ITERATIVE PORTFOLIO OPTIMIZATION

An essential tool for reliable and clean electricity planning

DECEMBER 2024

Elaine Hart
Sylvan Energy Analytics

ACKNOWLEDGEMENTS

The author would like to thank Ana Mileva, Priya Sreedharan, Taylor McNair, Ric O'Connell, and Gwen Spencer for input and feedback throughout the development of this paper; and Annie Dore of BeeSpring Designs for graphic design of this report.

SYLVAN ENERGY ANALYTICS

GridLAB

INTRODUCTION | LIMITATIONS OF CURRENT CAPACITY EXPANSION MODELS

Capacity expansion models are the primary method for developing resource portfolios in electricity planning exercises, including integrated resource plans and transmission plans. These models are complex and imperfect, primarily because investments in the electricity system have operational impacts that are both large in scale and increasingly require high resolution analysis to understand. Despite advances in computing, it remains computationally intractable to bring all of the complexity of the electricity system operations at high temporal and geographical resolution and across a wide range of conditions directly into the capacity expansion problem. And yet, utilities are increasingly expected to account for that complexity in their plans to meet reliability requirements and policy objectives at lowest cost.

In recent years, new approximations in capacity expansion models, such as effective load carrying capability (ELCC) curves and flexibility value adders have meaningfully advanced our ability to address increasing shares of renewables and energy storage within electricity planning. However, with increasingly complex capacity expansion models, new challenges have emerged. Planners may put too much trust in the optimality of portfolios built with capacity expansion models and the increasing complexity of the models makes it challenging for regulators and stakeholders to effectively scrutinize them. From a practical perspective, planning processes are constrained by both time and human resources. Ultimately, the significant effort put toward developing and scrutinizing increasingly complex capacity expansion model formulations is time and energy not spent engaging with the critical planning questions that we face in this environment of rapid change and vast uncertainty.

ITERATIVE PORTFOLIO OPTIMIZATION AT A GLANCE

In light of the limitations of capacity expansion models, planners are turning to roundtrip modeling to develop and evaluate portfolios that account for the most complex planning considerations, including resource adequacy and policy compliance. The term roundtrip modeling often refers to a two-step process of running a capacity expansion model to develop a portfolio and then running the portfolio through an evaluation model, for example a production cost or a resource adequacy model, to evaluate the performance of the portfolio with more accuracy and granularity. In some cases, roundtrip modeling may reveal that the portfolio does not achieve a planning constraint, for example a reliability or policy requirement. In this case, the planner may make adjustments to the starting portfolio — typically deciding on a limited set of resources to add to the portfolio — until it meets the planning constraint. While this approach ensures that the planning constraints are met, round trip modeling alone does not ensure that the resulting portfolio is least cost. Portfolio costs will depend on how the planner decides to adjust the portfolio to meet the planning constraints. So while roundtrip modeling can provide more precise information about the performance of a portfolio and help to ensure portfolios meet planning constraints, planners, stakeholders, and regulators are left wondering if the process may have missed lower cost solutions.

FIGURE 1. High level schematic of iterative portfolio optimization

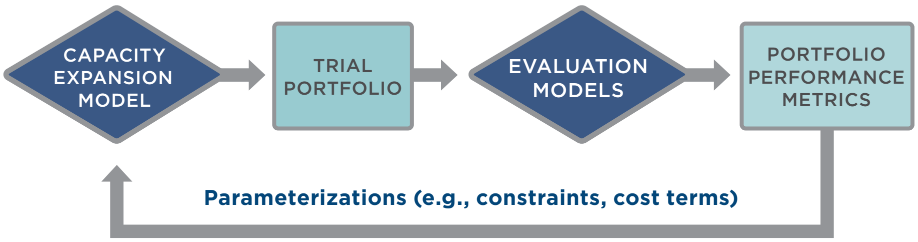

Iterative portfolio optimization is a type of roundtrip modeling in which mathematical rules are designed to systematically update the portfolio in each iteration in a manner that ultimately converges to the optimal portfolio that meets all planning constraints. Iterative portfolio optimization removes subjectivity from the process of updating portfolios and reduces the amount of detail that must be explicitly represented in the capacity expansion problem, removing the need for bespoke formulations to capture complex planning considerations. In an iterative optimization framework, capacity expansion models are not called upon to identify optimal portfolios, but simply to produce reasonable guesses, or trial portfolios, and to incorporate information from evaluation models to improve upon those guesses in each iteration. With iterative optimization, optimality is not claimed based on the cleverness of the capacity expansion model formulation; it is instead demonstrated through the relative performance of the trial portfolios as they converge toward the optimal solution.

