Pacific Northwest
NATIONAL LABORATORY

PNNL-40040

Integrated HPC-Enabled Co-Design and Decision Making Framework

Jennifer N. Pham
Sam Donald
Tonya Martin
John Wenskovitch
Milan Jain
Wei Wang
Palak Mattoo
Cliff Gyamfi
Himanshu Sharma
Rounak Meyur
Dustin L. Arendt
Ji Young Yun
Sumit Purohit

U.S. DEPARTMENT OF ENERGY

Prepared for the U.S. Department of Energy
Under contract DE-AC05-76RL01830

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor Battelle Memorial Institute, nor any of their employees, makes **any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights**. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from
the Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062
www.osti.gov
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
or (703) 605-6000
email: info@ntis.gov
Online ordering: http://www.ntis.gov

PNNL-40040

Integrated HPC-Enabled Co-Design and Decision Making Framework

Jennifer N. Pham
Sam Donald
Tonya Martin
John Wenskovitch
Milan Jain
Wei Wang
Palak Mattoo
Cliff Gyamfi
Himanshu Sharma
Rounak Meyur
Dustin L. Arendt
Ji Young Yun
Sumit Purohit

Prepared for
the U.S. Department of Energy
Under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

Acknowledgments

This research was supported by the E-COMP initiative at the Pacific Northwest National Laboratory (PNNL). The computational work for this research was performed using Research Computing at PNNL. PNNL is a multi-program national laboratory operated for the U.S. Department of Energy (DOE) by Battelle Memorial Institute under Contract No. DE AC05 76RL01830. The work used assistance of AI tools: Claude Code (Opus 4.6 from Anthropic; https://claude.ai) for coding assistance, and Gemini (Gemini 3.1 from Google; https://gemini.google.com) and Microsoft Copilot (GPT-5) for editing the manuscript. All generated code and final text were reviewed, edited, and approved by the author to ensure accuracy.

Contents

Acknowledgments iv

1.0 Introduction 1

2.0 pyMOODS Decision Analytics Framework 3
2.1 Background 3
2.2 Multi-Criteria Decision Making 4
2.2.1 Generalizers and Specializers 4
2.2.2 Visualization 7
2.3 Frontend Development 7
2.3.1 Software architecture 7
2.3.2 Software functionalities 9
2.3.3 Sample code snippet analysis 10
2.4 Case Studies 13
2.4.1 MTDC-ESS Co-Design 13
2.4.2 Large Electric Load Integration 14
2.4.3 Protection-aware Voltage Ride-Through Study 15
2.5 Evaluation 18
2.5.1 MTDC-ESS Co-Design 18
2.5.2 Large Electric Load Case Study 21
2.5.3 Protection-aware Voltage Ride-Through Study 23
2.5.4 Frontend evaluation 26

3.0 CAMEO: A Co-design Architecture for Multi-objective Energy System Optimization . . 28
3.1 CAMEO: Research Capabilities 28
3.2 CAMEO workflow for Large Electric Load (LEL) use-case 28
3.3 Generative AI-based assistant to automate workflow generation 29
3.4 Leveraging Large Language Models for optimization 29
3.5 Decision Support System Integration 31

4.0 Multi-Objective Co-Design for Data Center Energy Storage Planning for Providing Grid Flexibility 32
4.1 Problem Description 33
4.2 Model Formulation 34
4.2.1 BESS size limits 34
4.2.2 DC load operation constraints 35
4.2.3 BESS operation for demand response 36
4.2.4 BESS for fast load fluctuation smoothing 36

4.2.5 BESS state dynamic constraints . . . 36
4.2.6 DC PCC voltage equations . . . 37
4.2.7 DC cooling power equation . . . 37
4.2.8 Objective function . . . 37
4.3 Data Center Case Study . . . 38
4.3.1 Use-case setting . . . 38
4.3.2 Demand response scenario only . . . 40
4.3.3 Voltage sag scenario . . . 41
4.3.4 Integrating demand response and voltage sag scenarios . . . 42
4.3.5 Multi-objective optimization with both demand response and voltage sag scenarios . . . 46

5.0 Conclusion . . . 47

6.0 References . . . 48

Figures

1 Parallel coordinates visualization of the rank matrix with visual encodings for specializations (circles) and tradeoffs (squares). Axes represent simultaneous objectives, not sequential stages. 5

2 Visualization of rank matrix specializations and tradeoffs using (a) cluster heatmap of rank matrix color encoding specialization (blue square) and tradeoffs (red squares). Solutions in the upper half tend to specialize in r_3, r_5, and r_1 and tradeoff the other objectives–the pattern reverses in the lower half, and (b) an Euler diagram of blue-filled specialization and red-outlined tradeoff sets. Solutions $i = \{3, 8, 9\}$ belong to the r_2 specialization and r_4, r_5 tradeoff sets. 6

3 pyMOODS system architecture. A React frontend hosting the visual decision-support layer communicates with a Python backend through a Flask REST API; the data access module loads, caches, and validates CSV and JSON inputs, and the backend performs data processing together with the machine learning and rank-analysis routines. 8

4 Evolution of the pyMOODS dashboard from early Dash-based prototypes (a–c) to the current React-based implementation (d). 11

5 JSON schema of input data. In this example, the categories “input parameters” and “control inputs” are empty. However, these categories do appear in other formulations. 12

6 Parallel coordinates plot showing the relative performance ranks of all design configurations across physical and economic objectives. 19

7 Cluster heatmap grouping similar operational strategies at the Tesla node. 19

8 Euler diagram mapping operational strategies at the Tesla node. 20

9 Parallel coordinates visualization of the rank matrix for the Large Electric Load (LEL) case study across three objective weighting scenarios: reliability-focused, balanced, and cost-focused. 22

10 Worst case scenario demonstrating slaking for static cap and time-graded ceiling . . . 23

11 Static cap protection trace feasibility boundary 24

12 Time-graded protection trace feasibility boundary 25

13 CAMEO workflow for Large Electric Load use-case 29

14 Multi-agent framework to automate workflow generation 30

15 Generative AI-based assistant to automate workflow generation 30

16 CAMEO generated summary for Decision Support Systems 31

17 Data center grid connection and power supply architecture. The root node is connected to the main grid through the PCC. The total DC load and the aggregate charging and discharging power of all BESS units are represented at the root node. 39

18 Data center IT load profile over a 48-hour period at 1-minute resolution. 40

19 Data center net load in the demand-response-only case. The periods during which demand response events are requested are marked in pink. 41

20 BESS energy-capacity allocation and SoC trajectory in the demand-response-only case. The energy capacity allocated for demand response is marked in blue. The yellow regions represent the upper and lower 10% of the total BESS energy capacity reserved to accommodate energy exchange for fast fluctuation smoothing. . 41

21 PCC and data center root-node voltages in the voltage-sag-only case. The interval from Sep. 1 at 3:10 am to 3:50 am is enlarged. 42

22 BESS SoC trajectory in the voltage-sag-only case. The yellow regions represent the upper and lower 10% of the total BESS energy capacity reserved to accommodate energy exchange for fast fluctuation smoothing. 42

23 PCC and data center root-node voltages in the voltage-sag-only case with voltage sag last for one time step. The interval from Sep. 1 at 3:10 am to 3:50 am is enlarged. 43

24 Data center net-load and BESS SoC trajectories under extreme demand response scenarios. During the second DR event, the peak-load reduction is less than 20% because the BESS has insufficient time to fully recharge. 44

25 Randomly generated scenario configurations and corresponding BESS sizes from the design results. The upper subplot illustrates the occurrence times and types of the two events in each scenario for all configurations. The lower subplot shows the BESS rated-power and energy-capacity results for each configuration. 45

26 Pareto front for the multiple objectives. The x-axis represents demand response capacity, the y-axis represents voltage support capacity, and the color of the dots indicates the total BESS cost. 46

Tables

1 Non-parametric characterization of generalizability, specialization, and tradeoffs. Black ▲: specialization; red ▼: tradeoff; gray: neither. 4

2 Feature status in pyMOODS v0.0.3. Stages (a)–(d) refer to the development iterations shown in Figure 4. *Implemented* denotes a capability present in the archived v0.0.3 release; *Prototype only* a capability that existed in an earlier iteration but is not part of the current release; *Experimental* code present in the repository but not exposed through the interface; and *Planned* a capability under development. 9

3 Characteristics of operation modes 14

4 Characteristics of location scenarios 15

5 MTDC–ESS co-design results. Solutions are ordered by decreasing net profit. 18

6 Reproducible performance evaluation across formulations and tasks. Each task was repeated three times under identical conditions; runtime is wall-clock latency per task. The exact API call sequence for each task, and the dimensionality-reduction backend, are given in the accompanying text. 27

7 BESS Size Solutions and Cost for Single-Objective Cases 40

1.0 Introduction

Large-scale infrastructure planning requires an analysis of assets under competing economic, operational, and reliability objectives. These decisions must also remain useful across uncertain future conditions. Co-design models can jointly optimize infrastructure investments and operational controls, but their outputs are often sets of non-dominated alternatives rather than a single design that is best on every criterion (Krishnan et al. 2016; Dvorkin et al. 2018; Cho, Li, and Grossmann 2022; Miettinen 1999). As the number of objectives and scenarios grows, the decision problem changes from identifying an optimum to understanding which compromises are acceptable. A design that performs well in one scenario may perform poorly in another. For example, joint planning of high-voltage direct-current cable capacity and energy storage can be evaluated through aggregate net profit. Although useful for comparing overall economic value, that aggregate can obscure the individual contributions of day-ahead, real-time, and reserve markets, as well as the cost of the underlying design choices. A decision-maker therefore needs to inspect both the summary metric and the objective- and scenario-level evidence from which it is formed.

To address the computational and analytical demands of complex energy system co-design, this report presents an integrated suite of frameworks and optimization models designed to streamline workflow execution, facilitate scenario management, and enhance multi-criteria decision-making. Specifically, this work details three interconnected projects that collectively advance the life-cycle of energy infrastructure planning:

pyMOODS Decision Analytics Framework: Multi-objective optimization produces candidate designs, but stakeholders need robust methods to interpret these candidates post-optimization. pyMOODS supports this post-optimization stage (Pham, Jain, and Mattoo 2026). It loads candidate solutions and case-study metadata, computes complementary multi-criteria decision-making (MCDM) characterizations, and presents coordinated views of the solution and decision spaces. Rather than treating a numerical ranking as a final decision, the framework makes it possible to ask whether a highly ranked design is broadly balanced, unusually strong on a particular objective, or sensitive to a change in preferences or scenario. The goal is to help domain experts and stakeholders interpret the consequences of a choice while retaining responsibility for the choice itself (Osika et al. 2023).

CAMEO (Co-design Architecture for Multi-objective Energy System Optimization): Complex energy system co-design requires computational workflows that can manage numerous interdependent components, support a wide range of optimization methods, and scale effectively across high-performance computing environments. Conventional solutions often rely on specialized, monolithic tools that are difficult to reuse or extend, making them poorly suited for heterogeneous co-design tasks involving multiple scenarios and optimization formulations. We developed CAMEO (Co-design Architecture for Multi-objective Energy System Optimization), designed to overcome these limitations. It provides a modular workflow framework that structures co-design problems as Directed Acyclic Graphs (DAGs), where each computational stage communicates through standardized input-output interfaces. With JSON-based workflow specifications, CAMEO decomposes complex studies into reusable components such as data loaders, scenario generators, optimization solvers, and result summarizers. Its architecture supports multiple optimization paradigms through containerized execution environments and integrates seamlessly with HPC systems via Nextflow. This modular, plug-and-play design allows users to modify optimization strategies without altering the underlying workflow logic, while standardized interfaces ensure compatibility across diverse tools and execution platforms.

Multi-Objective Data Center Energy Storage Formulation: The rapid expansion of artificial intelligence and large-scale computational workloads has introduced significant new electricity demands and grid integration challenges. To demonstrate the application of advanced

co-design principles, this report introduces an optimization-based framework for sizing battery energy storage systems (BESSs) in grid-interactive data centers. This formulation explicitly balances internal data center constraints with external grid-support requirements, simultaneously optimizing for demand response capabilities, fast load fluctuation smoothing, and point-of-common-coupling (PCC) voltage support during main-grid faults. By integrating these factors into a unified planning model, decision-makers can evaluate the multidimensional trade-offs between BESS investment costs and operational grid flexibility.

2.0 pyMOODS Decision Analytics Framework

2.1 Background

MCDM provides several families of methods for evaluating conflicting objectives. Weighted value or utility models aggregate normalized criteria; the Analytic Hierarchy Process (AHP) elicits pairwise judgments; and TOPSIS ranks alternatives by distance from ideal and anti-ideal points. Outranking methods such as ELECTRE and PROMETHEE use pairwise relations and may represent indifference, preference, or veto thresholds, while Stochastic Multicriteria Acceptability Analysis (SMAA) evaluates acceptability across uncertain preference information (Cinelli et al. 2020). These methods offer preference-sensitive rankings, partial orders, or robustness assessments, but require choices about weights, thresholds, normalization, compensation, or uncertainty models. Aggregation choices can influence rankings (Triantaphyllou and Baig 2005), and some methods also depend on the candidate set (Cinelli et al. 2020). No method is therefore a universal baseline; the appropriate choice depends on the available preference information and desired recommendation.

