Modeling WECC electric grid wildfire resilience with security-constrained optimal power flow

Runako Gentles
Civil and Environmental Engineering
Stanford, CA
runako@stanford.edu

Elbert Gong
Computational and Mathematical Engineering
Stanford, CA
elbygong@stanford.edu

***Abstract*—We analyze the Western-US electrical grid's wildfire resilience using security-constrained and "N-k" contingency analysis of a DC optimal power flow model. We model load shedding and system cost as a function of fire-induced line outages. We identify specific lines and regions that are especially vulnerable to fire risk or important for system stability.**

I. Introduction

As climate change increases the risk of natural disasters, such as wildfires, the grid must be prepared to handle more intense disruptions. For instance, the Jan 2025 Los Angeles wildfires led to large-scale power outages [1]. The challenge is exacerbated by the growth in renewable energy sources, which are distributed, intermittent, and less predictable than fossil-fuel generation. In California, peak clean power supply is projected to more than double in the coming decades [1]. As the complexity of keeping the grid in balance increases, system operators are searching for more sophisticated long-term and real-time contingency analysis tools.

We build an operational power flow model using the PyPSA-USA package (Python for Power Systems Analysis). PyPSA-USA integrates various data sources to build a model of transmission lines, loads, generation, and storage. PyPSA can perform *optimal power flow* (OPF) to allocate power production across generators to minimize cost, subject to certain feasibility constraints. To assess the grid's resilience to fire-based outages, we add *security constraints*, which stipulate that the system must stay feasible even if k lines go out. We also perform *contingency analysis*, which re-solves the optimization problem supposing that k lines have gone out.

For large k, only a small fraction of the $\binom{n}{k}$ possible contingencies can analyzed, due to computational limits. And not all contingencies are equally important. System operators need a way to determine which contingency scenarios are especially likely. Thus, we build a simple model of wildfire outage risk based on historical US wildfire data. We use this data to assign a "burn score" to each transmission line, and we perform a sensitivity analysis of system cost as k increases. We also perform a joint analysis of fire risk and system cost in several specific high-risk burn zones, especially the 5th-highest-risk zone near the California-Oregon border.

We show how security-constrained OPF and contingency OPF provide two perspectives for looking at resilience– the former from an *a priori* perspective and the latter from an *a posteriori* perspective. With security-constrained OPF, we identify an inflection point in the system cost curve due to a single "transmission line that broke the camel's back" in Oregon.

Insights from these types of models can be useful for day-ahead planning and for fast (non-instantaneous) response to unplanned disruptions. A wildfire is difficult to predict more than days in advance, but once it begins, it can last for several days. However, steady-state power flow techniques are inadequate for modeling harmonic effects or high-frequency dynamics in the seconds after an unplanned disruption. Ultimately, we will need many flavors of models at different spatial/temporal resolutions as we adapt to a changing grid and a changing climate.

II. Literature Review

Contingency analysis is a well-established method for studying power system security and reliability. Operators commonly run "N-1" contingencies (one line outage), "N-1-1" (two independent outages), and "N-k" (multiple simultaneous outages) [16]. There are two popular methods for selecting the most important contingencies: ranking methods and screening methods [16]. Contingency analysis with DC power flow is much less computationally intensive than AC analysis without sacrificing too much model performance [4].

There is growing evidence that wildfire risk is increasing due to climate change [13]. Wildfire heat affects the thermal capacity of transmission lines, and the fires themselves cause physical damage to electrical equipment [2]. Lines operating outside voltage capacity limits are more likely to arc and set fire to nearby vegetation, causing a self-reinforcing feedback loop [2]. Haze from wildfires also limits solar PV generation [13].

Several organizations are coming up with new ways to incorporate climate data into existing power modeling tools. For instance, the Probabilistic Resource Adequacy Suite simulates outage events to quantify the risk of unserved load [14]. Sayarshad et al. integrate OPF analysis with a dynamic fire behavior model and a model of line heating and cooling [3]. Taylor et al. modify the OPF objective function to consider the ability of customers to cope with power outages, factoring in both wildfire risk and social vulnerability [17].