A GENERALIZED APPROACH

In this paper, we present a generalized iterative portfolio optimization approach that can be used to represent a wide range of complex considerations within capacity expansion models, including policy and reliability requirements as well as planning considerations with high geographic or temporal resolution. The approach is unique in that it guarantees convergence for complex constraints or cost terms, provided that they meet some mathematical criteria,[1] so the planner can be certain that the resulting portfolio is the lowest cost portfolio that meets all of the constraints with the same level of accuracy provided by the more complex evaluation models.

This generalized approach to iterative portfolio optimization guarantees optimality (e.g., least cost) and does not require the planner to develop bespoke approximations for complex planning considerations (e.g., ELCC curves and policy compliance metrics) in the capacity expansion problem.

The approach (a type of cutting plane technique) involves bounding each planning parameter (e.g., expected unserved energy, GHG emissions, renewable or clean energy delivered, variable costs, etc.) with a set of linear constraints. At the outset, the planner may begin with little to no information about how each planning parameter is affected by resource decisions. In each iteration, evaluation models are used to calculate the parameter values for the trial portfolio and to derive linear constraints that approximate the planning parameters in the neighborhood of the

1 Specifically, the constraints and cost terms must be convex to guarantee convergence. This requirement is no more restrictive than current methods, which generally require the constraints and cost terms in the capacity expansion problem to be convex.

trial portfolio. The capacity expansion model is then re-run with the additional constraints. As the process continues, the linear constraints more closely approximate the planning parameters in the neighborhood of their optimal values, and the portfolio converges toward the optimal solution.

The process is shown illustratively with one constrained planning parameter, F, and one decision variable, x, in Figure 2. With each iteration, the collection of linear constraints more closely approximates the underlying (unknown) relationship between F and x. While the example shows the approach in one dimension for simplicity, it can be applied to represent any number of functions (or planning parameters) and any number of planning variables (e.g., resource additions) in the capacity expansion problem. The approach can also apply both to planning parameters that are constrained and planning parameters that are penalized with cost terms in the portfolio optimization problem, provided that they are convex.