MCDM is well established in energy planning, where technical, economic, environmental, and social criteria must be considered together (Kumar et al. 2017). Multi-criteria and stochastic frameworks also support cooperative planning (Tian, Wang, and Wang 2023) and ranking loads (Jain et al. 2021). These applications motivate comparison with preference-based MCDM, but the present work addresses a different stage: parameter-free, post-hoc examination of solution sets before a decision-maker supplies a preference model.

Parallel to theory, visualization extensively supports multi-criteria decision-making (MOO/MCDM), including foundational scatter plot matrices (Chambers 2018), parallel coordinates (Inselberg 2008), heatmaps, and lower-dimensional latent space projectionssuch as self-organizing maps (Kohonen 2012) and RadVis (Hoffman, Grinstein, and Pinkney 1999). Because taxonomies of desirable frontier properties (Tušar and Filipič 2015) and data-driven use (Filipič and Tušar 2018) reveal that no single technique is universally optimal, modern approaches leverage multiple coordinated views (Munzner 2025). Tools like LineUp (Gratzl et al. 2013), SRVis (Weng et al. 2018), Podium (Wall et al. 2017), and frameworks by Smedberg and Bandaru (2023) synthesize parallel coordinates, spatial projections, and network diagrams to help decision-makers explore objective weightings and learned decision rules. Pareto or nondominated-set analysis retains alternatives for which no other alternative is at least as good on every objective and strictly better on at least one objective (Tušar and Filipič 2015). It requires no weights and preserves objective magnitudes, but it may leave a large set without an internal ordering. Weighted aggregation, AHP, and TOPSIS can instead recommend a preference-sensitive ordering; outranking methods can express thresholds, vetoes, or incomparability; and SMAA can characterize robustness to uncertain preferences. Our framework does none of those tasks. It converts each objective to ranks and derives relative generalizers, specializers, and trade-offs without eliciting preferences or selecting a preferred alternative. Its contribution is this explicit structural characterization, not exclusive access to insights that tabular inspection of the objective-performance matrix, Pareto views, or other MCDM methods could also reveal. These methods are therefore complementary rather than task-equivalent performance baselines. The extent to which the proposed filtering reduces decision-maker consideration remains to be quantified.

2.2 Multi-Criteria Decision Making

We aim to support decision-makers' post-hoc evaluation of multi-criteria, multi-scenario solution sets through three goals, published in Arendt et al. (2027):

Goal 1: Reduce the number of alternatives requiring detailed consideration.

Goal 2: Describe the useful and distinctive properties of retained alternatives.

Goal 3: Map structural similarities and trade-offs across criteria and scenario realizations.

We pursue these goals using a non-parametric method that does not require decision-makers to specify objective weights. The remainder of this section defines MCDM-TAGS in detail.

2.2.1 Generalizers and Specializers

We characterize the *generalizability* of a given solution relative to others in the solution set using a technique similar to the game theory principle of minimizing maximum loss or "minimax" from Tadelis (2013). We first compute the rank matrix R_F where r_{ij} is the rank of solution i for objective j. The generalizability of a given solution is then its maximum (worst) rank across all objectives, i.e. $g(i) = \max_j \{r_{ij}\}$. We sort solutions in the rank matrix by this worst rank, breaking ties by considering the next-worst rank sequentially. Table 1 shows an example where solutions are evaluated on r_j objectives and sorted based on their worst rank shown under column g.

Table 1: Non-parametric characterization of generalizability, specialization, and tradeoffs. Black ▲: specialization; red ▼: tradeoff; gray: neither.

	r_1	r_2	r_3	r_4	r_5	g
1 [a]	6	4	2	4	7	7
2	4 ▲	5 ▼	4 ▼	8 ▼	5 ▲	8,5
3	2 ▲	3 ▲	7 ▼	7	8 ▼	8,7
4	1 ▲	8 ▼	8 ▼	5	3 ▲	8,8
5 [b]	5	9 ▼	3	2 ▲	2 ▲	9,5
6	9 ▼	6	1 ▲	3	6	9,6
7 [c]	8	7	9 ▼	6	4	9,8
8	7	2 ▲	5	9 ▼	9 ▼	9,9
9	3	1 ▲	6	10 ▼	10 ▼	10,10,6
10 [d]	10 ▼	10 ▼	10 ▼	1 ▲	1 ▲	10,10,10

[a] Most-generalizable solution.
[b] Example tradeoff: specializes in r_4 and r_5 and trades off in r_2.
[c] No specialization.
[d] Least-generalizable solution.

Starting with this sorted rank matrix, a solution *specializes* in an objective if its rank for that objective is strictly better than the performance of all solutions having better generalizability (i.e., all ranks above it in the matrix). Conversely, a solution exhibits a *tradeoff* for a given objective when its performance is worse than all solutions having better generalizability. To be precise, for row $i > 1$, let the preceding column minimum be $m_{ij} = \min_{k<i} r_{kj}$ and the preceding column

maximum be $M_{ij} = \max_{k<i} r_{kj}$. A rank is marked as a specialization when $r_{ij} < m_{ij}$ and as a tradeoff when $r_{ij} > M_{ij}$. Thus, the marks can be computed in one pass using cumulative column minima and maxima. Identifying the specializer solutions and the specific objectives in which they specialize supports **Goal 2** by describing the useful and distinctive properties of retained alternatives.

Row 1 is the unmarked, most-generalizable solution because it has no preceding solution. In the tie-free example, row 2 compares only with row 1: its $r_1 = 4$ is a specialization because $4 < 6$, whereas its $r_2 = 5$ is a tradeoff because $5 > 4$. Consequently, every row-2 objective in this example receives one of the two marks. Later rows compare with all preceding rows and can contain either mark or neither. For example, row 5 specializes in r_4 and r_5, trades off in r_2, and is neither in r_1 and r_3. Row 7 has no specialization.

Exact ties are uncommon in our case studies, but they can occur when objective values or sorted rank profiles coincide. Exact duplicate objective profiles should be grouped before applying the method, and any remaining ties require a consistent ordering rule.

This parameter-less approach can filter solutions without a specialization, supporting **Goal 1**. The number of solutions to consider can be further reduced by finding the smallest set of solutions such that each objective is specialized in by at least one solution. This is equivalent to computing a minimum set cover on the specialization matrix (Korte and Vygen 2008).

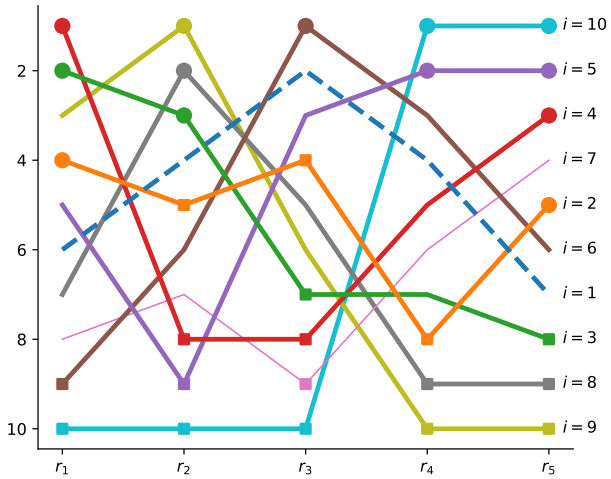
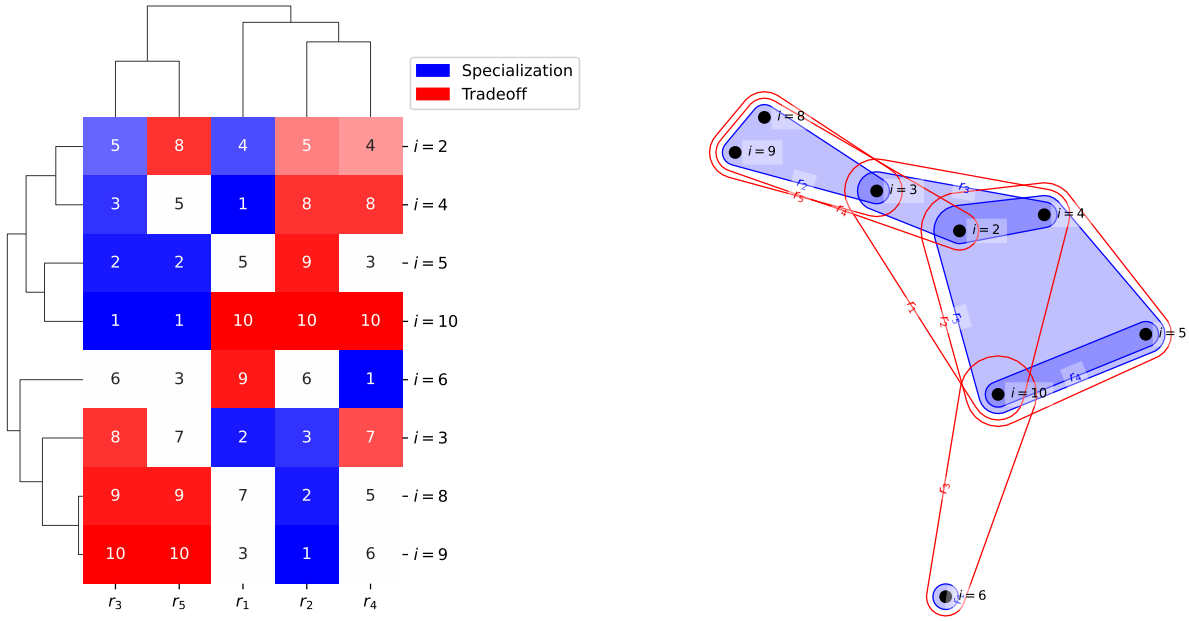


Figure 1: Parallel coordinates visualization of the rank matrix with visual encodings for specializations (circles) and tradeoffs (squares). Axes represent simultaneous objectives, not sequential stages.

The approach is computationally tractable for large datasets. For n solutions and m objectives, assume a dense, complete numeric performance matrix and comparison-based sorting. Rank construction costs $O(mn \log n)$ time and $O(mn)$ stored space. For the stated tie-breaking rule, sorting each solution's rank vector costs $O(nm \log m)$ and lexicographically ordering the vectors costs $O(mn \log n)$ in the worst case. Cumulative minima and maxima then identify all marks in $O(mn)$ time, using $O(m)$ streamed working space or $O(mn)$ when cumulative and indicator matrices are retained. The full retained characterization therefore costs $O(mn(\log n + \log m))$ time and $O(mn)$ space, or $O(n \log n)$ time and $O(n)$ space for fixed m. The optional exact minimum-set-cover reduction is NP-hard; a standard greedy approximation has an $H_m \leq 1 + \ln m$ guarantee and naive $O(nm^2)$ time and $O(nm)$ space (Korte and Vygen 2008).



(a) Cluster heatmap visualization

(b) Euler diagram

Figure 2: Visualization of rank matrix specializations and tradeoffs using (a) cluster heatmap of rank matrix color encoding specialization (blue square) and tradeoffs (red squares). Solutions in the upper half tend to specialize in r_3, r_5, and r_1 and tradeoff the other objectives–the pattern reverses in the lower half, and (b) an Euler diagram of blue-filled specialization and red-outlined tradeoff sets. Solutions $i = \{3, 8, 9\}$ belong to the r_2 specialization and r_4, r_5 tradeoff sets.

2.2.2 Visualization

To communicate the identified characteristics and structural relationships, we map the sorted rank matrix R_F, along with its derived specializations and tradeoffs, into three distinct visual encodings that support the stated analytical goals.

Parallel coordinates. This view traces solution ranks across simultaneous, unordered objectives (Figure 1); the inverted y-axis encodes rank, circles specializations, and squares tradeoffs. The most-generalizable solution ($i = 1$) is dashed and solutions without specializations ($i = 7$) are thinner. For static rendering, the implementation recommends an order from Pearson correlations of objective-rank columns among specializers: it sets negative correlations and the diagonal to zero, then applies graph-Laplacian spectral seriation to the resulting nonnegative adjacency matrix (Atkins, Boman, and Hendrickson 1998). This is a heuristic, not a unique or optimal order. The implementation also accepts the explicit domain orders used in the case-study figures. Reordering changes adjacency and salience, not ranks, generalizability, or specialization and tradeoff classifications.

Cluster heatmap. Hierarchical clustering to reorder the rank matrix, revealing groups with similar characteristics (Figure 2a) and supporting **Goals 2 and 3**. Its colormap encodes specialization as blue and tradeoffs as red; it omits the most-generalizable solution and solutions without specializations.

Euler diagram. This visualization is designed to arrange solutions and objectives in a way that reveals structural relationships among groups sharing the same characteristics, thus supporting **Goals 2 and 3**. Euler diagrams (Alsallakh et al. 2014) are suitable for visualizing sets. The diagram represents solutions as points and sets as hulls containing member solutions. We transform the annotated rank matrix in Table 1 into two binary indicator matrices. The specialization matrix marks cells identified by black up-carets, while the tradeoff matrix marks cells identified by red down-carets. For each objective, the marked solutions in its column form that objective's specialization or tradeoff set, respectively. Multiple solutions can therefore belong to the same set, and one solution can belong to sets for several objectives. The Euler diagram visualizes these sets and their intersections.