However, these models are likely to be subject to significant bias due to the scarcity of damage data and the variations in power grid assets across different geographic regions [2]. Also, there are relatively few studies of wildfires and floods (including climate change effects) compared to earthquakes and wind hazards [2]. In practice, utilities' decisions are often based on expert judgment rather than sophisticated weather-based risk models [15]. The translation of an OPF solution into actual planning and operating decisions is inherently subjective and not clear-cut [15].

III. Problem Setup / Discussion

A. Power flow model

A power flow model treats the grid as a graph of buses connected by transmission lines. The full AC power flow equation (derived from Kirchhoff's laws) is

$$P_i = \sum_k |V_i||V_k|(G_{ik}\cos(\theta_i - \theta_k) + B_{ik}\sin(\theta_i - \theta_k))$$

where index k sums over all buses connected to bus i. P, $|V|$, and θ are bus real power, voltage magnitudes, and phase angles respectively. G and B are line conductance and susceptance.

If we apply some approximations that are commonly used in high-voltage transmission systems, we have the linearized DC power flow equation

$$P_i = \sum_k B_{ik}(\theta_i - \theta_k) \tag{1}$$

which can be solved much more efficiently using iterative methods from linear algebra. Thus, given a set of known quantities (i.e. the P_i), we can efficiently solve for unknowns (i.e. the θ_i).

In real life, we control the power production of the generators, and we can coordinate them to minimize some cost function subject to the constraint imposed by (1). This is the linear OPF problem. In PyPSA, the cost function is the sum of the capital costs of each component and the generation costs of most components. For our *operational analysis* (as opposed to *capacity expansion*), we hold capital expenditures fixed and focus on operational expenditures, modeled as a linear function of power generated at each bus i with marginal price c_i.

$$\min \sum_i c_i P_i$$

The PyPSA model also includes various constraints related to capacity, reliability, and temporal coupling. With convex constraints and a linear objective function, the OPF can be solved using methods from convex optimization, like the interior point algorithm.

PyPSA tries to find a feasible solution if one exists; however, if the system is infeasible, then it introduces load-shedding (rolling blackouts) to relax the load constraints. Load-shedding is modeled as a quasi-generator at each bus with a very high marginal cost of \$1000/MW (so it is the generator of last resort on the merit order curve).

Security-Constrained OPF adds a constraint that the system must stay feasible even if k transmission lines go out. For a line c, the constraint is

$$|f_b + BODF_{bc} f_c| \leq F_b \ \forall b$$

where f_b is the line flow along branch b (a function of line susceptance and θ), F_b is the capacity, and $BODF_{bc}$ is the branch-outage distribution factor: the amount that f_b would increase if c experienced an outage [5]. Intuitively, if line c is known to be at-risk, we want to leave it some "breathing room" so that the system wouldn't collapse if it went out.

In *contingency OPF*, we re-solve the optimization problem supposing that k lines have gone out. This is a stricter constraint than security-constrained OPF, as we are prevented entirely from sending power on these lines– F_c is effectively 0.

B. Fire model

Our fire dataset partitions the US into grid cells of size 300 mi^2. It contains a historical record of whether or not there was a wildfire during each day in each grid cell. For grid cell i, we estimate its fire probability as

$$P_{\text{fire}}(i) = \frac{\text{\# days with fire}}{\text{total \# days}} = \frac{\text{\# days with fire}}{3652}$$

We use only the 10 most recent years of data for better alignment with modern climate conditions. The recent data is also more accurate thanks to modern remote sensing technology.

For a transmission line between buses i and k, we use Bresenham's line algorithm [7] to identify the grid cells covering the line segment between the centroids of i and k. Then we calculate the line's *burn score* by simply summing all the probabilities of the grid cells. This burn score is not a probability per se, since it can be greater than 1, but it is a better heuristic than averaging all the probabilities, since longer lines cover more area and are inherently more exposed to burn risk.