FIGURE 2. Iterative portfolio optimization technique in one dimension

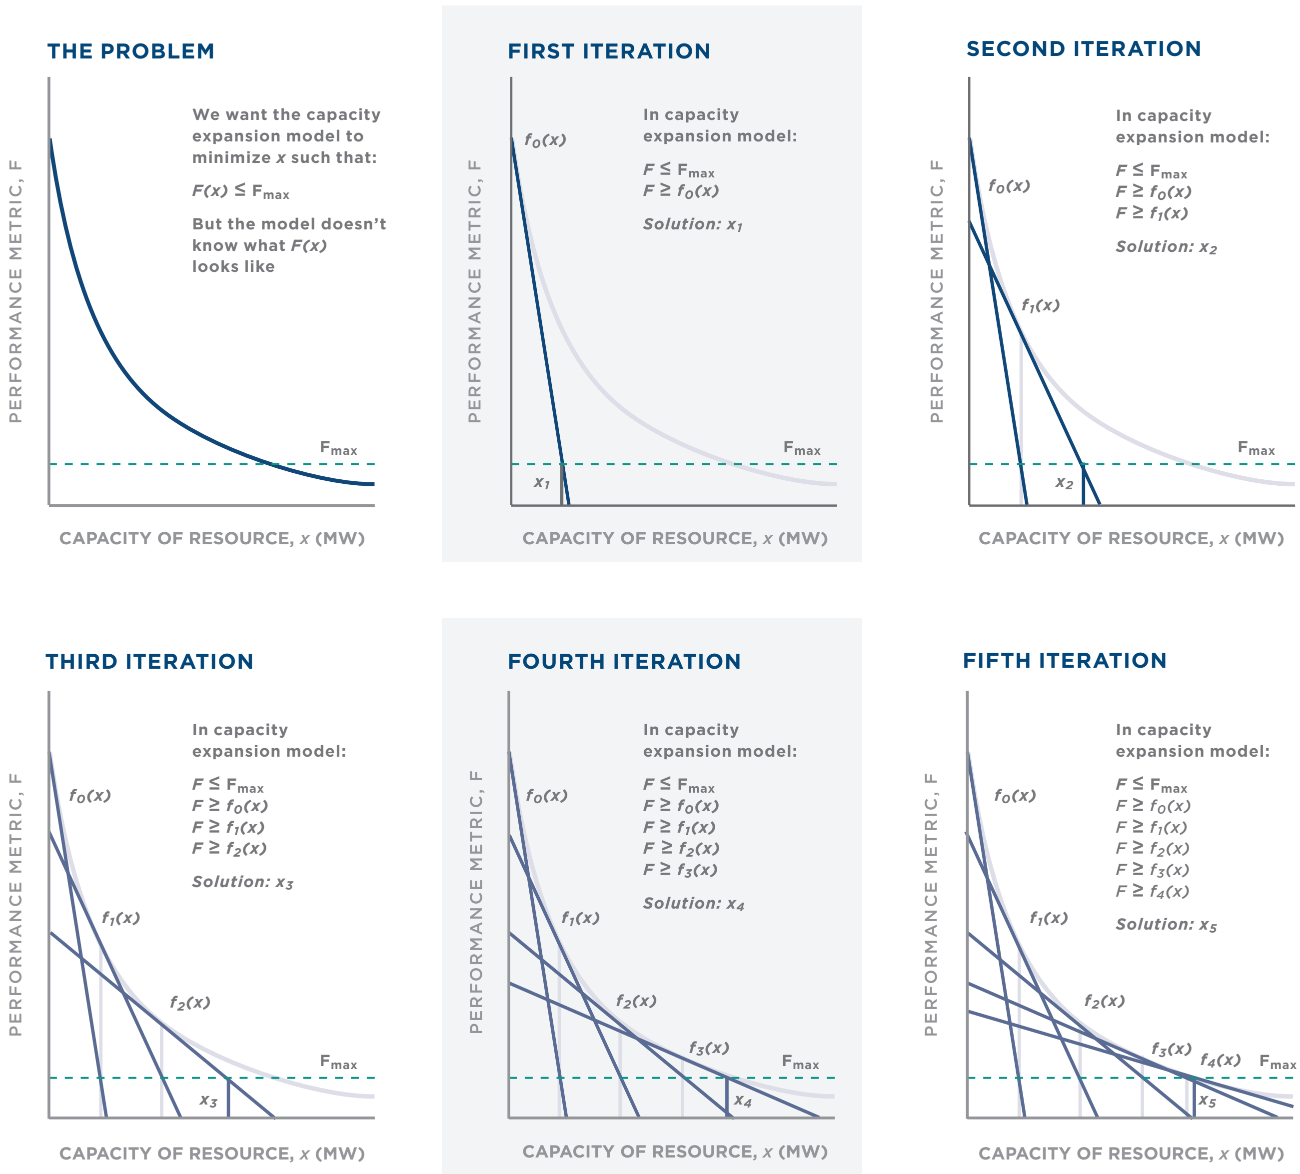

DEMONSTRATION OF ITERATIVE PORTFOLIO OPTIMIZATION FOR A HIGHLY CHALLENGING PLANNING PROBLEM

We demonstrate the generalized iterative optimization approach by solving an especially challenging problem: developing a fully optimized 100% clean portfolio that achieves a resource adequacy requirement without the use of clean firm capacity. The model considers an illustrative representation of the Western United States in 2035,[2] assuming that all fossil fuel resources have been retired. The model is allowed to select from 5 different wind resources, 5 different solar resources, and 5 different storage resources (independently selecting the MW and MWh of storage), resulting in 20 total investment variables.[3] The capacity expansion problem uses 52 representative days to develop the initial portfolio and to estimate lost production tax credits (PTCs) due to renewable curtailment. The algorithm must find the least cost portfolio that achieves a target amount of expected unserved energy (EUE) across 100 years of potential weather and hydro conditions. Both the capacity expansion and resource adequacy models were developed in GridPath[4], Sylvan's open source electricity modeling system.

