Emissions-aware electricity network expansion planning via implicit differentiation

Anthony Degleris
Department of Electrical Engineering
Stanford University
degleris@stanford.edu

Lucas Fuentes Valenzuela
Department of Electrical Engineering
Stanford University
lucasfv@stanford.edu

Abbas El Gamal
Department of Electrical Engineering
Stanford University
abbas@ee.stanford.edu

Ram Rajagopal
Department of Civil Engineering
Stanford University
ramr@stanford.edu

Abstract

We consider a variant of the classical problem of designing or expanding an electricity network. Instead of minimizing only investment and production costs, however, we seek to minimize some mixture of cost and greenhouse gas emissions, even if the underlying dispatch model does not tax emissions. This enables grid planners to *directly* minimize consumption-based emissions, regardless of whether or not the carbon market incorporates a carbon tax. We solve this problem using gradient descent with *implicit differentiation*, a technique recently popularized in machine learning. To demonstrate the method, we optimize transmission and storage resources on the IEEE 14-bus test network and compare our solution to one generated by standard planning with a carbon tax. Our solution significantly reduces emissions for the same levelized cost of electricity.

1 Introduction

In order to mitigate the escalating climate crisis, electricity networks must be redesigned to enable increased renewable penetration and reduce greenhouse gas emissions. This will involve making long-term structural changes to the grid that may drastically influence day-to-day grid operations. Because electricity grid investments are expensive and infrequent, choosing the right design decisions is a fundamental problem in power systems.

Expansion planning is the problem of optimally designing or modifying an infrastructure network to meet demand growth alongside cost and security objectives. In expansion planning problems, a grid planner optimizes structural variables, e.g., power line capacity, to minimize the sum of the *investment cost*—the cost to modify the network—and the *operation cost*—the cost to operate the grid given the chosen network variables. The latter is determined by solving the *dispatch problem* (also called the production cost model), in which electricity generators are dispatched to meet demand at minimal cost while satisfying physical constraints.

Traditionally, *monetary costs* have been the primary focus of expansion planning studies. Although recent studies have introduced emissions-related penalties in the planning objective [1, 2, 3, 4], these works only model emissions insofar as how they affect monetary costs via a carbon tax or cap. Importantly, this assumes the underlying dispatch problem uses a carbon tax or cap when dispatching generators, which is often not the case in practice. In contrast, real world grid planners may be tasked

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2021.

with modifying the network to drastically reduce emissions (say, to comply with a state mandate), even if the corresponding dispatch model does not tax or cap emissions.

Contribution. We formulate the problem of minimizing a mixture of carbon emissions and monetary costs during planning. We solve this problem using gradient descent via *implicit differentiation*, a popular technique in machine learning for differentiating through optimization problems. We compare our method to planning with a carbon tax on the IEEE 14-bus network, and show that, at the same levelized cost of electricity, emissions are lower when emissions are included in the planning objective rather than through a carbon tax.

1.1 Related work

Expansion planning. Expansion planning problems have been a focus of power systems research for several decades [5, 6]. Much of this research has focused on building better models for electricity grid operation [7, 8, 9, 10, 11], dealing with uncertainty [7, 12, 13, 14], and developing scalable algorithms [15, 16, 17, 18, 19, 20]; see [6] for a comprehensive overview. Recently, several papers have considered expansion planning with carbon emissions by incorporating a carbon tax or cap to the electricity market dispatch process [1, 2, 3, 4] or through carbon equity constraints [21].

Implicit differentiation and bi-level optimization. Implicit differentiation is a method for differentiating the *solution map* $x^*(\theta)$ of a parameterized system of equations $K(x, \theta) = 0$. Implicit differentiation is not new—it goes back at least to Cauchy [22]—but has recently been popularized in machine learning [23, 24, 25]. Since any convex optimization problem can also be formulated as solving a system of nonlinear equations, e.g., the KKT conditions, this technique can be used to differentiate *through* optimization problems. This often occurs in bi-level optimization [26, 27, 28], where an optimization problem is embedded in another problem. In power systems, implicit differentiation has recently been used for sensitivity analysis [29], network discovery [30], and reinforcement learning [24], although applications of the technique can be found at least as far back as [31].

2 Expansion planning

In this section, we describe the expansion planning problem, and its variant with a carbon tax.