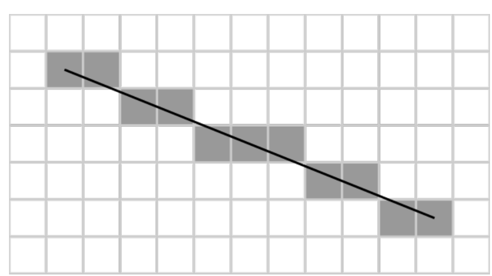

Fig. 1. Bresenham's line algorithm.

We rank-order lines by burn score to prioritize them for contingency analysis. For instance, for an $N-k$ analysis with $k = 3$, we select the three lines with the highest burn score (i.e. the three yellowest lines in figure 2).

This approach, while simple, assumes that each line fails independently, so it doesn't capture the correlations between lines close together in space. Thus, we also use another

Fig. 2. Transmission line burn scores.

approach based on thresholding and connected-component labeling to identify high-risk clusters in space.

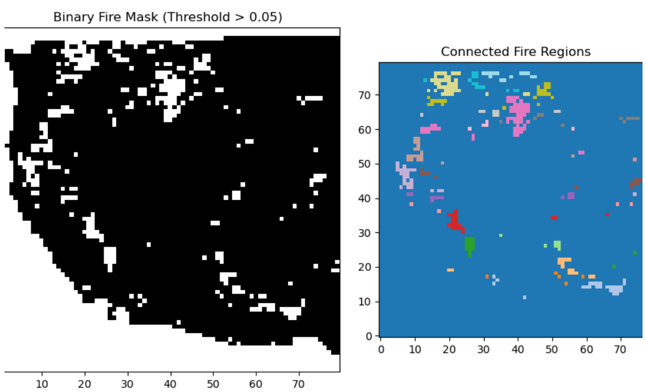

Fig. 3. Burn clusters in space.

We compute a cluster's burn score by summing over its grid cells. See figure 4.

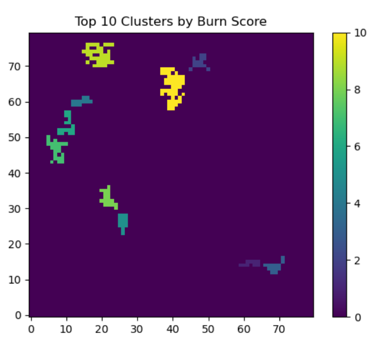

Fig. 4. Top 10 clusters by burn score.

Our simple model is limited because it doesn't account for the evolution of a fire in space and time. Moreover, if the PyPSA model's clusters are too coarse, the centroid-connecting line segments don't exactly correspond to actual physical power lines. (And power lines are not necessarily straight lines.) Nevertheless, this model serves as a first approximation of fire risk at sub-network granularity.

IV. Methodology

We obtain the wildfire dataset from our mentor Tao Sun. It synthesizes multiple source datasets, including InFORM, the Fire Program Analysis fire-occurrence database (FPA-FOD), WFIGS, InterAgencyFirePerimeterHistory, and GeoMAC, removing apparent duplicates by cross-checking overlapping incidents with matching or near-matching ignition dates, fire names, or bounding polygons.

We build the power flow model with the PyPSA-USA open source package [6]. The model topology uses the Texas A&M University Synthetic Nodal Network [8]. The demand profile uses the National Renewable Energy Laboratory's Electrification Futures Study [9].

TABLE I
PyPSA-USA Model Parameters

Parameter	**Value**
Interconnect	Western
# buses	200
# renewable gen zones	270
Δt	1 hour
Planning horizon	2022
Capacity expansion	disabled
Load-shedding	enabled
Rolling horizon	enabled

The TAMU network has thousands of nodes; however, PyPSA-USA uses the K-means algorithm to cluster them together into hundreds or tens of buses. Choosing the number of buses is a key modeling decision. Given the granularity of the fire dataset, we want to have as granular of a nodal network as possible. But we are constrained by compute resources. We use Slurm to schedule batch jobs on Stanford's Farmshare cluster computing "rye" nodes with 100GB RAM and 4 hrs runtime [10]. But we find it difficult to run with the desired number of buses (400). And we find it difficult to run security-constrained OPF, with its additional constraints, on even 200 buses. So we experiment with some workarounds:

- Using 200 buses, or even just 50 (for prototyping)
- Optimizing over 336 snapshots (2 weeks) instead of 8760 snapshots (1 year), or even just a single snapshot (ignoring time-coupling constraints entirely).
- Optimizing with a rolling horizon (horizon = 48 hrs, overlap = 24 hrs).
- Disabling capacity expansion.
- Enabling parallel compute (8 CPUs per task).
- Using the Gurobi solver instead of HiGHS.