The algorithm converged to the least cost portfolio that met the target EUE (to within 1%) in 364 iterations. Figure 3 shows the composition and Figure 4 shows the cost[5] and reliability performance of all of trial portfolios. The trial portfolios demonstrate the highly non-linear relationship between cost and reliability (shown in log scale in the figures), which makes this such a challenging problem.

FIGURE 3. Trial portfolio compositions

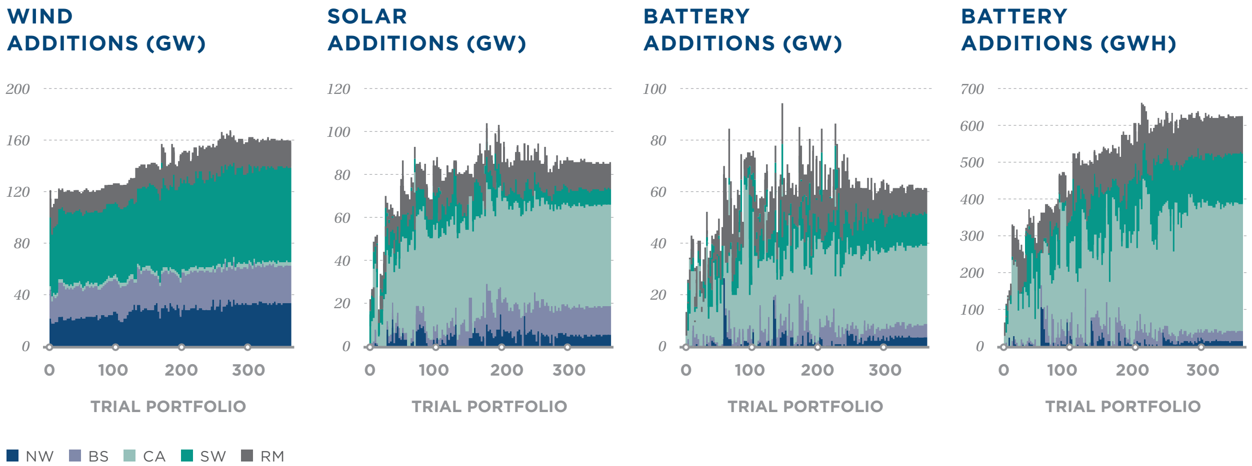

2 Loads and existing and planned non-emitting resources were based on the representation of the Western US in GridLab's Moonshot Study (gridlab.org/moonshot-study)

3 Resource costs were estimated based on NREL's Annual Technology Baseline (ATB) and resource potentials were estimated by screening NREL wind and solar supply curves.

4 www.gridpath.io

5 Portfolio costs include fixed costs associated with selected resources and total variable costs, net of lost PTCs due to curtailment.