Figure 2b fills specialization-set hulls blue and outlines tradeoff-set hulls red. An objective will usually appear as two different hulls, one for the solutions that specialize in it and one for the solutions that trade off performance in it. Hull containment and intersection encode binary membership only; they do not encode the magnitude of a rank improvement or loss. Magnitudes and complete rank profiles remain available in the annotated rank matrix and parallel coordinates plot.

2.3 Frontend Development

2.3.1 Software architecture

pyMOODS (Mattoo et al. 2026) couples a computational backend with an interactive frontend so that decision-makers can explore multi-objective optimization results without re-running the underlying optimization. The frontend is a React and MUI application MUI 2014 rendering linked interactive plots built with React Plotly. User actions issue asynchronous requests to a Flask REST API, which mediates all frontend–backend communication. Behind that API, a data access module loads, validates, and caches datasets held as CSV and JSON files, and a data processing module performs transformation, filtering, and post-processing before invoking the analytical routines described below. Results return through the API to the client, which updates the visualizations. Because interaction state is held in the frontend while analysis runs in the

backend, computation does not block the interface. Figure 3 shows the resulting component structure.

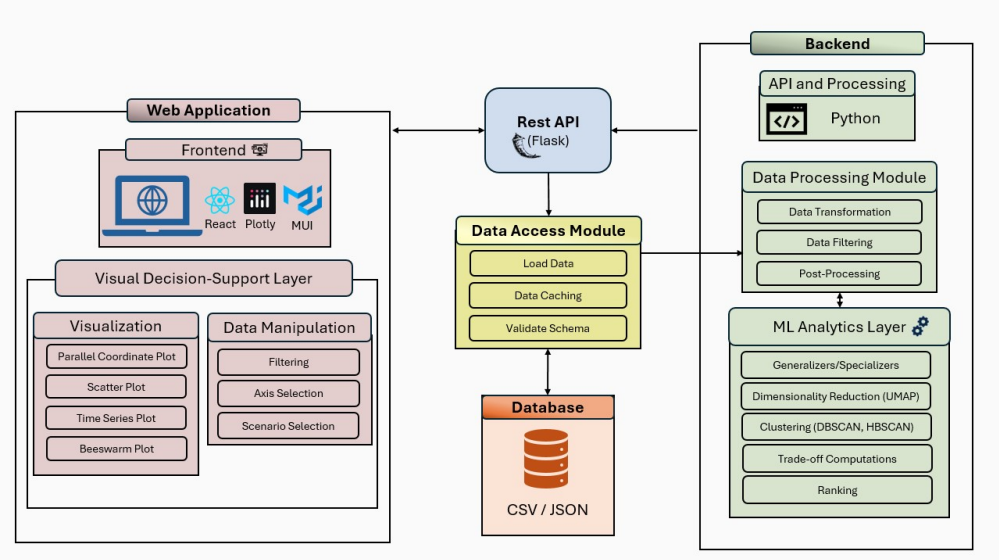


Figure 3: pyMOODS system architecture. A React frontend hosting the visual decision-support layer communicates with a Python backend through a Flask REST API; the data access module loads, caches, and validates CSV and JSON inputs, and the backend performs data processing together with the machine learning and rank-analysis routines.

Input schema: A problem formulation is supplied as a single JSON configuration file (Figure 5) that links to the data files and assigns each column to one of five categories:

- *Input parameters:* Fixed attributes (e.g., location or operating conditions);
- *Hyperparameters:* Configurable, non-optimized variables (e.g., technology choice);
- *Decision variables:* Variables subject to optimization;
- *Objective functions:* Quantities to be minimized or maximized; and
- *Control inputs:* Intermediate timeseries monitoring system states.

This schema lets one application serve formulations whose objectives, decision variables, and hyperparameters differ in number and meaning, rather than assuming a fixed problem structure.

Backend analytics. On loading a formulation, the backend computes the quantities the interface depends on, using two classes of method that differ in kind. Structuring the solution space uses unsupervised machine learning: a Uniform Manifold Approximation and Projection (UMAP) McInnes, Healy, and Melville 2020 embedding projects the high-dimensional objective and decision space into two dimensions, and density-based HDBSCAN clustering groups solutions within it. These are the framework's machine learning components. Trade-off interpretation instead uses a deterministic, parameter-free rank analysis: solutions are ranked on each objective, a solution's *generalizability* is its worst rank across all objectives, and a solution *specializes* in an objective when its rank there is strictly better than that of every more generalizable solution. Because this characterization derives entirely from the rank matrix, it introduces no trained model, tuned threshold, or random seed, and identical inputs always

produce identical generalizer and specializer assignments. Solutions can additionally be ordered by an aggregate score reflecting user-specified objective weights, providing a familiar mechanism for expressing stakeholder priorities.

The backend developed from a single-process Plotly Dash Inc. 2015 prototype into this decoupled service, adding the JSON schema and the asynchronous API before the current release.

Table 2: Feature status in pyMOODS v0.0.3. Stages (a)–(d) refer to the development iterations shown in Figure 4. *Implemented* denotes a capability present in the archived v0.0.3 release; *Prototype only* a capability that existed in an earlier iteration but is not part of the current release; *Experimental* code present in the repository but not exposed through the interface; and *Planned* a capability under development.

Component	Capability	Status (v0.0.3)	Stage
Backend	JSON configuration schema (five variable categories)	Implemented	(c)
	Asynchronous Flask REST API	Implemented	(c)
	UMAP latent-space embedding	Implemented	(d)
	Unsupervised solution clustering (HDBSCAN)	Implemented	(a)
	User-weighted sum ranking	Implemented	(b)
	Generalizer/specializer rank analysis	Implemented	(d)
	Trade-off lattice computation	Experimental	—
	Additional MCDM methods for comparison	Planned	—
	AI chatbot assistant	Planned	—
Frontend	Parallel coordinates plot	Implemented	(a)
	Scatter plot (objective pairs and UMAP view)	Implemented	(a)
	Beeswarm plot	Implemented	(d)
	User-weighted ranking table	Implemented	(d)
	Time-series panel	Implemented	(d)
	React and MUI application	Implemented	(d)
	Interactive decision-space sliders	Prototype only	(a)
	Radar chart	Prototype only	(b)
	Right-side filter and plot-control panel	Prototype only	(c)

2.3.2 Software functionalities

Problem configuration. A user selects a formulation and loads it from its JSON specification without modifying code, and a summary then reports the dataset's composition, including the

counts of decision variables, objective functions, and solutions. Because the schema is declarative, one deployment serves formulations that differ in the number and meaning of their variables.

Solution-space exploration. The framework reduces a high-dimensional front to a navigable two-dimensional layout in which groups of similar solutions are visible at a glance and individual candidates can be located. A user may instead plot any pair of objectives or decision variables when a specific relationship is of interest, and may restrict the working set to solutions sharing a chosen hyperparameter configuration, such as a technology option or case study.

Preference specification and ranking. A user assigns weights to the objectives and obtains an immediate re-ranking of the whole solution set. Because ranking is computed from the stored front rather than by re-solving, stakeholders with differing priorities can interrogate the same results interactively, and the consequence of a change in priorities is visible without repeating the optimization.

Trade-off interpretation. Beyond ordering solutions, pyMOODS reports what each one concedes. The generalizer and specializer characterization identifies, for every candidate, the objectives on which it outperforms all more broadly balanced alternatives and those on which it gives ground. This answers a question that ranking alone cannot: not merely which solutions score well under a given weighting, but which compromise each one represents. A selected candidate can then be placed within the distribution of all solutions for each objective and decision variable, showing whether its performance is typical or extreme.

Scenario and temporal analysis. For formulations that carry control inputs, the behaviour of a selected solution over time can be examined, so that dynamic performance can be compared across scenarios rather than only through aggregate objective values.

Coordinated interaction. Selection propagates across views: choosing a solution in one panel updates the others, so a candidate can be examined from several perspectives without losing context or cross-referencing identifiers by hand.

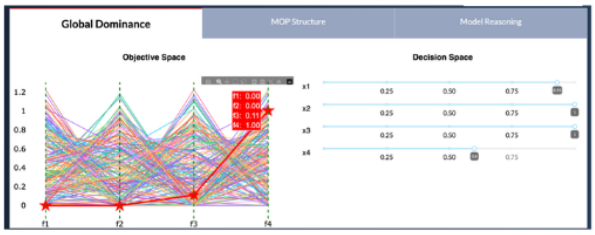
Current implementation. Figure 4 traces how these functionalities were realized, from the initial Dash prototype (a), through iterations that added high-dimensional exploration (b) and structured input with asynchronous communication (c), to the current React application (d). In the present release the left sidebar carries problem configuration, objective weighting, and hyperparameter filtering; the central workspace carries the UMAP-based projection and the parallel coordinates view that encodes generalizers and specializers; and the lower panels carry the weighted ranking table, the beeswarm distributions, and the time-series view. Figure 4d shows that interface, and Table 2 records which capabilities are implemented in v0.0.3, which existed only in earlier prototypes, and which remain experimental or planned.

2.3.3 Sample code snippet analysis

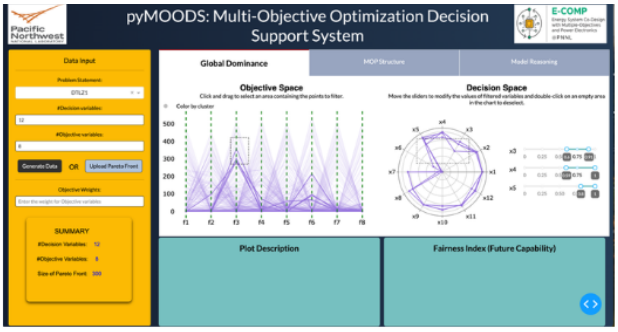
A problem formulation is supplied to pyMOODS as a single JSON descriptor that names the data file and declares the hyperparameters, decision variables, and objective functions. Figure 5 shows this structure for the MoCoDo formulation used in Section 3. The frontend then retrieves the parsed formulation and its solution set through a single call, `GET /api/solutions?use_case=MoCoDo_v3`, and no further configuration is required.

The pyMOODS dashboard includes automated frontend tests using Playwright. These tests verify that the main dashboard components load correctly, the solution table can be sorted, and the generalizer–specializer controls update the displayed results.

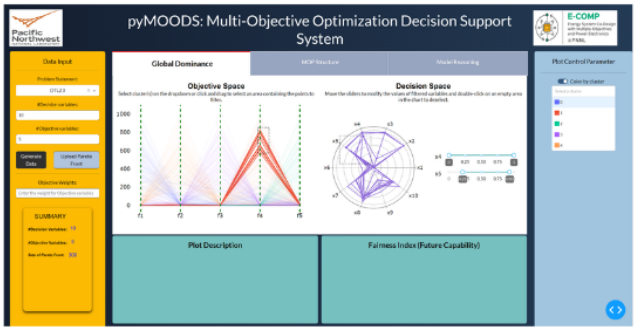
The tests use simulated backend responses to evaluate communication between the user interface and API without requiring the full backend to run. This provides repeatable validation of the main user workflows. Dedicated automated tests for the backend calculations are not



(a) Dash-based evolution from low-D Pareto views to parallel coordinates



(b) Parallel coordinates and radar charts for high-D filtering



(c) Adding plotting options



(d) React redesign with backend-generated

Figure 4: Evolution of the pyMOODS dashboard from early Dash-based prototypes (a–c) to the current React-based implementation (d).