Contingency analysis is not as difficult to run as security-constrained OPF– we simply set the max capacity of the affected lines to 0.

The OPF is not guaranteed to find a solution satisfying all its constraints. (Even with no outages, the 200-bus network is slightly infeasible, due to oddities in the source data.) Thus, for feasibility, it is necessary to enable load-shedding. The

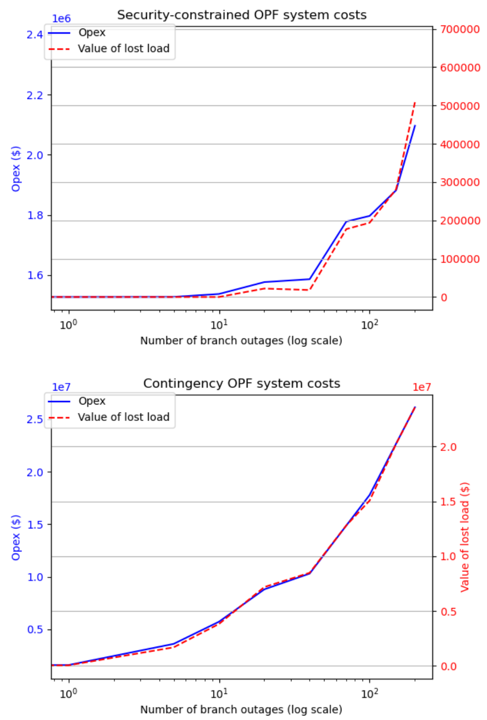

Fig. 5. Comparing system costs between security-constrained OPF and line removal OPF.

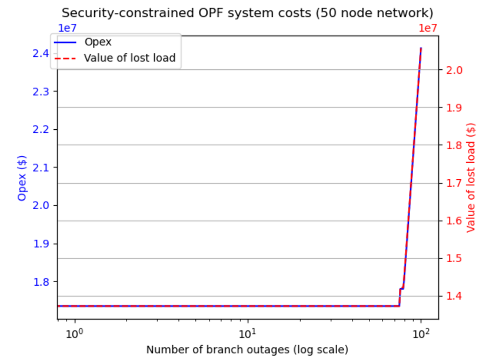

Fig. 6. Security-constrained OPF for a smaller 50-bus network.

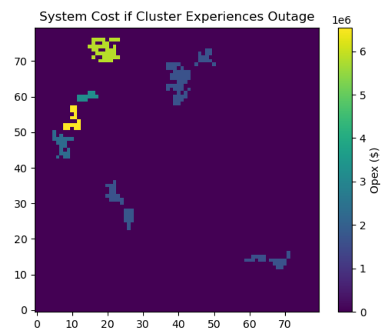

Fig. 7. For each cluster in figure 4, we report the system cost if all the lines touching that cluster go out.

baseline network with no outages has 6.1 MW of average load shed. To quantify the adverse impact of an outage, we measure the amount of additional load shed (MW) and the resultant value of lost load ($).

V. Results

A. *Security-constrained OPF vs contingency OPF*

In figure 5, we perform a sensitivity analysis of system costs as a function of k (the number of outages) for a single snapshot in time. We look at system operating costs, and we see that the primary contributor to the cost increase is the value of lost load. We also observe that security-constrained OPF imposes a weaker constraint on the system than contingency OPF; with security-constrained OPF, we don't observe much of an impact until we hit an inflection point at $k \approx 10$. This means that the grid can be optimized in a way that is resilient to losing 10 lines, without incurring too much additional cost. But with contingency OPF, we see that actually losing even one line incurs a sizable $2m in operating expenses.