Dispatch problem. Given an electricity network parameterized by $\theta \in \mathbf{R}^k$, the *dispatch problem* minimizes electricity cost while satisfying demand, and is given by,

$$\begin{array}{ll} \text{minimize} & \sum_{t=1}^{T} f(u_t; x_t) \\ \text{subject to} & u \in \mathcal{C}(x, \theta). \end{array} \tag{1}$$

where the variables $u = (u_1, \ldots, u_T) \in \mathbf{R}^n$ may represent, for example, the amount of energy to be produced by each generator at time t. The parameters $x = (x_1, \ldots, x_T) \in \mathbf{R}^{Tm}$ are external data, e.g., electricity demand. Physical constraints, e.g., that transmission lines do not violate thermal limits, are represented by the convex set $\mathcal{C}(x, \theta) \subset \mathbf{R}^{Tn}$. Optimization-based dispatch models are common in practice and used to dispatch generators in many major balancing authorities across the United States [32, 33]. We denote the objective function by $J(u) = \sum_t f(u_t; x_t)$ and let $u^*(\theta) : \mathbf{R}^k \to \mathbf{R}^{Tn}$ be the optimal solution of (1) as a function of θ. We call $J(u^*(\theta))$ the *operation cost* of θ.

Expansion planning problem. In the expansion planning problem, we optimize the parameters of an electricity network to minimize cost. For instance, transmission lines capacities can be increased between two nodes in order to reduce congestion. In this problem, we minimize the sum of the *investment cost* of the network, denoted by $Q(\theta)$, and the operation cost $J(u^*(\theta))$. The investment cost represents the monetary cost required to modify or expand the network, for example, by reinforcing transmissions lines or building electricity storage. We assume $Q(\theta)$ has a simple form, e.g., $Q(\theta) = c^T\theta$. We consider continuous planning problems where the parameters of the network take on continuous values.[1] The *expansion planning problem* is then,

$$\begin{array}{ll} \text{minimize} & Q(\theta) + J(u^*(\theta)) \\ \text{subject to} & \theta \in \mathcal{S}, \end{array} \tag{2}$$

where the variable is $\theta \in \mathbf{R}^k$, and $\mathcal{S} \subset \mathbf{R}^k$ is a convex set. Problem (2) can be formulated as a joint optimization problem over u and θ, then solved with standard optimization techniques. Often, the bi-level structure of (2) is exploited in the solution method [34]. For example, fixing θ and solving the dual of (1) can be used to generate a lower bound on $J(u)$ (this method is known as *Bender's decomposition*; [6, 15, 35]).

Carbon tax. Let the emissions at time t be given by $h(u_t; x_t)$. Often, h is modeled with a simple function, e.g., $h(u_t; x_t) = d^T u$. Then the dispatch problem with a carbon tax of rate λ would correspond to solving (1), only with the objective $\sum_t f(u_t; x_t) + \lambda \sum_t h(u_t; x_t)$. Previous papers [1, 2, 3, 4] have considered expansion planning when the dispatch problem includes a carbon tax (or cap) as a tool for reducing emissions.

3 Emissions-aware expansion planning

In the emissions-aware expansion planning problem, we directly design the network so that the dispatch process naturally results in lower emissions. Let $E(u) = \sum_t h(u_t^*(\theta); x_t)$ be the total emissions produced by the network. Then, the *emissions-aware expansion planning problem* is,

$$\begin{array}{ll} \text{minimize} & Q(\theta) + J(u^*(\theta)) + \lambda E(u^*(\theta)) \\ \text{subject to} & \theta \in \mathcal{S}, \end{array} \tag{3}$$

where the variable is $\theta \in \mathbf{R}^k$. We call the parameter $\lambda \geq 0$ the *emissions penalty*, which determines how much we weight emissions relative to monetary costs. We can interpret λ as the monetary value of reducing one unit of emissions; for example, $\lambda = 200$ means we price a metric ton of CO2 at \$200.

Comparison to planning with a carbon tax. By adding a carbon tax to the dispatch problem in standard expansion planning, we price emissions during operation. In contrast, the emissions-aware expansion planning *directly* incorporates emissions in the planning objective. Although these methods are not mutually exclusive (we can have an emissions-aware expansion planning problem using a dispatch model with a carbon tax), they represent two distinct approaches to reducing emissions: the former reduces emissions by raising the price of electricity produced by emissions-intensive generators, and the latter by encouraging investment decisions that promote using cleaner resources.