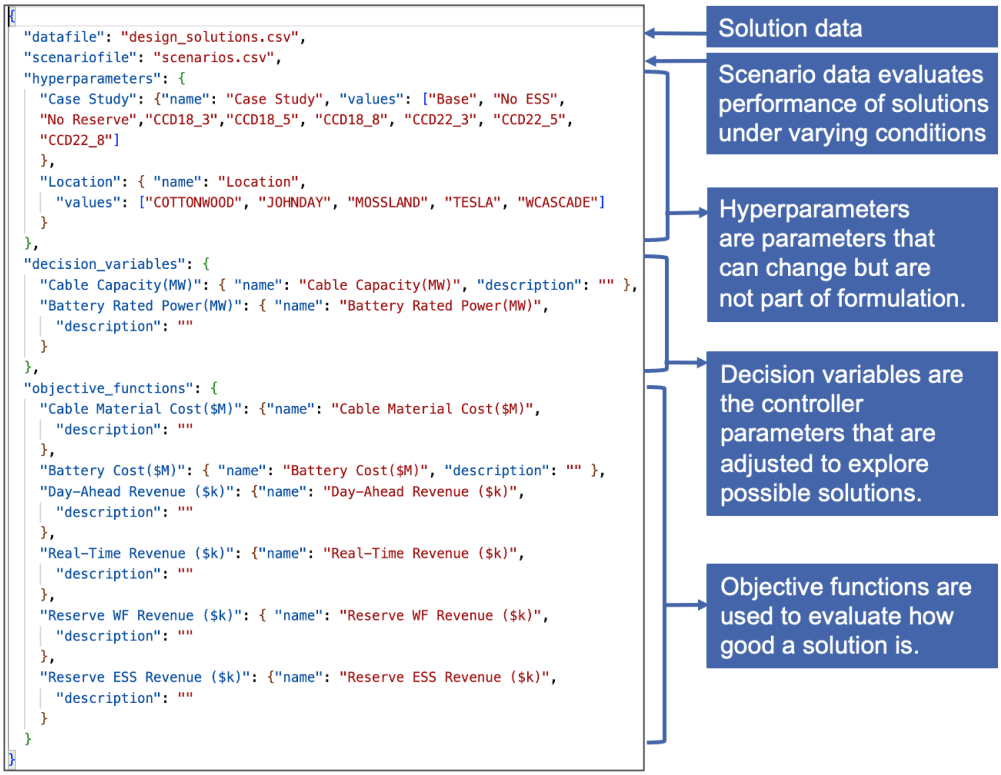


Figure 5: JSON schema of input data. In this example, the categories “input parameters” and “control inputs” are empty. However, these categories do appear in other formulations.

currently included and represent an area for future development. On loading a formulation, the backend checks that the JSON descriptor parses, that it declares a data file, and that the referenced CSV and any scenario file exist, returning a diagnostic error response rather than failing silently when a check does not pass. Datasets are read from the local filesystem of the machine running the application, and the backend exposes no upload endpoint, so the deployment introduces no server-side upload surface.

2.4 Case Studies

2.4.1 MTDC-ESS Co-Design

Sharma et al. (2025) proposed MTDC-ESS co-design and analyzed the joint sizing of HVDC export cables and onshore battery energy storage systems (ESS), together with the tuning of droop control parameters for connecting large non-conventional generators to the grid. The objective is to identify the optimal combination of HVDC cable capacity, ESS power rating, and operational control parameters for frequency response, to maximize the expected revenues from the market minus capital costs, while satisfying power delivery, frequency support, and operational constraints under uncertainty in resource output and electricity prices.

2.4.1.1 System design and formulation

The system consists of five potential generation sites radially connected to designated onshore points of interconnection (POIs) via HVDC export cables. The upper limit for the lithium-ion energy storage is set to 5% of the generation rated capacity to provide both short-term balancing through energy shifting and primary frequency reserve capability. The system participates in day-ahead (DA) and real-time (RT) energy markets, as well as ancillary service markets for upward and downward reserves, using historical data from the California ISO.

The study formulates the problem as a multi-timescale, multi-stage stochastic optimization, with design decisions made prior to market realization, DA market allocations determined once day-ahead prices are known, and RT adjustments based on actual resource output and real-time market prices. Design variables include HVDC cable capacity and ESS size, while control variables include droop gain parameters for both the generation units and ESS, ESS charge/discharge operation and state of charge, and HVDC power flow from offshore sites to onshore POIs. The objective function accounts for capital costs, operational revenues, and reserve participation payments. For additional details about the formulation, please refer to the paper from Sharma et al. (2025).

2.4.1.2 Case variants

The study considers different inflation rates, different historical market years to represent a range of price environments, and variations in operational configuration, such as with or without ESS deployment and with or without reserve market participation.

For each case, the optimizer determines the HVDC cable capacity, ESS size, and control parameters best suited to maximize net profit while satisfying operational and reliability constraints.

Table 3: Characteristics of operation modes

Attribute	Standard	High redundancy	AI/ML
Primary use	General computing; web services	Critical services; financial services	AI training; HPC workloads
Base load	Moderate	High	Variable
Peak load	Moderate	High	Very high
Variability	Low	Very low	Very high
Load pattern	Steady; slight daytime rise	Nearly constant	Spiky compute bursts
Redundancy	Standard	Extensive	Moderate
Grid challenge	Routine capacity planning	Extreme reliability	Rapid-load management

2.4.2 Large Electric Load Integration

The LEL use case determines optimal transmission grid expansion strategies to reliably support large-scale data center installations with power demands ranging between 100 and 200 MW. The objective is to determine the optimal set of transmission lines to expand or construct to ensure N-1 reliability standards for the data center while minimizing total expansion costs.

2.4.2.1 System design and formulation

Expanding the grid to add a new hyperscale data center is formulated as a mixed-integer linear programming problem. The power grid is a network graph of interconnected buses and transmission lines, encompassing both existing infrastructure and potential candidate routes. The objective function minimizes the weighted sum of two competing terms: (1) the total capital investment for grid expansion, which includes both a fixed cost for initiating a line upgrade and a variable cost proportional to the added capacity, and (2) the expected financial penalty of system-wide load shedding. The primary design decisions involve binary variables dictating whether to construct or reinforce specific lines, paired with continuous variables that determine the exact magnitude of apparent power capacity (MVA) installed. To ensure physical feasibility, the optimization is bound by a DC power flow approximation that links bus voltage angles to active power flows via transmission line susceptance.

2.4.2.2 Case variants

Data centers exhibit distinct operational characteristics (Table 3) that significantly impact power grid planning requirements and infrastructure design. *Standard operation* data centers, serving general computing, maintain moderate loads with low variability and steady consumption

Table 4: Characteristics of location scenarios

Attribute	Urban edge	Industrial park	Rural tech hub
Context	Metropolitan periphery	Mixed industrial zone	Rural tech area
Grid strength	Very strong	Moderate	Weak
Corridors	Limited	Available	Abundant
Congestion	High	Moderate	Low
Load distance	Close	Moderate	Far
Investment	Reinforce existing lines	Selective expansion and upgrades	New transmission lines

patterns enabling consistent capacity planning. *High redundancy* data centers for critical services operate with high base loads, low variability, and nearly flat consumption, presenting extreme reliability challenges. *AI/ML-focused* data centers demonstrate variable base loads with very high peak demands and consumption spikes during computational bursts, creating a challenging scenario for grid operators managing rapid load changes.

The data center location significantly influences grid expansion strategies and infrastructure investment requirements due to varying regional characteristics and existing grid capabilities (see Table 4). *Urban Edge* locations at metropolitan peripheries have transmission infrastructure and close proximity to load centers but face high congestion and limited transmission corridors, necessitating reinforcement of existing lines. *Industrial Park* locations offer moderate grid strength with available corridors and moderate congestion, enabling selective expansion and upgrade approaches. *Rural Tech Hub* locations present weak existing infrastructure but abundant corridors and low congestion, requiring new transmission construction due to distance from load centers and limited existing capacity.

2.4.3 Protection-aware Voltage Ride-Through Study

To support the provision of ancillary services in data center operation, there have been efforts to ensure standardized grid integration and smooth operations during fault events. Voltage ride-through (VRT) studies for large loads such as data centers focus on internal voltage control and UPS coordination but do not model the relay response to the resulting fault signatures as inferred by Xie, Cui, and Wierman (2025). VRT requirements and adaptive relay protection have evolved as largely separate research streams. However, there is a critical gap at their intersection for large electric loads. Existing work reveals that VRT control strategies for inverter-based resources (IBRs) are developed without explicit coordination with protective relay algorithms, while adaptive relay schemes treat ride-through behavior as a disturbance rather than a co-design function.

Uninterrupted Power Supply (UPS) systems, Battery Energy Storage Systems (BESS) and other power electronic interfaces in the data center can help keep them connected to the power grid during voltage disturbances. However, some converter actions that support VRT also change the currents observed by protective relays and can cause nuisance tripping. Protection and VRT can impose competing requirements. Greater UPS/BESS support can improve VRT performance while pushing relay current toward pickup. With increasing data center grid integration, this problem can no longer be treated only as a load management problem. This study proposes a co-design synthesis, clearly demarcating feasibility and non-feasibility regions. Wayne et al. (2026) formulate VRT as a signal temporal logic (STL) controller synthesis problem where temporary requirements such as remaining connected, maintaining voltage and active power floors during disturbance and restoring demand during recovery are to be satisfied. This study builds on that framework by adding a protection clause to the dispatch per-timestep while evaluating relay placement, time dial settings, and coordination margins.

2.4.3.1 Methodology

The ERCOT NOGRR282 worst-case staircase trace, which is, a fixed code-mandated envelope with bands (0.80–0.90 pu, 2.00 s), (0.50–0.80 pu, 0.50 s), (0.35–0.50 pu, 0.25 s), (0.00–0.35 pu, 0.15 s) is the standard used in this study. A protection clause is added as a co-optimization feature. Protection is separated into three objectives:

$$\varphi_{\text{prot},r} = \varphi_{\text{security},r} \wedge \varphi_{\text{dependability},r} \wedge \varphi_{\text{coordination}}$$

The mature protection objective requires a genuine fault to be cleared and the correct primary relay to operate before its backup. In the current implementation, these are design-time targets rather than runtime constraints. The dependability and coordination objectives are defined as:

$$\varphi_{\text{dependability},r} = \square\big[F_r \rightarrow \Diamond_{[0,T_{\text{max}}]}\, TRIP_r\big]$$

$$\varphi_{\text{coordination}} = \square\big[F_z \rightarrow (t_b - t_p \geq CTI)\big]$$

In the current study, security is enforced at VRT runtime. The dependability and coordination properties will be verified and checked during design implementation. Generally, in a dimensionless form, relay security can be expressed as:

$$\rho_r(t) = \frac{I_r(t)}{I_{\text{pickup},r}}, \qquad \text{requirement: } \rho_r(t) \leq \rho_r^{\text{max}} < 1$$

An alternate protection-security model is the accumulated inverse-time operating integral.

$$\Lambda_r(t) = \int_{t_0}^{t} \frac{d\tau}{T_{\text{op},r}(I_r(\tau))}, \qquad \text{trip when } \Lambda_r(t) \geq 1$$

$$\varphi_{\text{security},r} = \square_{[t_f,\, t_r + T_{\text{rec}}]}\big[\Lambda_r(t) < 1\big]$$

Protection security is implemented in two different ways:

- Static magnitude cap that builds protection limits.Margin is defined as

$$\kappa = 1 - \rho^{max}$$

$$I(P, Q, v_i) <= \kappa * I_{pickup(state)}$$

$I(P, Q, v_i)$ is a fixed current ceiling per VRT state,used for smooth-sag and ERCOT sweeps.

- Time-graded IEC operate-integral ceiling builds time-graded limits. It tracks $\Lambda_r(t)$ directly and derives a per-step current ceiling that shrinks as TDS is set higher. Its floor,TDS_{min}, is the time-graded analogue of κ/ρ^{max}. Stage-2 LP minimizes each relay's TDS subject to TDS $\geq$ TDS_{min} and a 0.30-s coordination-time interval at a reference fault current. Importantly, this LP is solved by VRT state and setting basis; it is not re-solved for each facility loading or each individual sag event.

$$\text{TDS} \geq \text{TDS}_{\text{min}}, \qquad t_{\text{backup}}(I) - t_{\text{primary}}(I) \geq \text{CTI}$$

2.4.3.2 Formulation

The model facility is a 10-converter fleet behind a 34.5-kV medium-voltage bus: two 20-MVA BESS converters and eight 25-MVA UPS converters, each serving a 20-MW compute node, connected to the grid through a 138/34.5-kV interface. The disturbance trace used is the ERCOT NOGRR282 worst-case staircase trace. The controller classifies each timestep as NORMAL, $MILD_{VRT}$, $DEEP_{VRT}$, or RECOVERY using grid voltage and whether P and/or Q has to be dispatched from the converter fleet to sustain ride through. The state determines the applicable VRT floor and protection limits. Wayne et al. (2026)'s SolVRT casts LVRT ride-through as a Signal Temporal Logic (STL) specification evaluated over the sag trajectory, then synthesizes a per-timestep dispatch through a convex QCQP that maximizes the specification's robustness margin. Core per-timestep dispatch problem , minimizing dispatch cost subject to:

$$\begin{aligned}
&P^2 + Q^2 \leq S^2_{\text{rated}} && \text{(converter rating)}\\
&P \leq P_{\text{series-UPS,cap}} && \text{(series-UPS power cap)}\\
&I(P, Q, v_i) \leq I_{\text{ceil}}(\text{state}) && \text{(protection constraint } \phi_{\text{prot}})\\
&\text{SOC} \geq \text{SOC}_{\text{min}} && \text{(discharge floor)}\\
&v_i \leq v_{\text{max}}
\end{aligned}$$

plus, per operating state, a floor on either terminal voltage $(v_i \geq v_{trip}, MILD_{VRT}/DEEP_{VRT})$ or active power $(P \geq target_{floor}, RECOVERY)$

We implement a two-stage lexicographic slack. The voltage/power floors can conflict with φ_{prot} under deep sags, the solver introduces nonnegative slacks ξP, ξV and solves in two stages.

- Stage A — minimize infeasibility:

$$\min_{\xi P, \xi V \geq 0} \ \xi P w P + \xi V w V \quad \text{s.t. hard_cons}, \ v_i \geq v_{\text{trip}} - \xi_V, \ P \geq \text{target_floor} - \xi P$$

- Stage B — with ξP^*, ξV^* fixed from Stage A, minimizes the original dispatch cost subject to the same hard constraints plus the now-fixed floors.