With the smaller 50-bus network, with security-constrained OPF, we observe a very sharp inflection point in the cost curve (see figure 6). The 76th line (out of 107) is like the proverbial straw that broke the camel's back. This is much more pronounced than in the 200-bus network; the overall network dynamics appear to be highly dependent on the number of buses. Thus, coarse models may be of limited usefulness, and we should really be running with 400 or more buses.

B. *Contingency analysis of specific regions*

In figure 7, we perform a contingency analysis on each of the top 10 clusters by burn score. The cluster with the highest burn score (in Montana) has a relatively low system cost, while the clusters with the 5th- and 2nd-highest burn scores (in the Pacific Northwest and Northern California/Oregon) have the highest system costs. These Western clusters are closer to population centers and hence touch more lines. when deciding where to invest in preventative measures, a system planner would need to trade off between these two factors (burn risk and system cost).

We look in detail at the California-Oregon cluster, which scores highly on both factors. The CA-OR cluster is in a dense part of the grid– it touches 18 lines in our 200-bus (500-line) network. These lines are largely responsible for transferring the PNW's rich hydro resources down to major population centers in NorCal. If they went out, we would incur a sizable $6m in system costs. We would also experience major load shedding in this region, as shown in figure 8.

In figure 9, we analyze the generation profiles over a 2 week period in July during the height of fire season. Compared to the base case (which has negligible load-shedding), the contingency case has a non-negligible chunk of load-shedding, which stays roughly constant over time. We don't see drastically different load-shedding profiles at different times of day.

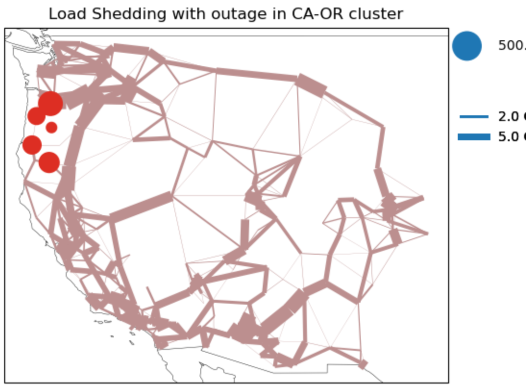

Fig. 8. Bubble sizes denote load-shedding and line widths denote apparent capacity.

Therefore, if we were to optimize on just a single snapshot in time, we would significantly lighten the computational burden without losing too much information.

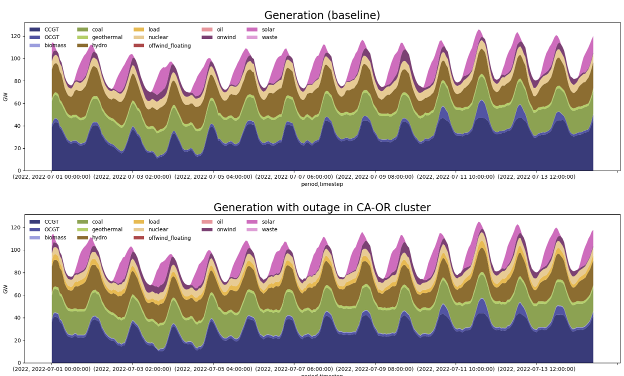

Fig. 9. Generation profiles over time. The bottom graph has a thin golden band denoting load-shedding, which is roughly constant over time.

VI. Discussion

It is sometimes difficult to tell whether our results are truly signal or noise. For instance, if we enable capacity expansion with a historical planning horizon of 2022, the results don't make sense: the marginal prices are negative and dominated by a single bus near Las Vegas, and the generation is almost entirely nuclear. It appears that the source data is not well-suited for backward-looking capacity expansion. We disable capacity expansion for the remainder of our analysis.

We also see suspicious results on a California-only network: massive load-shedding in all the buses on the border of the state. We think this is because California's energy supply cannot match its own demand; the state is reliant on imports, so an isolated model of the state is highly erroneous. (We could fix this in the future by using PyPSA-USA's `focus_weights` parameter to create a full WECC model with 80% of the buses in California. But for now, we just look at all of WECC.)