FIGURE 4. Trial portfolio cost and reliability performance

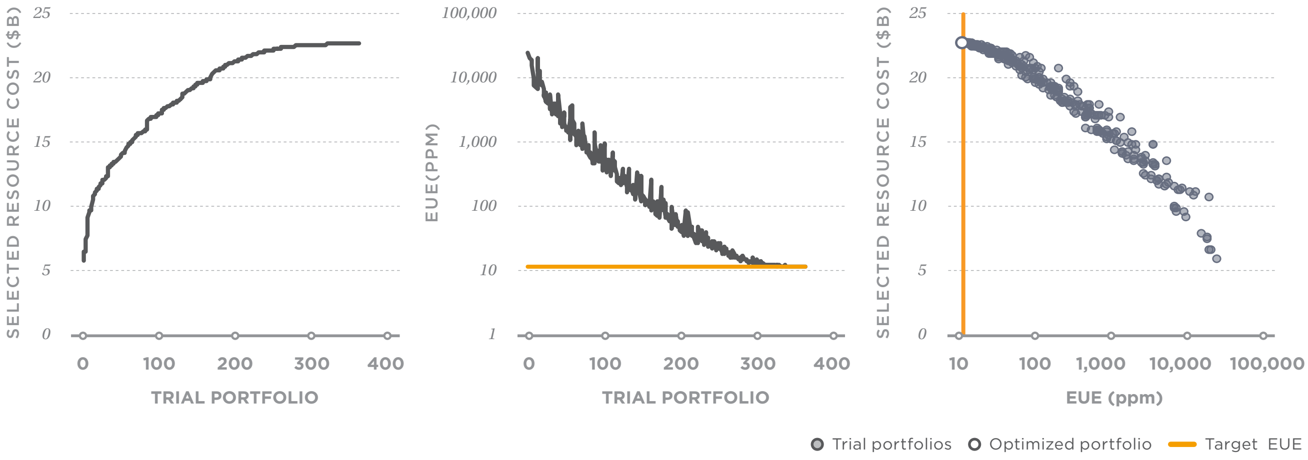

For context, we also developed a portfolio using a more traditional approach – identifying the capacity needed to meet the resource adequacy requirement (in this case the EUE target) and estimating the contribution of each resource option to meeting that capacity need with one-dimensional ELCC curves (curves that relate a resource's capacity contribution to the amount of the resource that is added to the portfolio). We developed ELCC curves for each of the 5 wind and solar resources and for 4-hr and 10-hr battery storage resources and incorporated those curves into the same capacity expansion model that we used to demonstrate the iterative portfolio optimization algorithm. Building the ELCC curves required running 194 resource adequacy simulations in advance of running the capacity expansion model. Despite meeting the capacity requirement in the capacity expansion model, the resulting portfolio was far from meeting the reliability requirement when it was tested in the resource adequacy model — it experienced 1,000 times as much unserved energy as the portfolio developed through iterative portfolio optimization. This result highlights the limitations of single-pass capacity expansion modeling. Because the portfolio does not meet a key planning constraint, a planner would need to add resources to the portfolio to meet reliability requirements and the composition and cost of the final portfolio would depend on the subjective decisions the planner makes on which resources to add. Final portfolio costs would be necessarily higher (or no lower than) the cost of the portfolio developed through iterative portfolio optimization.

The poor performance of the portfolio developing using one-dimensional ELCC curves is largely attributable to the fact that one-dimensional ELCC curves test individual resources one at a time and do not capture interactions between resource options (i.e., portfolio effects). To capture portfolio effects, planners sometimes test combinations of resource options to build multidimensional ELCC surfaces. However, multidimensional ELCC surfaces quickly become impractical as more resource options are considered. For this particular problem, developing a multidimensional ELCC surface to represent portfolio effects between the 20 investment variables would require on the order of 10^{20} resource adequacy simulations and would take about one million times the age of the universe to run.[6] In short, identifying a truly optimal portfolio for this system is not possible with ELCC curves or ELCC surfaces. This problem requires an iterative approach.

6 This assumes 10 increments of additions for each resource and one-hour runtimes for each RA simulation.

ITERATIVE PORTFOLIO OPTIMIZATION FOR MULTI-OBJECTIVE PLANNING

Importantly, iterative portfolio optimization does not relieve the planner from the fundamentally challenging task of balancing competing priorities and planning under uncertainty. What it does provide is more confidence that planning considerations and electricity system dynamics are being accurately accounted for in the development of portfolios and that no stone has been left unturned in searching for the best portfolios that meet the planning criteria.

Iterative portfolio optimization provides confidence that no stone has been left unturned in searching for the best portfolios that meet the planning criteria.

We demonstrate how iterative portfolio optimization can be used to support multi-objective planning by developing three alternative portfolios that have different requirements and apply different weights to key planning metrics, including reliability, GHG emissions, and land use (see Table 1 for the design parameters for each portfolio).

Portfolios A-C all require zero GHG emissions, but Portfolio D allows the model to select new natural gas plants.[7] Evaluating Portfolio D requires more detail in the evaluation model for each trial portfolio, including economic dispatch and GHG emissions accounting across the 100 years of weather and hydro conditions tested for resource adequacy.[8] While we demonstrate the approach by applying a cost penalty to GHG emissions, the same approach can be used to develop GHG-constrained resource adequate portfolios.