Stage A is feasible by construction (P=Q=0) always satisfies $hard_{cons}$), so the two-stage solver never fails by itself. It degrades gracefully to a nonzero ξ rather than returning infeasible. Crucially, φ_{prot} is never slacked in either stage: it is a hard constraint throughout, reflecting that, for a given ACTIVE relay setting group, all protection properties are hard constraints. The relay setting itself is an engineering design variable, not an immutable floor. Adaptive setting-group selection is what actually changes which hard constraint applies, and that selection is a supervisory-layer decision, not a relaxation of φ_{prot} inside the optimization.

2.5 Evaluation

The case studies were developed and interpreted iteratively by the interdisciplinary author team, several of whom have domain expertise in the represented applications. These experts assessed the technical plausibility and domain relevance of the generalizers, specializers, and trade-offs surfaced by MCDM-TAGS; the resulting findings are reported in the case-study analysis. This expert-informed, author-led assessment was informal and was not an independent usability or controlled user study.

Table 5: MTDC–ESS co-design results. Solutions are ordered by decreasing net profit.

Future realization	Total cost ($M)	Expected revenue ($M)	Net profit ($M)
CCD22-5	1136.35	12046.48	10910.13
CCD22-3	1136.35	12046.30	10909.95
CCD22-8	1104.18	12007.77	10903.60
No Reserve	1055.91	5244.36	4188.46
CCD18-5	1104.17	5233.06	4128.89
CCD18-3	1104.17	5232.99	4128.82
CCD18-8	1102.60	5231.03	4128.42
No ESS	1055.89	5146.72	4090.83
Base	1149.40	5067.40	3918.00

2.5.1 MTDC-ESS Co-Design

The MTDC-ESS formulation produces alternatives across scenario realizations rather than a Pareto front. We use it to show how economic assumptions and physical constraints create tradeoffs.

The PCP (Figure 6) exposes cross-objective rank patterns across capital costs and multi-market revenues at every location. The Tesla heatmap (Figure 7) groups recurring specialization and tradeoff patterns, while the Euler diagram (Figure 8) makes their shared memberships and intersections explicit. Together, the views connect individual rank trajectories with recurring operational strategies.

Balanced infrastructure sizing can hedge market volatility—In Table 5, we observe that *Net Profit* heavily skewed results toward the extreme price spikes of 2022. The PCP makes the contrast visible: 2018 configurations (e.g., CCD18_3 and CCD18_5) avoid extreme ranks across objectives, whereas 2022 configurations gain top revenue ranks at the cost of poor Battery Cost ranks from aggressive ESS sizing. The 2018 configurations therefore hedge competing objectives rather than being universally optimal. This finding is not exclusive to MCDM-TAGS: tabular inspection of the objective-performance matrix or a PCP can reveal the pattern; TOPSIS can evaluate it under specified weights and normalization, and SMAA can test acceptability across uncertain weights. MCDM-TAGS instead provides parameter-free screening before those preferences are available, identifying configurations that avoid poor ranks rather than recommending a preferred alternative. MCDM-TAGS also isolates the WCASCADE

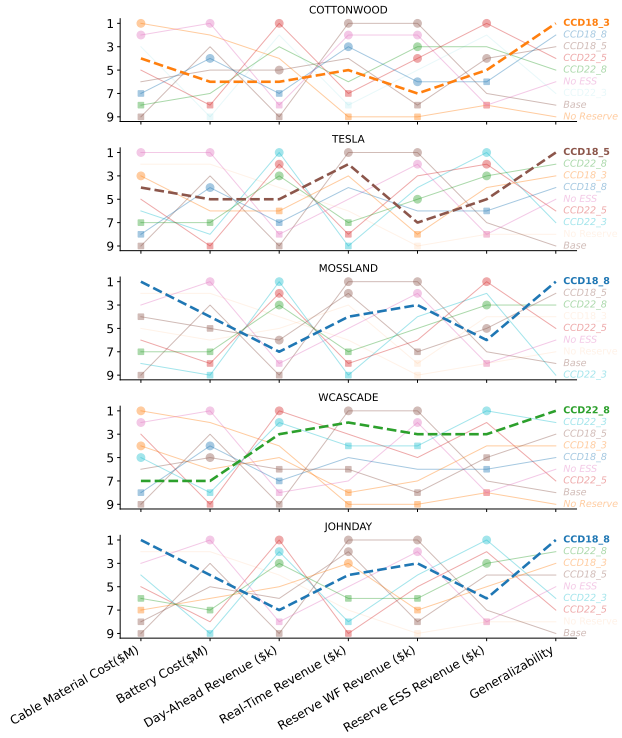


Figure 6: Parallel coordinates plot showing the relative performance ranks of all design configurations across physical and economic objectives.

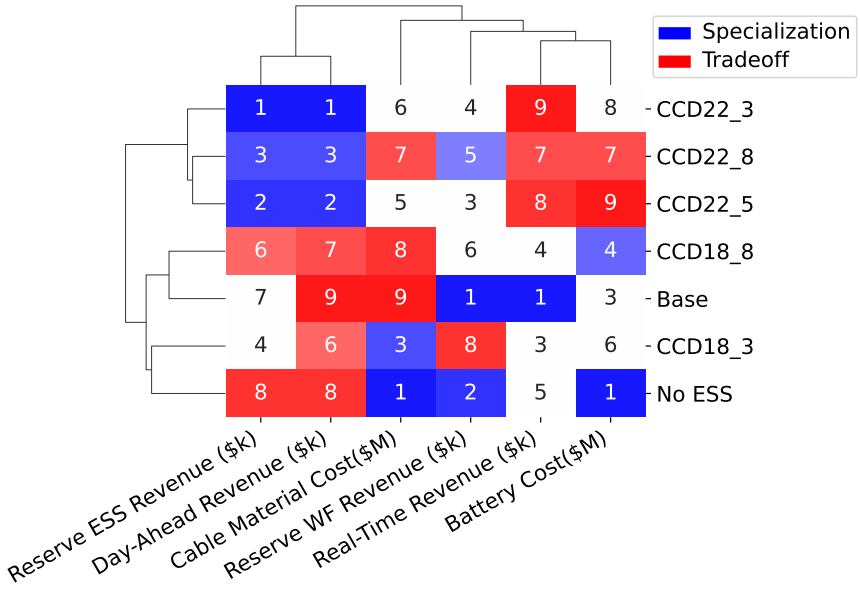


Figure 7: Cluster heatmap grouping similar operational strategies at the Tesla node.

anomaly, the only site where CCD22_8 is the generalizer. At this smallest site (1500 MW), the 5% ESS cap limits the battery to 75 MW. The smaller absolute capital penalty, combined with 8% inflation, lets WCASCADE exploit 2022 volatility without the Battery Cost penalty observed at larger sites.

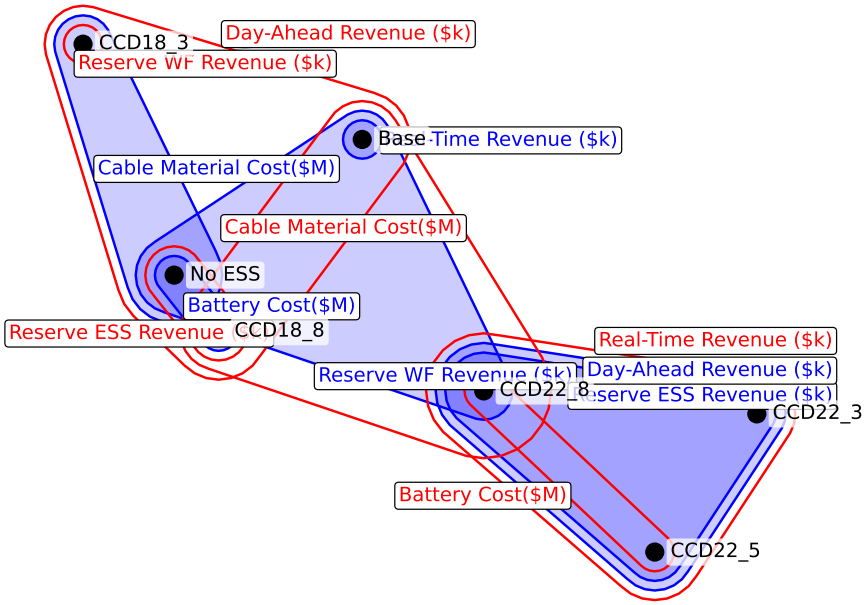


Figure 8: Euler diagram mapping operational strategies at the Tesla node.

Aggressive day-ahead commitments sacrifice real-time flexibility—Across all locations, 2022 configurations (especially CCD22_3 and CCD22_5) specialize in Day-Ahead (DA) revenue at the expense of Real-Time (RT) performance. The tradeoff is visible in the PCP and becomes more evident in the cluster heatmap plot for Tesla interconnect, where these aggressive configurations form a tight, dark blue specialization cluster within the DA Revenue column. However, tracing this exact same group to the RT Revenue column reveals a severe, dark red tradeoff cluster. The Euler diagram structurally confirms this mutual exclusivity. The nodes for CCD22_3 and CCD22_5 sit entirely within the blue specialization hull for DA Revenue, which itself is completely enveloped by the red tradeoff hull for RT Revenue. To dominate the DA market, the optimizer forward-commits generation and battery discharge, leaving little energy headroom for RT clearing or forecast error. WCASCADE remains the exception because its capacity limit prevents the battery from being drained in this way.
Rigid control strategies incur severe opportunity costs in energy markets—PCP also highlights key challenges of traditional, fixed-rule operational strategies across different scenario realizations. For instance, in *Base* case, the formulation was hard-coded to provide maximum primary frequency support. While the parallel coordinates show that it succeeded in its narrow objective, the respective tradeoff surfaces in supplementary visualizations. In the Tesla heatmap (Figure 7), the *Base* configuration forms an isolated specialization cluster (blue) in Reserve WF Revenue and Real-Time Revenue, which perfectly correlates with a severe tradeoff cluster (red) in the Day-Ahead (DA) Revenue column. The Euler diagram geometrically confirms this financial lockout. The *Base* node sits isolated deep within the blue-filled specialization hull for Real-Time Revenue, with tradeoff in Day-Ahead Revenue. Reserving excessive capacity for grid support excludes the plant from the bulk energy market.

Ancillary service allocation forces a strict equipment tradeoff—The Euler diagram also reveals how the system allocates grid support. The specialization hulls (blue lines) for Reserve WF Revenue and Reserve ESS Revenue are largely disjointed, with one notable intersection: CCD22_8. The WF and the ESS are usually not used for providing significant reserve simultaneously. For aggressive configurations (CCD22_3, CCD22_5), higher price and lower inflation rate result in higher operation revenue, incentivizing installing larger battery. This allows the ESS to provide frequency reserves and lets wind turbines to generate more power. However, CCD22_8 breaks this tradeoff. Because its 8% inflation reduces overall current value of operation revenue, the cost of large ESS cannot be covered. With relatively smaller ESS, CCD22_8 uniquely occupies the intersection of both sets, successfully providing top-tier frequency support from both the battery and the turbines simultaneously.

2.5.2 Large Electric Load Case Study

For the large electric load integration formulation, we demonstrate how MCDM-TAGS exposes how geographic constraints and compute profiles force critical tradeoffs. The PCP in Figure 9 places ranks for expansion cost, expected data-center load shed, and generalizability side by side for each weighting scenario. Its trajectories reveal whether balanced generalizer performance avoids poor ranks across all objectives or whether a reliability gain is paired with a cost tradeoff.

The industrial park with a standard compute profile acts as the strategic generalizer—Across the three PCP panels, the `industrial_park–standard` trajectory demonstrates highly balanced performance by avoiding extreme costs and poor reliability. It bypasses the heavy grid congestion of the `urban_edge`, where existing power lines operate near maximum capacity and force expensive upgrades for even minor capacity additions. Simultaneously, it benefits from stronger foundational infrastructure than a `rural_tech_hub`, saving the immense costs if building entirely new, long-distance power lines from scratch. Operationally, a `standard` data center demands steady, predictable power, sidestepping the massive energy spikes caused by AI/ML workloads. Because of these moderate baseline conditions, the grid does not require massive, pre-emptive expansions. Instead, the optimizer selects only a few targeted upgrades, achieving the lowest average expansion cost of $1.95 million (adding just 5 transmission lines). This ensures high reliability while entirely avoiding the $4+ million investments required to support demanding AI workloads in crowded areas.

AI compute loads incur a reliability premium—The PCP reveals the reliability premium for AI/ML-focused loads, pairing better load-shed ranks with worse expansion-cost ranks. AI/ML data centers create massive, unpredictable power spikes driven by intense computational bursts. To guarantee stable power delivery and prevent load shedding during these volatile peaks, the optimizer upgrades capacity, driving up expansion costs. This financial tradeoff becomes worse at the `urban_edge`. Metropolitan grids are heavily congested, and restricted physical space prevents the construction of new transmission corridors. To safely accommodate sudden AI power surges, operators are forced to upgrade existing infrastructure. Consequently, minimizing load shed (down to 15.69 MW) at the urban edge incurs a massive premium, driving average expansion costs to $4.02M.