We think the DC linear power flow assumptions (uniform voltage magnitude, small phase angles, and no line resistance) are likely not contributing to significant model error. These assumptions have been well-validated for high-voltage transmission systems [4], and the PyPSA-USA source data is better calibrated for DC power flow than for AC, according to its maintainers.

Our fire model results also seem intuitively reasonable given basic knowledge of US climate conditions.

Thus, while there are major sources of model error, we are reasonably confident that our model results are at least directionally correct. As k increases, load-shedding and system costs increase. And our model's geographic behavior also makes intuitive sense. Northern California has a high population density and industrial base that consumes a lot of power; if we simulate an outage in the California-Oregon cluster, load-shedding goes up significantly in that region.

The contrast between security-constrained OPF and contingency OPF also makes sense, since security-constrained OPF imposes a weaker constraint on the system than contingency OPF. Both analyses can be useful. Security-constrained analysis helps the operator make a slight adjustment to their OPF, incorporating prior knowledge of fire risk. Contingency analysis tells the operator what to do when a fire has actually occurred.

It's worth noting that operators typically don't run contingency analysis for $k > 20$ lines, because that sort of highly-unlikely event is way out of the model's distribution. The further away we are from the model's baseline, the larger a grain of salt we have to take the results with.

VII. Conclusion / Next Steps

We integrate historical wildfire data into an OPF model to simulate the effect of outages on system stability, performance, and cost. While our model serves useful insights, it contains many simplifying assumptions that would need to be revisited in order to productionize it for use on the control room floor.

Firstly, we are limited by computational constraints. With more compute, we could definitely improve the model accuracy by running with a higher spatial granularity (i.e. 400 instead of 200 buses). We could also run at longer time scales. And we could run capacity expansion planning (which significantly increases the number of variables) to project future additions of renewables and storage.

Secondly, we could have a better framework for balancing fire-risk concerns against cost concerns, either by adding fire-risk constraints into the OPF model or adding a fire-risk term into the objective function.

Thirdly, our fire model is extremely simplified. It doesn't capture the dynamic evolution of a fire through space and time. It doesn't capture the intensity of a fire, which affects the likelihood of damaging a neighboring power line. And it doesn't touch other important climate risks like heat waves, flash floods, and wind storms.

Lastly, there are other types of rare disruptions for which real-time analysis is more important than power flow. Frequency disturbances and harmonic oscillations play a major role in outages like the Iberian peninsula blackout of April 2025 [11]. Power flow is a quasi-static model, and it cannot model transient effects at the sub-second scale [12].

The power grid is one of humanity's most complex systems, and no single model can capture all of that complexity. (As George Box said, all models are wrong but some are useful.) The limitations in our current model point to promising directions for future research.

VIII. Acknowledgements

We thank Tao Sun for providing the fire dataset and brainstorming approaches to modeling fire risk. We thank Kamran Tehranchi for extensive help debugging the various quirks of PyPSA-USA. We thank Brad Rittenhouse from the Sherlock team for support with Farmshare cluster computing. Lastly, we thank Prof. Rajagopal and the CEE272r teaching staff for a great quarter!

References

[1] California Independent System Operator. 2022–2026 Strategic Plan. 16 May 2022, https://www.caiso.com/documents/2022-2026-strategic-plan.pdf.

[2] Karagiannakis, G., Panteli, M., and Argyroudis, S. "Fragility Modeling of Power Grid Infrastructure for Addressing Climate Change Risks and Adaptation." arXiv, 25 Apr. 2025, arXiv:2504.20056. https://arxiv.org/abs/2504.20056.