3.1 Implicit gradient descent

Unlike (2), problem (3) cannot be converted to a joint optimization problem. Specifically, since in (2) the only dependence on $u^*(\theta)$ is through the objective of (1), we can replace $u^*(\theta)$ with the variable u and still attain the same solution. However, we cannot do the same for $E(u^*(\theta)) = \sum_t h(u_t^*(\theta); x_t)$, since $u^*(\theta)$ is *not* the result of minimizing $E(u)$. Intuitively, the challenge is that (1) and (3) are attempting to optimize for different parameters. This is consistent with the real world—the goal of the dispatch problem is to minimize monetary costs, but network designers often wish to consider additional emissions-related objectives (for example, because of a state-issued mandate).

Instead, we use a projected gradient descent algorithm to solve (3). We first initialize $\theta^{(0)}$ randomly. Then, at each iteration, we solve (1) with $\theta = \theta^{(k)}$ and update the network parameters to be

$$\theta^{(k+1)} := \mathbf{proj}_{\mathcal{S}} \left(\theta^{(k)} - \eta_k \nabla L(\theta) \right),$$

where η_k is the step size at iteration k and $L(\theta) = Q(\theta) + J(\theta) + \lambda E(\theta)$. The projection operator $\mathbf{proj}_{\mathcal{S}}$ enforces the constraints on θ. Using the chain rule, one can show that $\nabla L(\theta) = \nabla Q(\theta) + Du^*(\theta)^T \left(\nabla_u J(u^*(\theta)) + \lambda \nabla_u E(u^*(\theta)) \right)$, where $Du^*(\theta) \in \mathbf{R}^{nT \times k}$ is the Jacobian of the solution map $u^*(\theta)$. The main technical challenge lies in computing the Jacobian of the solution map, since it

[1] Often, expansion upgrades are modeled as integer variables. Handling integer variables in the emissions-aware expansion planning problem is left as future work.

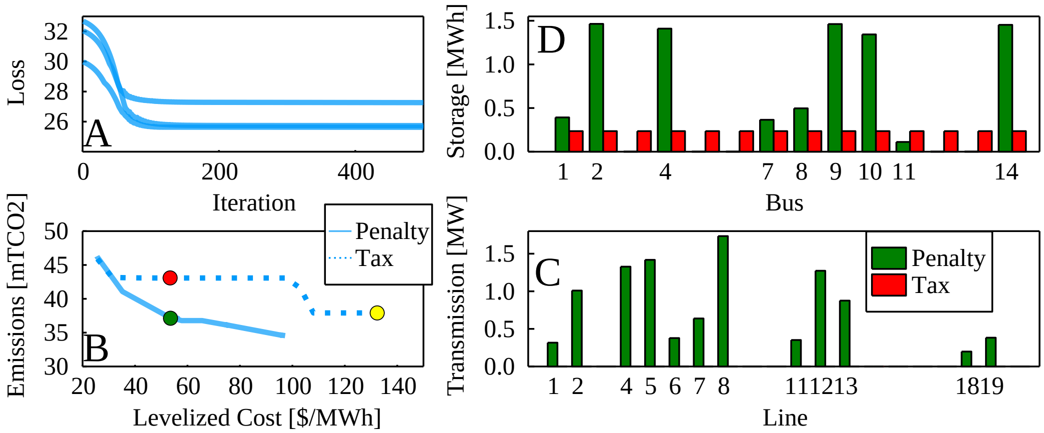

Figure 1: Results from jointly planning line and storage capacities on the IEEE 14-bus network. (A) Objective value of (3) for $\lambda = 500$ at each iteration of projected gradient descent, starting from three random initializations. (B) Levelized electricity costs and total emissions for varying emissions prices, using both the penalty and carbon tax model. The green circle is an emissions penalty of $\lambda = 500$, and the red and yellow circles are carbon taxes of $\lambda = 125$ and $\lambda = 500$, respectively. (C) Network modifications after planning with a penalty of $\lambda = 500$ or a tax of $\lambda = 125$. Both solutions achieve a similar levelized cost, but the emissions-aware problem has notably lower emissions.

usually does not have a clear analytical formula. To do this, we note that $u_t^*(\theta)$ is the solution map of a non-linear system of equations in u and θ (namely, the KKT conditions of (1)). Therefore, we can apply the implicit function theorem to differentiate the solution map and compute $Du_t^*(\theta)$ (as long as the solution to (1) exists uniquely, and under reasonable technical conditions). Additional technical details are available in Appendix A.