Reliability challenges of high redundancy—`High_redundancy` configurations struggle to achieve optimal load shedding metrics (Figure 9). Designed for critical applications, these facilities maintain high, flat baseline loads with minimal variability (Table 3). This rigid demand leaves the grid with almost no operational headroom during contingencies, forcing the optimizer to mandate expensive, high-capacity upgrades to avoid severe load-shedding penalties. This inflexibility is severely amplified when the facility is sited in a `rural_tech_hub`. Isolated from

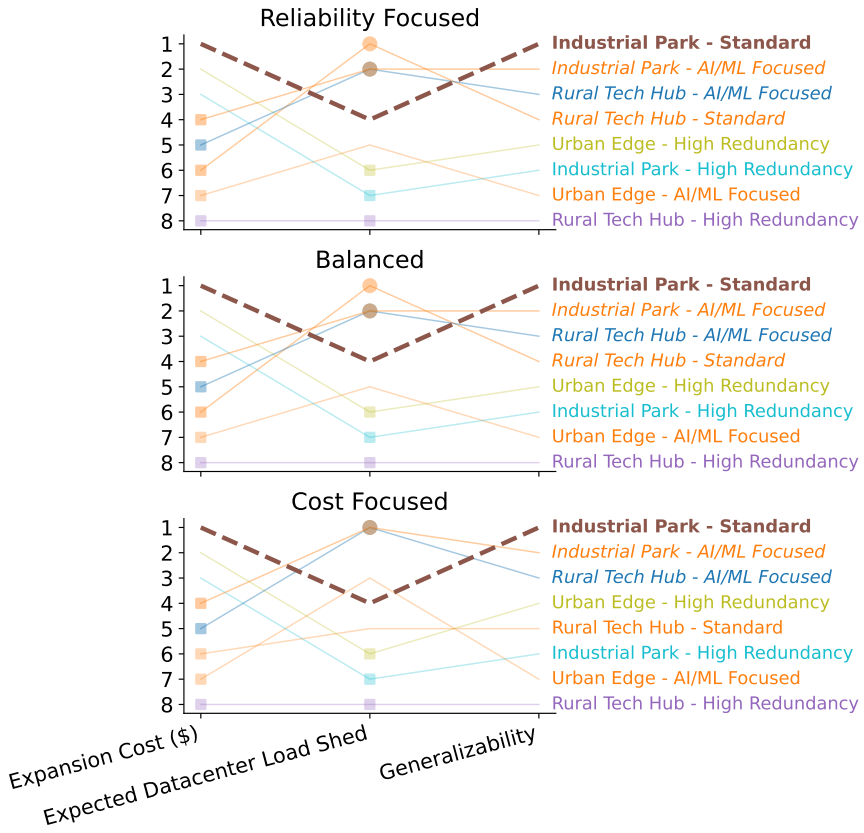


Figure 9: Parallel coordinates visualization of the rank matrix for the Large Electric Load (LEL) case study across three objective weighting scenarios: reliability-focused, balanced, and cost-focused.

major load centers and constrained by weak existing infrastructure, rural grids must construct entirely new transmission lines to guarantee continuous power delivery. Ultimately, even with a massive $4.78 million average expansion investment, the system cannot fully overcome these geographic limitations, resulting in a poor expected load shedding of over 28 MW.

Across both case studies, cost is one component of a broader value proposition rather than the sole decision criterion. In the MTDC-ESS case, the modeled alternatives exchange capital cost for market-revenue opportunities, frequency-support capability, operational flexibility, and performance across uncertain market conditions. In the LEL case, transmission investment purchases reliability value by reducing expected load shedding under contingencies. MCDM-TAGS exposes these relationships without assigning stakeholder-specific values or selecting a preferred alternative. Determining whether the modeled benefits justify their costs still requires stakeholder preferences and, where appropriate, additional measures such as the monetary consequences of interruptions; those judgments are outside the evidence provided by these case studies.

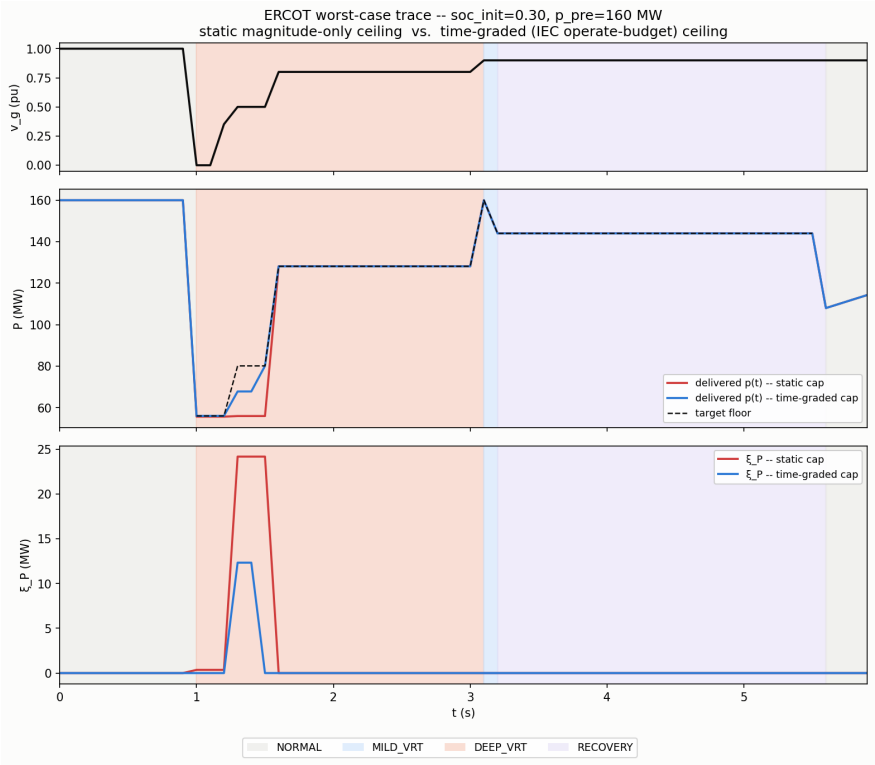


Figure 10: Worst case scenario demonstrating slaking for static cap and time-graded ceiling

2.5.3 Protection-aware Voltage Ride-Through Study

Comparing the static κ cap with the time-graded IEC operate-integral ceiling by sweeping initial SOC × facility loading P_{pre} × TDS_{MIN} against the real ERCOT trace gives a materially better view of how the co-design synthesis will look like. 48/64 points fully feasible, 16/64 feasible-if-relaxed, 0/64 infeasible regardless — versus 0/64 fully feasible under the static cap. Dominant relay among violated points: 100% UPS primary marker shape encodes the dominant binding relay group (BESS primary, UPS primary, aggregate backup, or none). On-plot labels give the specific relay and its own coordinated TDS at the most severe points. The

relay-coordination layer is overlaid directly on the feasibility surface.

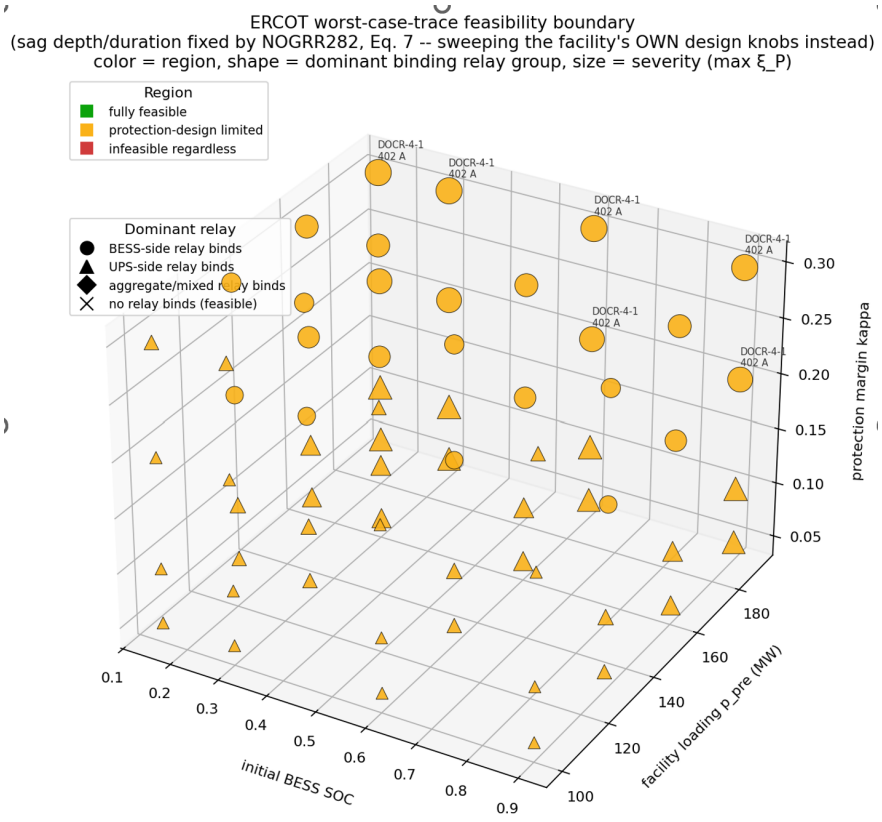


Figure 11: Static cap protection trace feasibility boundary

For the smooth-sag sweep, only 16/80 cases are fully feasible under the static protection ceiling, while 64/80 are protection-design limited and none remain infeasible after removing the protection ceiling. Under the actual ERCOT worst-case trace, the static ceiling produces 0/64 fully feasible cases. This demonstrates that the selected protection model can dominate the apparent VRT boundary.

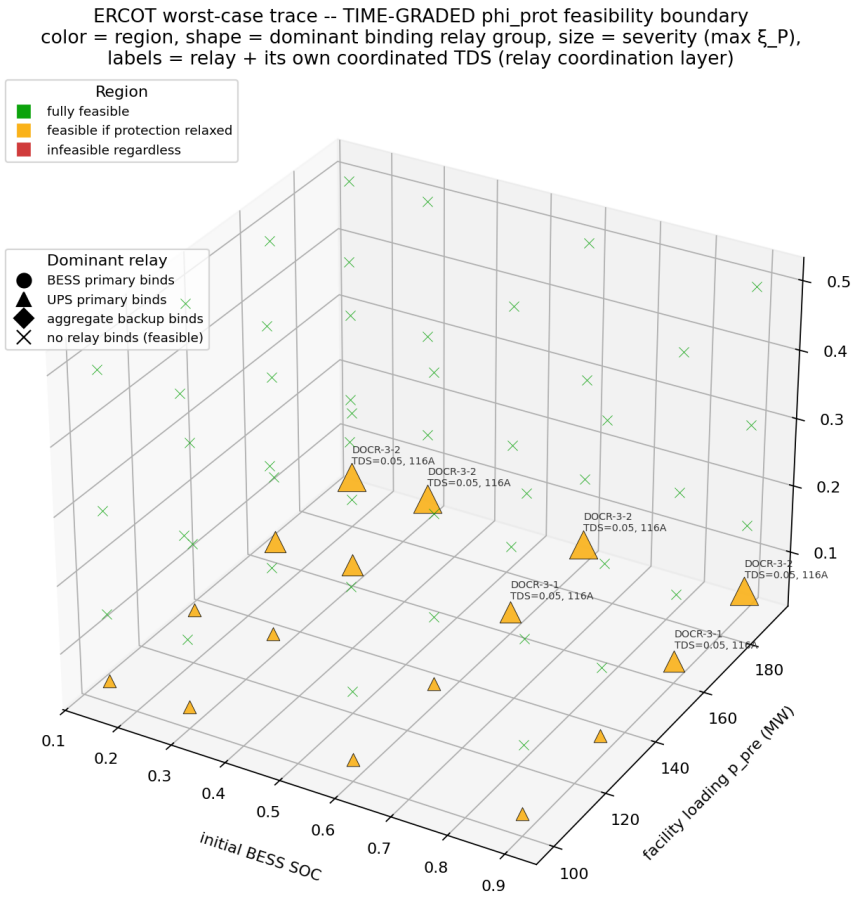


Figure 12: Time-graded protection trace feasibility boundary

$DEEP_{VRT}$ and RECOVERY fail for different reasons. During $DEEP_{VRT}$, the electrical demand is approximately 4.5-4.7 times the fixed pickup across the tested loading range. Because this ratio is nearly independent of facility loading, the problem is identified as a design-basis mismatch. During RECOVERY, the demand-to-pickup ratio rises from approximately 0.43 to 1.12 as facility loading increases from 100 to 190 MW, revealing a genuine loading-dependent limitation. Comparatively, the time-graded protection feasibility boundary provides a much improved result still with the UPS primary relays still not binding and infeasible at certain loading and battery state of charge. However, for $\kappa \geq 0.2$, all relays bind and can be seen as safe operating boundary.

The work is still in progress and hasn't been evaluated for multi-criteria decision making. The collected results were synthesized with near real values using the LinDistFlow model of a data center. Further work is needed to simulate by design, fault scenarios that can validate and improve the results achieved. Below are the list of tasks to be performed.

