," *J. R. Stat. Soc. B*, vol. 63, no. 3, pp. 425–464, Sep. 2001.

[14] A. I. J. Forrester, A. Sóbester, and A. J. Keane, "Multi-fidelity optimisation via surrogate modelling," *Proc. R. Soc. A*, vol. 463, no. 2088, pp. 3251–3269, Dec. 2007.

[15] IEEE Power & Energy Society, Distribution System Analysis Subcommittee, "IEEE PES Test Feeder," 2018. [Online]. Available: https://cmte.ieee.org/pes-testfeeders/resources/. Accessed: Aug. 10, 2025.

[16] U.S. Geological Survey, "The National Map – Data Delivery," U.S. Dept. of the Interior. [Online]. Available: https://www.usgs.gov/the-national-map-data-delivery. Accessed: Jun. 17, 2025.

[17] LANDFIRE, "LANDFIRE Program," U.S. Department of the Interior, Geological Survey, and U.S. Department of Agriculture. [Online]. Available: https://www.landfire.gov. Accessed: Jun. 19, 2025.

[18] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual." 2025. [Online]. Available: https://www.gurobi.com. Accessed: Nov. 5, 2025.

979-8-3315-8138-1