4 Numerical experiments

To demonstrate the proposed method, we solve an emissions-aware expansion planning problem (without a carbon tax) on a modified IEEE 14-bus network with 50% renewable penetration. In this problem, the grid planner may upgrade line capacities and add storage resources at each node. We compare our solutions to those generated via standard expansion planning in a dispatch market with a carbon tax.

To make the problem realistic, we use demand and renewable time series from historical California ISO data. We model both line and storage upgrades as continuous variables. For the dispatch problem, we use a dynamic transport model with quadratic generator costs $f(u_t; x_t) = \alpha^T u_t + u_t^T \mathbf{diag}(\beta) u_t$. We model generator emissions as linear, $h(u_t; x_t) = \gamma^T u_t$. Storage resources are optimally dispatched to minimize total cost. We use a linear investment cost $Q(\theta) = c^T \theta$ for upgrading transmission and storage capacity. Additional details are available in Appendix B.

Results. We solve the emissions-aware expansion planning problem with 20 distinct emissions penalties ranging in $\lambda \in [0, 2000]$. Similarly, we solve the standard expansion planning problem using a dispatch model with 20 distinct carbon taxes ranging in $\lambda \in [0, 500]$. Results are displayed in Figure 1. In general, the emissions-aware problem is not convex, and different initializations may converge to different solutions, as shown in Panel A. Notably, as displayed in Panel B, standard planning in the carbon tax setting leads to a significantly higher emissions rate per levelized cost of electricity.[2] This is because the emissions-aware problem directly reduces emissions through additional grid investment (see Panel C), whereas the standard problem with a carbon tax simply raises the price of electricity from high-emitting generators, passing on electricity costs to consumers.

Discussion. These preliminary experiments highlight two important points. First, even in electricity markets without carbon taxes, grid planners can significantly reduce emissions by intelligently

[2]Levelized cost is the total monetary cost (investment plus operation) divided by the total generation.

investing in transmission and storage capacity. Second, because a carbon tax raises the price of electricity, the emissions reduction from grid investment often outperforms that from a tax at the same levelized cost of electricity. In practice, we expect a grid planner to use a mix of both policies (the revenue from a carbon tax, for example, could be used to fund investments that reduce emissions).

5 Conclusion

We introduce the emissions-aware expansion planning problem and solve it with projected gradient descent using implicit differentiation. Our empirical results suggest that directly incorporating emissions during planning outperforms standard planning in markets with a carbon tax. The proposed method could be used by grid planners to make grid expansion decisions that directly reduce greenhouse gas emissions. Future work includes incorporating uncertainty and integer constraints.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0021110.

References

[1] Vahid Asgharian and Morad Abdelaziz. A linear programing model for coordinated low-carbon generation and transmission expansion planning. In *IEEE Electrical Power and Energy Conference*, pages 1–6, 10 2017.

[2] Russell Bent and G Loren Toole. Grid expansion planning for carbon emissions reduction. In *IEEE Power and Energy Society General Meeting*, pages 1–8, 7 2012.

[3] Fabian Neumann and Tom Brown. Heuristics for Transmission Expansion Planning in Low-Carbon Energy System Models. In *International Conference on the European Energy Market*, pages 1–8, 9 2019.

[4] Jing Qiu, Zhaoyang Dong, Junhua Zhao, Ke Meng, Fengji Luo, Kit Po Wong, and Chao Lu. A low-carbon oriented probabilistic approach for transmission expansion planning. *Journal of Modern Power Systems and Clean Energy*, 3(1):14–23, 3 2015.

[5] Mario V F Pereira and Leontina M V G Pinto. Application Of Sensitivity Analysis Of Load Supplying Capability To Interactive Transmission Expansion Planning. *IEEE Transactions on Power Apparatus and Systems*, PAS-104(2):381–389, 2 1985.

[6] Sara Lumbreras and Andrés Ramos. The new challenges to transmission expansion planning. Survey of recent practice and literature review. *Electric Power Systems Research*, 134:19–29, 5 2016.

[7] Zhenyu Zhuo, Ershun Du, Ning Zhang, Kang Chongqing, , Qing Xia, and Zhidong Wang. Incorporating Massive Scenarios in Transmission Expansion Planning With High Renewable Energy Penetration. *IEEE Transactions on Power Systems*, 35(2):1061–1074, 3 2020.