- Simulate a real fault-current model. Replace the $IBR_{FAULTMULTIPLE}$ = 1.5× rated place-holder with a converter-current-injection-based model consistent with IEEE 2800-2022 (1.1–1.2× rating for grid-following/grid-forming IBRs under fault), including sequence quantities (I_1, I_2, I_0, V_1, V_2, V_0).

- Simulate dependability and coordination from design-time-verified to runtime-monitored. Both

are currently checked offline (via the Stage-2 LP and post-hoc trip-time inspection); neither is enforced online inside the solver.

- Validate a real VRT-vs-fault discriminator before any adaptive setting-group switch is deployed. Classified state signal cannot yet distinguish a legitimate VRT event from a genuine internal fault, and a balanced 3-phase bolted fault defeats sequence-unbalance-based discrimination specifically. This report keeps the setting-group-switch concept and its dispatch-feasibility evidence but does not recommend enabling it until this is resolved.

- A full multi-location, multi-fault-type short-circuit study (3-phase, line-line, line-line-ground, line-ground, at multiple points on the feeder and MV bus) would be required before the protection settings in this report can be deployment ready.

2.5.4 Frontend evaluation

Performance was evaluated on a local deployment (frontend: `localhost:3000`, backend: `localhost:8080`) using two built-in formulations available in the dashboard interface: `Cameo_datacenter` and `MoCoDo_v2`. Each measurement was repeated three times under identical conditions on macOS 26.6.2 (Apple M2 Pro, 16 GB RAM). For each formulation, we profiled three representative user tasks using wall-clock timing: (i) formulation initialization via `GET /api/init?use_case=<formulation>`; (ii) the dashboard load sequence via `GET /api/init`, `GET /api/objective`, `GET /api/solutions`, and `GET /api/lmp`; and (iii) a filter/selection interaction via `GET /api/init`, a filtered `GET /api/solutions` request (first available filter value), and `GET /api/specializers?min_specializers=7`. Table 6 reports the results; values are the mean and min–max runtime across repetitions. Formulation size descriptors (no. of objectives, no. of decision variables, and # solutions) are reported for context. Dimensionality reduction uses CPU UMAP (via the `umap-learn` API), i.e., no GPU-specific UMAP backend is assumed.

Table 6: Reproducible performance evaluation across formulations and tasks. Each task was repeated three times under identical conditions; runtime is wall-clock latency per task. The exact API call sequence for each task, and the dimensionality-reduction backend, are given in the accompanying text.

	Cameo_datacenter	**MoCoDo_v2**
Objectives	3	1
Decision variables	1	2
Solutions	360	160
Runs per task	3	3
Load formulation		
Runtime avg (s)	0.005	0.004
Runtime min–max (s)	0.004–0.005	0.004–0.005
Load dashboard		
Runtime avg (s)	3.838	0.475
Runtime min–max (s)	3.813–3.852	0.467–0.488
Filter/select options		
Runtime avg (s)	3.804	0.460
Runtime min–max (s)	3.801–3.808	0.456–0.463

3.0 CAMEO: A Co-design Architecture for Multi-objective Energy System Optimization

3.1 CAMEO: Research Capabilities

CAMEO streamlines complex energy-system co-design by turning domain-expert knowledge into automated, ready-to-run optimization workflows without requiring extensive software development and scripting. Users simply fill out structured forms, and CAMEO builds the full pipeline, sets up execution environments, and applies smart checks to ensure correctness (Meyur et al. 2025). Under the hood, CAMEO has developed the following key capabilities to modularize the process:

- Specification Language to Modularize Optimization Problem Execution: CAMEO's specification language allows experts to define optimization studies through simple JSON-based inputs instead of code. Each form entry is translated into modular workflow components—such as data loaders, scenario generators, solvers, and summarizers—linked together with standardized interfaces. This makes it easy to reconfigure studies, swap components, and maintain reproducibility without writing scripts.

- High-Performance Computing Workflow to Execute Optimization Problems: The specification language provides a standard approach to defining a wide variety of classes of optimization problem formulations. From the specification, CAMEO automatically builds a scalable HPC-ready Nextflow pipeline. CAMEO defines *stages* to describe each step in the pipeline. It sets up containerized environments, manages stage ordering and dependencies, and runs large multi-scenario studies across clusters, cloud platforms, or local machines. Experts can conduct extensive co-design analysis at scale without handling workflow mechanics.

- Generative AI-Based Assistant to Develop, Validate, and Explain Optimization Problems: CAMEO strives to help domain experts during the co-design life-cycle to expedite entity design and optimization via state of the art Large Language Models (LLMs). CAMEO includes an AI assistant that helps users create optimization studies accurately and confidently. It guides form completion, identifies inconsistencies, explains workflow structure, and ensures specifications are logically sound before execution. The assistant reduces errors, increases transparency, and makes advanced optimization workflows accessible to non-programmers.

3.2 CAMEO workflow for Large Electric Load (LEL) use-case

Meyur et al. provide additional details of the platform and demonstrate synthetic distribution network generation and reliability-aware data center integration use cases (Meyur et al. 2025). The project also developed a Large Electric Load workflow to demonstrate scenarios defined in 4.0. The workflow is structured to run nine parallel processes, each representing a unique combination of demand response (DR) and voltage ride-through (VRT) signal settings. The process begins by reading hour block settings for DR (such as morning, afternoon, and midnight configurations) and 20-minute block settings for VRT. These settings are used to generate three sets of DR signals and three sets of VRT signals, resulting in nine possible combinations. Each combination is paired with an hourly load profile CSV file, which serves as input for solving a multi-objective design optimization problem. The workflow automates the generation of DR and VRT signals, optimization scripts, and dependencies, ultimately producing nine optimal designs—one for each scenario.

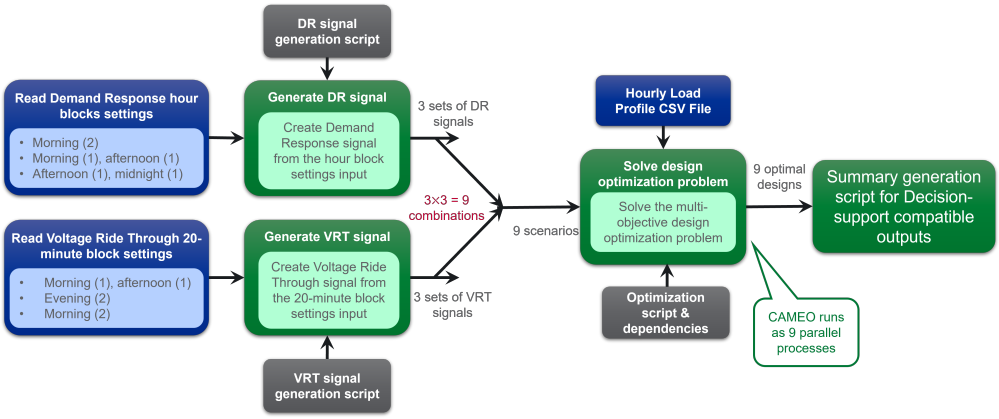


Figure 13: CAMEO workflow for Large Electric Load use-case

Figure 13 emphasizes CAMEO's ability to streamline complex optimization tasks by automating signal generation, scenario management, and summary output creation for decision support. After the optimization problems are solved, a summary generation script produces outputs compatible with decision-support systems, as shown in Figure 16, enabling stakeholders to quickly interpret results and make informed decisions. This approach not only accelerates the analysis of multiple operational scenarios but also ensures consistency and reproducibility in the optimization process. By integrating these automated workflows, CAMEO empowers domain experts to efficiently explore a wide range of design solutions, supporting robust planning and operational strategies for large electric loads.

3.3 Generative AI-based assistant to automate workflow generation

Large language models (LLMs) present a novel opportunity to leverage their advanced reasoning and code generation capabilities to automate the workflow generation process. The project builds upon the existing workflow generation pipeline and develops a multi-agent framework to automate the end-to-end co-design life-cycle. The resulting AI-assistant has a swarm of autonomous agents with a human-in-the-loop design that enables an iterative review loop, as shown in Figure 14. This assistant enables domain experts to generate, validate, and explain optimization problems without writing code, streamlining the process from form input to executable pipelines. By leveraging AI, CAMEO ensures that workflows are correctly structured, dependencies are managed, and mistakes are minimized, making it easier for users to focus on decision-making rather than technical implementation. This capability is deployed as a web-application with an easy-to-use chat interface, coupled with a graph-based visualization as shown in Figure 15. The application is central to CAMEO's mission of transforming expert engineering knowledge into automated, scalable solutions for large-scale infrastructure planning and operations.

3.4 Leveraging Large Language Models for optimization

Large language models (LLMs) also present a novel opportunity to incorporate human expertise directly into analytical optimization processes. This project explores the integration of modern LLMs with classical Bayesian optimization techniques to solve engineering design challenges.