TABLE 1. Constraints and penalties applied to develop portfolios[9]

PORTFOLIO	A	B	C	D
EUE	≤ 0.10 customer-hrs/yr	$20,000/MWh	$20,000/MWh	$20,000/MWh
GHGs	= 0 $mtCO_2$	= 0 $mtCO_2$	= 0 $mtCO_2$	~ $200/$mtCO_2$
Direct Land Use	$0/acre	$0/acre	$10,000/acre	$10,000/acre
Indirect Land Use	$0/acre	$0/acre	$2,500/acre	$0/acre

The compositions of the resulting portfolios are shown in Figure 5 and their associated planning metrics are listed in Table 2. This exercise highlights the tradeoffs between the various planning objectives and the options for balancing these objectives. When using iterative portfolio optimization, the planner can be sure that each portfolio best balances the objectives they've set for it and that the differences between portfolios reflect real trade-offs, not artifacts of whatever simplifications have been employed in the capacity expansion model.

7 Recall that all existing natural gas plants were removed from the starting portfolio, so additions represent total natural gas capacity on the system, not incremental capacity relative to the current system.

8 In developing Portfolio D, the evaluation model was also used to calculate the cost impacts of lost PTCs due to curtailment across the 100 years of weather and hydro conditions for each trial portfolio. In Portfolios A-C, lost PTCs were estimated in the capacity expansion problem using representative days.

9 Average customer-hrs/yr are calculated by dividing the EUE by the weather-averaged annual load and multiplying by 8760 hrs/yr. This reflects the average number of hours per year each customer would experience an outage if outages were evenly distributed across customers. 0.10 customer-hrs/yr is equivalent to 11.4 ppm.

FIGURE 5. Optimized portfolio compositions, including existing and planned resources

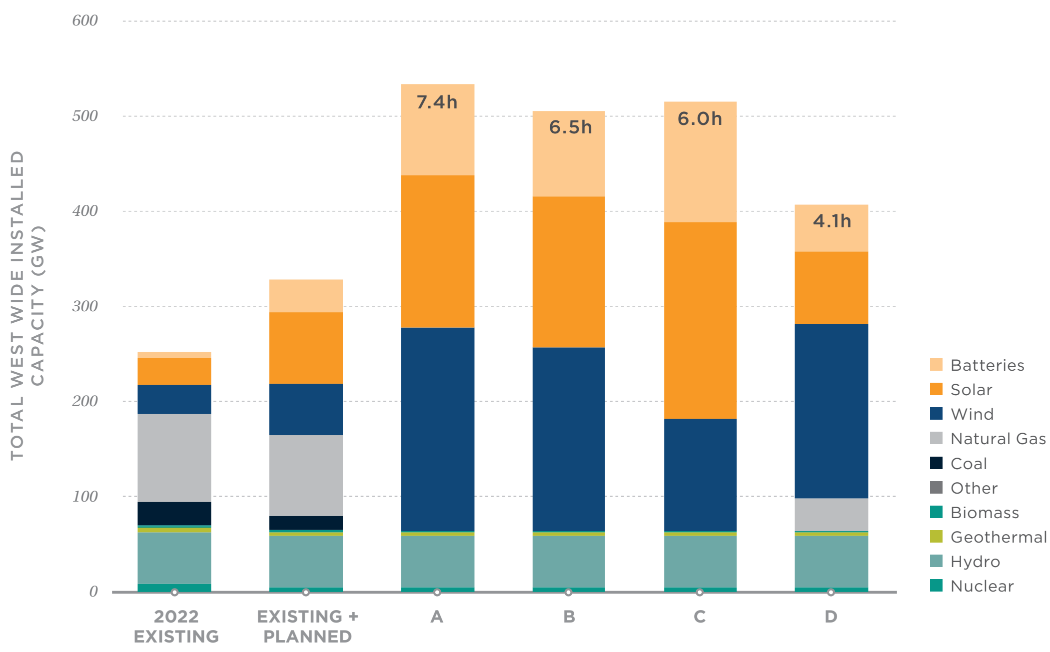

TABLE 2. Optimized portfolio performance metrics

PORTFOLIO	A	B	C	D
Portfolio Cost ($B/yr)	$22.7	$18.7	$21.6	$10.4
EUE (customer-hrs/year)	0.10	1.04	3.01	0.22
GHGs (mmtCO_2)	0.00	0.00	0.00	11.1
Direct Land Use (sq. miles)	1,840	1,761	1,896	1,086
Indirect Land Use (sq. miles)	38,370	34,599	20,981	32,354
Percent Clean (% of generation)	100%	100%	100%	97.5%