[8] Mohammad Majidi-Qadikolai and Ross Baldick. Integration of N-1 Contingency Analysis With Systematic Transmission Capacity Expansion Planning: ERCOT Case Study. *IEEE Transactions on Power Systems*, 31(3):2234–2245, 5 2016.

[9] Zechun Hu, Fang Zhang, and Baowei Li. Transmission expansion planning considering the deployment of energy storage systems. In *IEEE Power and Energy Society General Meeting*, pages 1–6. ieeexplore.ieee.org, 7 2012.

[10] Lina P Garces, Antonio J Conejo, Garcia-Bertrand Raquel, and Rubén Romero. A Bilevel Approach to Transmission Expansion Planning Within a Market Environment. *IEEE Transactions on Power Systems*, 24(3):1513–1522, 8 2009.

[11] Isidro G Sanchez, Ruben Romero, Jose R S Mantovani, and Marcos J Rider. Transmission-expansion planning using the DC model and nonlinear-programming technique. *IEE Proceedings-Generation, Transmission and Distribution*, 152(6):763–769, 2005.

[12] Vahid Khaligh and Amjad Anvari-Moghaddam. Stochastic expansion planning of gas and electricity networks: A decentralized-based approach. *Energy*, 186:115889, 11 2019.

[13] Birger Mo, Jan Hegge, and Ivar Wangensteen. Stochastic generation expansion planning by means of stochastic dynamic programming. *IEEE Transactions on Power Systems*, 6(2):662–668, 5 1991.

[14] Alireza Zare, C Y Chung, Junpeng Zhan, and Sherif Omar Faried. A Distributionally Robust Chance-Constrained MILP Model for Multistage Distribution System Planning With Uncertain Renewables and Loads. *IEEE Transactions on Power Systems*, 33(5):5248–5262, 9 2018.

[15] Mohammad Majidi-Qadikolai and Ross Baldick. A Generalized Decomposition Framework for Large-Scale Transmission Expansion Planning. *IEEE Transactions on Power Systems*, 33(2):1635–1649, 3 2018.

[16] Ibrahim Alhamrouni, Azhar Khairuddin, Ali Ferdavani, Khorasani, and Mohamed Salem. Transmission expansion planning using AC-based differential evolution algorithm. *IET Generation, Transmission and Distribution*, 8(10):1637–1644, 2014.

[17] Joshua A Taylor and Franz S Hover. Linear Relaxations for Transmission System Planning. *IEEE Transactions on Power Systems*, 26(4):2533–2538, 11 2011.

[18] Natalia Alguacil, Alexis L Motto, and Antonio J Conejo. Transmission expansion planning: a mixed-integer LP approach. *IEEE Transactions on Power Systems*, 18(3):1070–1077, 8 2003.

[19] Silvio Binato, Gerson C Oliveira, and Joao L Araujo. A greedy randomized adaptive search procedure for transmission expansion planning. *IEEE Transactions on Power Systems*, 16(2):247–253, 5 2001.

[20] Javier Contreras and Felix F Wu. A kernel-oriented algorithm for transmission expansion planning. *IEEE Transactions on Power Systems*, 15(4):1434–1440, 11 2000.

[21] Yanlong Sun, Chongqing Kang, Qing Xia, Qixin Chen, Ning Zhang, and Yaohua Cheng. Analysis of transmission expansion planning considering consumption-based carbon emission accounting. *Applied Energy*, 193:232–242, 2017.

[22] Asen L Dontchev and R Tyrrell Rockafellar. *Implicit Functions and Solution Mappings: A View from Variational Analysis*. Springer, New York, NY, 2014.

[23] Brandon Amos and J Zico Kolter. Optnet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning*, pages 136–145, 2017.

[24] Bingqing Chen, Priya Donti, Kyri Baker, J Zico Kolter, and Mario Berges. Enforcing Policy Feasibility Constraints through Differentiable Projection for Energy Optimization. *arXiv preprint arXiv*, 2021.

[25] Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and Zico Kolter. Differentiable convex optimization layers. *arXiv preprint arXiv:1910.12430*, 2019.

[26] Ankur Sinha, Pekka Malo, and Kalyanmoy Deb. A Review on Bilevel Optimization: From Classical to Evolutionary Approaches and Applications. *IEEE Transactions on Evolutionary Computation*, 22(2):276–295, 2018.