[3] Hamid R. Sayarshad, Romina Ghorbanloo, Evaluating the resilience of electrical power line outages caused by wildfires, Reliability Engineering & System Safety, Volume 240,

[4] P. Sood, D. J. Tylavsky and Y. Qi, "Improved dc network model for contingency analysis," 2014 North American Power Symposium (NAPS), Pullman, WA, USA, 2014, pp. 1-6, doi: 10.1109/NAPS.2014.6965414. https://ieeexplore.ieee.org/document/6965414 2023, 109588, ISSN 0951-8320, https://doi.org/10.1016/j.ress.2023.109588. (https://www.sciencedirect.com/science/article/pii/S0951832023005021)

[5] T. Brown, J. Hörsch, and D. Schlachtberger, "PyPSA: Python for power system analysis," J. Open Res. Softw., vol. 6, no. 1, p. 4, Jan. 2018, doi: 10.5334/jors.188.

[6] Tehranchi, Kamran and Barnes, Trevor and Frysztacki, Martha and Azevedo, Ines. PyPSA-USA: A flexible open-source energy system model and optimization tool for the United States (February 07, 2025). Available at SSRN: https://ssrn.com/abstract=5029120 or http://dx.doi.org/10.2139/ssrn.5029120

[7] Bresenham, J. E. "Algorithm for computer control of a digital plotter". IBM Systems Journal. 4 (1): 25–30, 1965. doi:10.1147/sj.41.0025.

[8] Xu, Yixing, et al. "US test system with high spatial and temporal resolution for renewable integration studies." 2020 IEEE Power & Energy Society General Meeting (PESGM). IEEE, 2020. https://arxiv.org/abs/2002.06155

[9] Zhou, Ella, and Trieu Mai. 2021. Electrification Futures Study: Operational Analysis of U.S. Power Systems with Increased Electrification and Demand-Side Flexibility. Golden, CO: National Renewable Energy Laboratory. NREL/TP-6A20-79094. https://www.nrel.gov/docs/fy21osti/79094.pdf.

[10] Stanford Farmshare documentation. Accessed May 2025. https://docs.farmshare.stanford.edu/slurm/

[11] Millard, R., Hall, B., Hancock, A., Bounds, A., Pitel, L., and Johnston, I. "How did Spain's electricity grid collapse?". Financial Times. Apr 29, 2025. https://www.ft.com/content/e922cda3-801d-40df-8455-5d3aeae34288

[12] Von Meier, S. "Curveball: Cognitive Challegnes in Modern Grids." Stanford Smart Grid Seminar. October 14, 2014. https://www.youtube.com/watch?v=ru7FzlQv2GI

[13] Josh Novacheck, Justin Sharp, Marty Schwarz, Paul Donohoo-Vallett, Zach Tzavelis, Grant Buster, and Michael Rosso. "The Evolving Role of Extreme Weather Events in the U.S. Power System with High Levels of Variable Renewable Energy". National Renewable Energy Laboratory. December 2021. https://docs.nrel.gov/docs/fy22osti/78394.pdf

[14] "Adapting Existing Energy Planning, Simulation, and Operational Models for Resilience Analysis". National Renewable Energy Laboratory. February 2020. https://docs.nrel.gov/docs/fy20osti/74241.pdf

[15] Von Meier, S. "Electric Power Systems: A Conceptual Introduction, 2nd Edition". June 2006. Wiley-IEEE Press. ISBN: 978-0-470-03642-6.

[16] Nazir, N., Gupta, N., Farishta, R. (2024). Contingency Analysis in Power System Studies: A Critical Review. In: Tomar, A., Mishra, S., Sood, Y.R., Kumar, P. (eds) Proceedings of 4th International Conference on Machine Learning, Advances in Computing, Renewable Energy and Communication. MARC 2023. Lecture Notes in Electrical Engineering, vol 1231. Springer, Singapore. https://doi.org/10.1007/978-981-97-5227-0_9.

[17] Sofia Taylor, Gabriela Setyawan, Bai Cui, Ahmed Zamzam, and Line A. Roald. 2023. Managing Wildfire Risk and Promoting Equity through Optimal Configuration of Networked Microgrids. In Proceedings of the 14th ACM International Conference on Future Energy Systems (e-Energy '23). Association for Computing Machinery, New York, NY, USA, 189–199. https://doi.org/10.1145/3575813.3595196