NEAR-OPTIMAL PORTFOLIOS

In developing the four optimized portfolios, the algorithm tested 1,075 trial portfolios.[10] These trial portfolios provide the planner, their stakeholders, and regulators with direct visibility into the optimality of the optimized portfolios and a range of near-optimal portfolios. This can provide important context for interpreting the optimal portfolios. As an example, Figure 6 shows the composition of 235 near-optimal portfolios that achieved a total cost (including cost penalties) within 10% of the total cost of Portfolio D. The near-optimal portfolios share quite a bit in common,

10 Some portfolios converged more quickly by "warm-starting" the algorithm with all of the linear constraints that were derived in developing previous portfolios.

including substantial development of wind beyond current plans and much less natural gas capacity than we have today. However, while all near-optimal portfolios select all 3 GW of available California Wind and at least 38 GW of New Mexico Wind, the geographical distribution of the rest of the wind additions vary considerably across the near-optimal portfolios. For this particular demonstration case, which has admittedly simplified assumptions with regard to locational resource costs and potentials, the total amount of wind added to the portfolio appears to matter more than where much of that wind is developed. In a case with more precise locationally-specific information, investigation of near-optimal portfolios in this manner might reveal more robust geographic differences and help planners identify whether or where to expand transmission to access particularly crucial renewable resources.

FIGURE 6. Near optimal portfolios compared to optimal Portfolio D

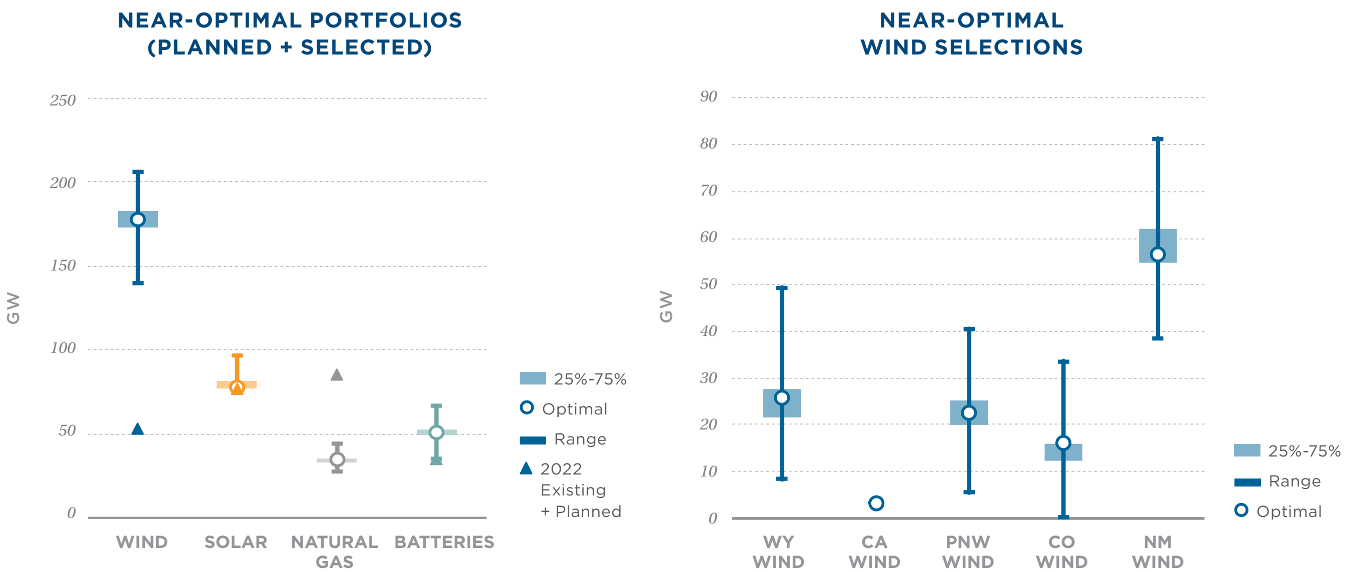

Takeaways and future work

Iterative portfolio optimization is a promising and practical technique for addressing the limitations of capacity expansion models. In this paper, we demonstrate its use for accurately capturing resource adequacy constraints, supporting multi-objective planning exercises, and investigating near-optimal portfolios. While we place specific focus on resource adequacy and greenhouse gas emissions, the approach could be expanded to encompass any number of planning metrics that are not well-suited to capacity expansion formulations, particularly those that require more granularity than is practical within capacity expansion model, including locational considerations, such as transmission constraints, distribution system constraints, and community impacts. The approach could also be used to solve for decision variables beyond generation and storage additions, including resource retirements, demand side management programs, and transmission expansion, with minimal changes to the capacity expansion formulation.