[27] Stephen Gould, Basura Fernando, Anoop Cherian, Peter Anderson, Rodrigo Santa Cruz, and Edison Guo. On Differentiating Parameterized Argmin and Argmax Problems with Application to Bi-level Optimization. (1):1–16, 2016.

[28] Shane T. Barratt and Stephen P. Boyd. Least squares auto-tuning. *Engineering Optimization*, pages 1–28, 2020.

[29] Fengyu Zhou, James Anderson, and Steven H Low. The Optimal Power Flow Operator: Theory and Computation. *IEEE Transactions on Control of Network Systems*, page 1, 2020.

[30] Priya L Donti, Inês Lima Azevedo, and J Zico Kolter. Inverse Optimal Power Flow: Assessing the Vulnerability of Power Grid Data. *phys2018.csail.mit.edu*, 2018.

[31] Paul R Gribik, Dariush Shirmohammadi, Arthur D Little, Shangyou Hao, and Chifong L Thomas. Optimal Power Flow Sensitivity Analysis. *IEEE Transactions on Power Systems*, 5(3):969–976, 1990.

[32] PJM. PJM Manual 12: Balancing Operations. 2021.

[33] United States Department of Energy. The Value of Economic Dispatch. Technical report, 2005.

[34] Xiaohu Zhang, Di Shi, Zhiwei Wang, Zhe Yu, Xinan Wang, Desong Bian, and Kevin Tomsovic. Bilevel optimization based transmission expansion planning considering phase shifting transformer. *North American Power Symposium*, 2017.

[35] Mario V F Pereira, Leontina M V G Pinto, Sergio H F Cunha, and Gerson C Oliveira. A Decomposition Approach To Automated Generation/Transmission Expansion Planning. *IEEE Transactions on Power Apparatus and Systems*, PAS-104(11):3074–3083, 11 1985.

[36] Akshay Agrawal, Shane Barratt, Stephen Boyd, Enzo Busseti, and Walaa M Moursi. Differentiating Through a Cone Program. 4 2019.

A Implicit differentiation

The *implicit function theorem* is a classical theorem from analysis. The theorem roughly states that given a system of equations $K(x, \theta) : \mathbf{R}^n \times \mathbf{R}^d \to \mathbf{R}^n$ and a point $(x_0, \theta_0) \in \mathbf{R}^{n+d}$ such that

- $K(x_0, \theta_0) = 0$,
- the partial Jacobian $D_x K(x_0, \theta_0) \in \mathbf{R}^{n \times n}$ is invertible,

then

- there exists a function $x^*(\theta) : \mathbf{R}^d \to \mathbf{R}^n$ such that $x^*(\theta_0) = x_0$ and $K(x^*(\theta), \theta) = 0$ in a local region around x_0
- the Jacobian of $x^*(\theta)$ is given by $Dx^*(\theta) = -D_x K(x^*(\theta), \theta)^{-1} D_\theta K(x^*(\theta), \theta)$.

The above theorem also requires several technical conditions, e.g., that $K(x, \theta)$ is differentiable. There are a plethora of references available on the implicit function theorem, and we refer the reader to [22] for additional details.

In the context of optimization, the implicit function theorem allows us to compute sensitivities of a *parametrized* convex optimization problem, i.e., an optimization problem that changes according to some real variables. This comes from the fact that the solution $x^*(\theta)$ to a parametrized optimization problem can be viewed as the solution to a system of non-linear equations $K(x^*(\theta), \theta) = 0$—namely, the KKT conditions of the problem. One can also use other fixed point conditions that guarantee optimality as the non-linear system. Notably, the condition that $K(x_0, \theta_0) = 0$ corresponds to the *existence* of a solution, and the condition that $D_x K(x_0, \theta_0)$ is invertible corresponds to the *uniqueness* of a solution. To learn more, we recommend the reader refer to [25, 36], or to [22] for a more detailed analysis.

B Experiment details

In this section, we describe our experimental setup in detail for completeness and reproduciblity. All code was written in Julia and will be made publicly available after submission.

Dispatch problem. We use a dynamic transport model for the dispatch problem, given by,

$$
\begin{array}{lll}
\text{minimize} & Z \sum_{t=1}^{T} \left(\alpha^T g_t + \frac{1}{2} g_t^T \mathbf{diagm}(\beta) g_t \right) & \\
\text{subject to} & A p_t = d_t - B g_t + \delta_t^c - \delta_t^d, & t = 1, \ldots, T, \\
& s_t = s_{t-1} + \eta_c \delta_t^c + (1/\eta_d) \delta_t^d, & t = 1, \ldots, T, \\
& 0 \le g_t \le g_{\max}, & t = 1, \ldots, T, \\
& |p_t| \le p_{\max}, & t = 1, \ldots, T, \\
& |\delta_t^c| \le \rho C, & t = 1, \ldots, T, \\
& |\delta_t^d| \le \rho C, & t = 1, \ldots, T, \\
& |s_t| \le C, & t = 1, \ldots, T,
\end{array}
\tag{4}
$$

with variables $(g_t, p_t, \delta_t^c, \delta_t^d, s_t) \in \mathbf{R}^{\ell+m+3n}$ for $t = 1, \ldots, T$. All inequalities are to be read elementwise. The parameters of the problem are the problem scaling $Z \in \mathbf{R}$, the time horizon $T \in \mathbf{Z}_+$, the charge and discharge efficiencies $\eta_c, \eta_d \in \mathbf{R}$, the storage (dis)charging rate $\rho \in \mathbf{R}$, the generation costs $\alpha, \beta \in \mathbf{R}^\ell$, the generation capacities $g_{\max} \in \mathbf{R}^\ell$, the line capacities $p_{\max} \in \mathbf{R}^m$, the storage capacity C, the network incidence matrix $A \in \mathbf{R}^{n \times m}$, and the node-generator map $B \in \mathbf{R}^{n \times \ell}$. We set the time horizon to $T = 24$ and the problem scaling to $Z = 5 \times 365 = 1825$, i.e., we simulate the grid hourly for 5 years. Note that this makes the assumption that every day of operation is the same. Although this is not the case in practice, developing and sampling grid operation scenarios is not the focus of this work, and many papers [7, 12] have already addressed this question.

Expansion planning variables. We let $\theta = (p_{\max}, C)$.

Parameters. We set the charge and discharge efficiencies to $\eta_c = \eta_d = 0.95$. The storage (dis)charge rate is set to $\rho = 0.25$. These parameters are chosen to be reasonable values for grid scale lithium-ion battery storage. We give each of the m lines a random length $a_i \in [40, 50]$ miles. We also estimate the cost per mile MWh as $a_0 = 3,000$, We set the cost of storage to be $b_0 = 350,000$ per MWh. The total investment cost is then $Q(\theta) = c^T\theta$, where $c = (a_0 a, b_0 \mathbf{1}_n)$. We set the emissions rates of the six generators to $h(g_t) = v^T g_t$, where $v = (0.35, 0.45, 1.1, 1.2, 1.3)$. These values were chosen to reflect likely fuel source of each generator, given their costs and capacities.

Network data. We use a modified version of the IEEE 14-bus network and parse the data using `PowerModels.jl`. The case file lacks thermal line limits, so `PowerModels.jl` sets these automatically. Generator prices α, β, generator and line capacities $g_{\max}, p_{\max}$, and the network and generator topologies A, B all come from this file. We also load nodal demands $\tilde{d} \in \mathbf{R}^n$, which we use as the peak demands for the network. Nodal demand is scaled by 2.5 to simulate demand growth, and nodal generation is also scaled so that $\mathbf{1}^T g_{\max} = \mathbf{1}^T \tilde{d}$.

Demand and renewable time series. We use demand and renewable data from CAISO, available for free online from their OASIS portal (`http://oasis.caiso.com/mrioasis`). We extract regional demand and solar hourly time series from July 1st, 2021. We normalize each of the $K = 5$ demand time series $d^{\text{rel},k} \in \mathbf{R}^T$ so that $\max(d^{\text{rel},k}) = 1$ and each of the $R = 3$ solar time series $u^{\text{rel},r}$ so that $\mathbf{1}^T u^{\text{rel},r} = 1$. Each node i is given a demand profile k_i and a solar profile r_i, both selected uniformly at random. Then the net demand at time t is set to be $(d_t)_i = \tilde{d}_i d_t^{\text{rel},k_i} - \omega \mathbf{1}^T(\tilde{d}_i d^{\text{rel},k_i}) u^{\text{rel},r_i}$, where $\omega = 0.5$ is the renewable penetration. In other words, demand at each node is the product of the normalized CAISO demand time series with the peak demand at the node, minus the current solar generation. At 50% renewable penetration, 50% of total demand is met by renewable resources. However, since this is spread out non-uniformly throughout the day, some nodes have very high or very low (even negative) net demand during some time periods.