FIGURE 7. Iterative portfolio optimization scheme that considers multiple parameterized planning metrics

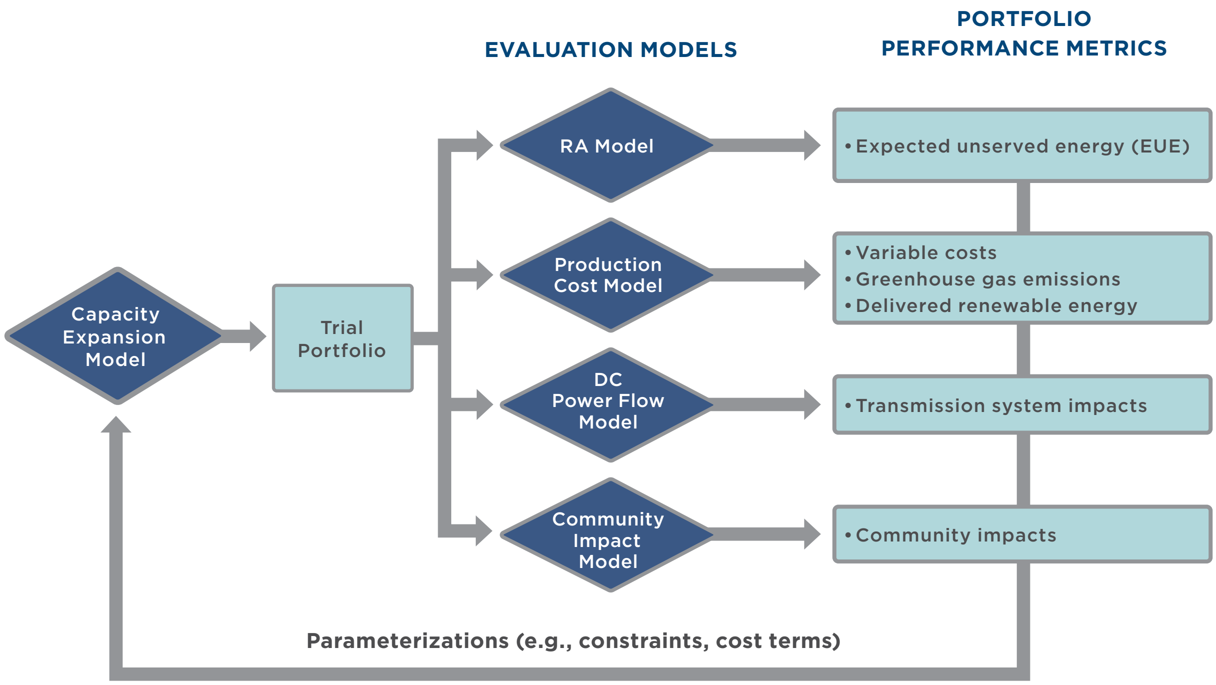

In addition to expanding the scope of questions explored through iterative portfolio optimization, future work could focus on improving convergence behavior for the generalized approach presented in this paper. More rapid convergence, through either improved initialization or changes to the algorithm itself, could promote the use of iterative portfolio optimization for larger systems with more decision variables, planning exercises that involve investment decisions in multiple years, and planning exercises that consider future uncertainties.

Sylvan Energy Analytics provides consulting services and software solutions for tackling today's most pressing and complex clean energy planning challenges. We specialize in integrated resource planning, capacity expansion and production cost modeling, resource adequacy, clean energy policy, and utility regulation.

sylvan.energy

GridLab is an innovative non-profit that provides technical grid expertise to enhance policy decision-making and to ensure a rapid transition to a reliable, cost effective, and low carbon future.

gridlab.org

Find all articles in this series at **gridlab.org/Innovate-modeling-series**