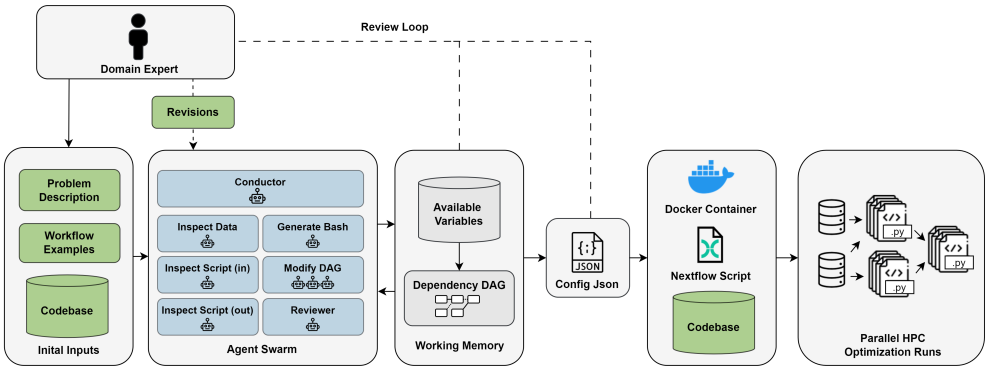


Figure 14: Multi-agent framework to automate workflow generation

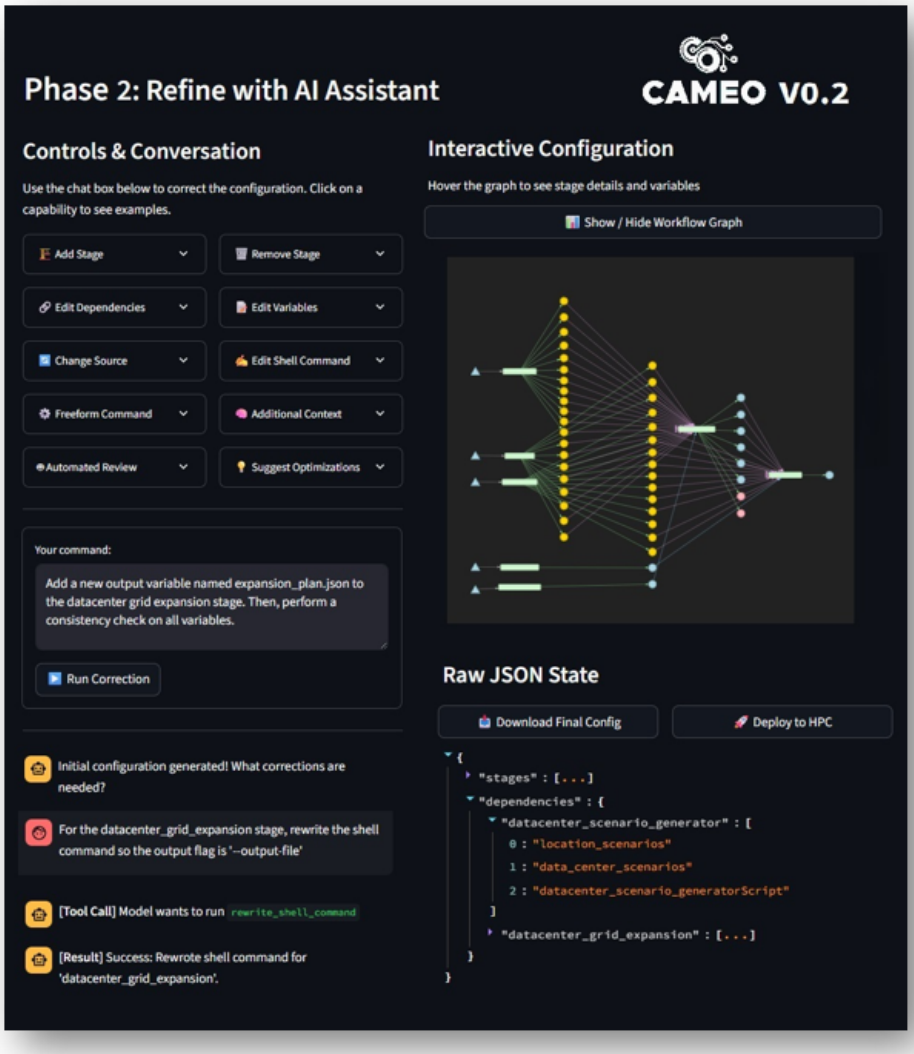


Figure 15: Generative AI-based assistant to automate workflow generation

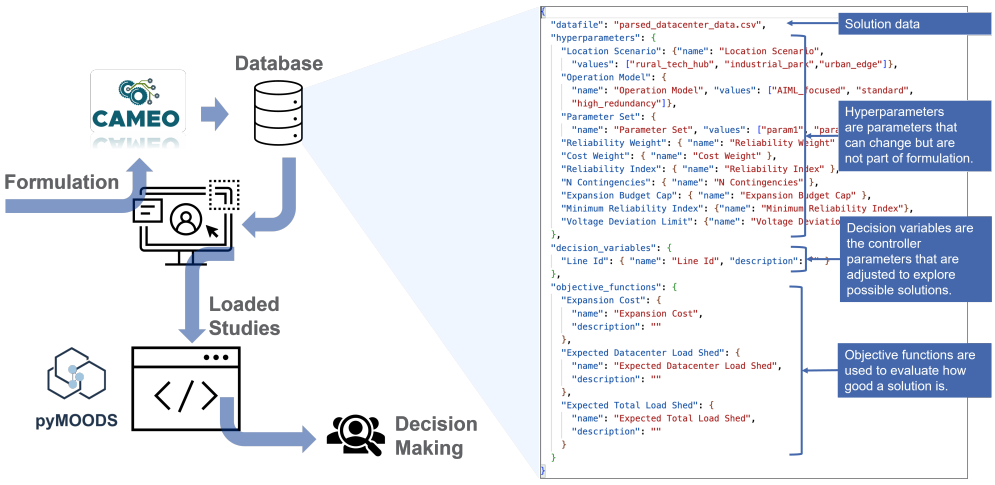


Figure 16: CAMEO generated summary for Decision Support Systems

Bayesian Optimization is used to efficiently optimize expensive, black-box functions by maintaining a surrogate model and iteratively selecting evaluation points, which is crucial when each experiment is costly. Traditionally, domain experts manually encode prior knowledge into parameter designs and constraints, but CAMEO leverages LLMs to automate and expedite this process—translating natural language domain expertise into quantitative signals for optimization. This integration enables early convergence in experimental campaigns, bridging the gap between human intuition and computational efficiency, and ultimately streamlining the design and optimization of complex engineering systems.

3.5 Decision Support System Integration

Integrated tool design and architecture enable how we link CAMEO's optimization studies to pyMOODS' decision support: it clarifies the roles of hyperparameters (tunable settings outside the mathematical formulation), decision variables (controller parameters adjusted to explore solutions), and objective functions (metrics that evaluate solution quality). The integrated capability also shows how solution data flows through core components—Formulation, Database, Loaded Studies, and Decision Making—to enable interactive analysis, as shown in Figure 16. Practically, CAMEO establishes the defined data schemas and executes studies; pyMOODS ingests those studies, presents them in an interface, and guides user interactions to synthesize tradeoffs and support near-real-time decisions for the use cases.

4.0 Multi-Objective Co-Design for Data Center Energy Storage Planning for Providing Grid Flexibility

Investment in artificial intelligence technologies has expanded rapidly across the United States, generating a substantial increase in electricity demand. According to the 2024 Report on U.S. Data Center Energy Use, data centers accounted for approximately 4.4% of total U.S. electricity consumption in 2023. This share is projected to rise to between 6.7% and 12% by 2028. The report further shows that annual data center electricity consumption grew from 58 TWh in 2014 to 176 TWh in 2023 and is expected to reach between 325 and 580 TWh by 2028 (Shehabi et al. 2024).

The continued expansion of large-scale data centers (DCs) is therefore introducing considerable new electricity demand and placing additional stress on power system infrastructure. In many regions, connecting a new project to the grid can take several years (Snider 2026). In addition, the power consumption of large data centers can vary substantially and rapidly as computational workloads change. At the same time, utilities and grid operators are increasingly requiring new loads to operate in a grid-friendly manner and contribute to reliable system operation. Data centers equipped with demand flexibility, controllable load-ramping capabilities, and the ability to remain connected during voltage sags resulting from main-grid faults can transform large loads from passive grid burdens into active grid resources (Chen et al. 2020; Xie, Cui, and Wierman 2025; Takci et al. 2025). Consequently, the ability to provide grid support proactively has become an important consideration in the design of new data center facilities.

Unlike data-center diesel backup generators, which are primarily intended for emergency operation and are subject to restrictions on non-emergency use (U.S. Environmental Protection Agency 2026), battery energy storage systems (BESSs) provide a practical means of facilitating faster data center grid integration by enabling flexible grid withdrawals while supporting multiple flexibility and grid-support functions. A properly sized BESS can reduce the net load during demand response events, suppress short- and medium-term power variations, comply with load-ramping requirements, and help maintain service continuity during upstream voltage disturbances by providing voltage ride-through (VRT) support.

Several studies have examined the use of energy storage to improve data center operation and grid participation. Mamun et al. (2016) developed a multi-objective model that maximizes monetary savings from peak shaving while minimizing battery degradation. Their results revealed a substantial Pareto trade-off between these objectives. Dabbagh et al. (2017) examined the coordinated use of energy storage and workload shifting to reduce data center peak demand and showed that BESSs can mitigate short-duration demand spikes while preserving service quality. Shi et al. (2016) demonstrated that BESSs can lower electricity costs by combining peak shaving with participation in ancillary service markets, including frequency regulation. Cupelli et al. (2018) introduced control strategies that allow data centers to participate in demand response programs by coordinating heating, ventilation, and air-conditioning systems, delay-tolerant information technology workloads, and battery storage. Azizi et al. (2026) proposed a grid-forming BESS framework capable of attenuating rapid power fluctuations, improving voltage and power quality at the grid interface, and enhancing data center reliability. Xie, Cui, and Wierman (2025) developed voltage controllers for the internal power distribution networks of data centers to improve low-voltage ride-through performance and prevent data center disconnection during main-grid faults. Liu, Shin, and Deka (2026) presented a co-optimization framework for jointly scheduling computing workloads and BESS dispatch under utility-imposed peak-load and ramp-rate constraints. Thompson et al. (2016) proposed a BESS sizing framework that reduces microgrid costs while improving reliability for

data centers. Cupelli, Barve, and Monti (2017) developed a methodology for optimally sizing BESSs that provide frequency containment reserve services while accounting for battery operating constraints and energy management strategies.

The aforementioned studies primarily address the control and dispatch of existing data center BESSs or the sizing of storage systems for individual applications. Although BESSs can provide multiple capabilities to data centers, each service imposes distinct technical requirements on the storage system. Demand response, for instance, depends on both available power and sustained energy capacity. Fluctuation smoothing is instead dominated by the required power rating, whereas VRT support requires sufficient apparent-power capability to enable reactive-power injection. Consequently, identifying a suitable BESS size that can support multiple applications represents an important planning challenge for grid-interactive data centers.

This chapter introduces an optimization-based framework for sizing BESSs in grid-interactive data centers while explicitly representing the requirements associated with multiple services. The framework incorporates demand response support, suppression of fast load fluctuations, reactive-power capability for quasi-static voltage support during point-of-common-coupling (PCC) voltage sags, and slower operational functions, including net-load stabilization and load-ramping limitation. By integrating these requirements into a unified planning model, the proposed approach determines the minimum BESS power and energy capacities required for the data center to satisfy both its internal operating constraints and the grid-support requirements imposed by external parties.

The main contributions of this work are as follows:

- We formulate a multi-capability BESS sizing problem for data centers that simultaneously accounts for demand response, fast load fluctuation smoothing, voltage support at the PCC, and load-ramping constraints.
- We differentiate between services that require energy consumption and those that primarily require power headroom, thereby providing a more representative description of BESS capability requirements.
- We establish a scenario-based optimization framework for sizing the BESS under diverse operating conditions, including demand response events and main-grid fault events.

The remainder of this chapter is organized as follows. Section 4.1 defines the problem and describes the required grid-support functionalities. Section 4.2 presents the mathematical formulation of the proposed BESS sizing framework. Section 4.3 provides the case studies and discusses the numerical results.

4.1 Problem Description

The objective is to determine the appropriate BESS size for enabling data centers to provide multiple capabilities, including demand response, fast load fluctuation smoothing, and voltage support during main-grid faults, while satisfying operational requirements and constraints. These capabilities are provided primarily through the BESS, whose rated power and energy capacity are allocated strategically among the different services. The data center also maintains a queue of delay-insensitive tasks, which provides an additional source of operational flexibility.

Demand response capacity denotes the data center's ability to reduce its power consumption over a specified period. This capability is represented as a percentage of the data center's total rated load and depends on the availability of both BESS power and energy capacity.

Fast load fluctuations in data centers are sudden and rapid variations in power demand resulting from changes in computational workloads. These fluctuations generally occur at frequencies above 1 Hz and can have substantial magnitudes, potentially reaching 80% of the total rated IT load. The BESS mitigates these variations by charging during periods of low demand and discharging during periods of high demand. In this study, the charging and discharging actions associated with this service are assumed to be balanced and to occur on a subsecond timescale. Consequently, they produce no net change in the BESS state of charge (SoC). By reserving the upper and lower 10% of the SoC range—that is, excluding these portions from normal operation and other services—the BESS can accommodate the energy exchange required for fast fluctuation smoothing. The power rating required for this function is nevertheless substantial, amounting to approximately 40%–50% of the data center's total rated load.

Voltage support capacity refers to the reactive-power capability needed to maintain the voltage at the data center root node within specified limits during an assumed PCC voltage sag. When the PCC voltage decreases, the BESS injects reactive power to raise the root-node voltage, thereby helping maintain normal operation and protect connected equipment. Since this service is primarily reactive, it does not require net charging or discharging of the BESS. It does, however, require sufficient inverter apparent-power capacity to deliver the necessary reactive power. The analysis therefore determines the apparent-power headroom that the BESS must reserve for quasi-static voltage support.

During normal daily operation, the BESS can regulate the data center's load over timescales ranging from minutes to hours and constrain load changes to remain within the specified ramping limits. Supporting these functions requires adequate BESS rated power and energy capacity, together with appropriate SoC management throughout the operating period.

The detailed architecture of the data center and its distribution network is not explicitly represented in the model. Instead, the analysis considers aggregated data center load and BESS operation at a one-minute resolution. Device-level controls and the transient system dynamics associated with fast fluctuation smoothing and VRT are likewise outside the scope of the model. Nevertheless, the formulation ensures that sufficient BESS power, energy capacity, and SoC are available to satisfy the requirements of these functions.

4.2 Model Formulation

A scenario-based framework is adopted to represent a range of possible operating conditions. Each scenario may include one or more demand response (DR) events, PCC voltage sag events, or a combination of both. By considering multiple scenarios, the model accounts for variations in the timing and duration of these events. This formulation enables the determination of a BESS size that is adequate across a broad set of plausible operating conditions and reflects the uncertainty and diversity of events that may arise in practice.

4.2.1 BESS size limits

The decision variables E^{BESS} and P^{BESS} denote the BESS energy capacity and rated power, respectively. The upper bounds on these variables are imposed by

$$0 \leq E^{\text{BESS}} \leq \overline{E}^{\text{BESS}} \tag{1}$$

$$0 \leq P^{\text{BESS}} \leq \overline{P}^{\text{BESS}} \tag{2}$$

Because each scenario represents a distinct operating condition, the following operating constraints and decision variables are defined for every scenario $s \in \mathcal{S}$. To avoid repeatedly listing the index ranges, we omit them below. Constraints involving decision variables indexed by t apply to all time periods $t \in \mathcal{T}$ within the optimization horizon and subject to the appropriate boundary conditions. Similarly, constraints involving variables indexed by s apply to all scenarios $s \in \mathcal{S}$.

4.2.2 DC load operation constraints

The DC uses coordinated BESS operation and load-queue management to process IT workloads within the required time while satisfying the applicable grid constraints.

$$p^{\mathsf{q}}_{t,s} = p^{\mathsf{q}}_{t-1,s} + d^{\mathsf{IT}}_{t} - p^{\mathsf{IT}}_{t,s} \tag{3}$$

$$p^{\mathsf{IT}}_{t,s} \geq \hat{d}^{\mathsf{IT}}_{t} \tag{4}$$

$$p^{\mathsf{IT}}_{t,s} + p^{\mathsf{Cool}}_{t,s} + p^{\mathsf{ch_{op}}}_{t} - p^{\mathsf{dis_{op}}}_{t,s} = p^{\mathsf{DC}}_{t,s} \tag{5}$$

$$-P^{\mathsf{ramp}} \leq p^{\mathsf{DC}}_{t,s} - p^{\mathsf{DC}}_{t-1,s} \leq P^{\mathsf{ramp}} \tag{6}$$

$$p^{\mathsf{DC}}_{t,s} \leq \overline{P}^{\mathsf{DC}} - \alpha P^{\mathsf{R}} D^{\mathsf{dr}}_{t,s} \tag{7}$$

$$p^{\mathsf{DC}}_{t,s} \leq (1+\delta) \sum_{\tau=t-T^{\mathsf{s}}}^{t+T^{\mathsf{s}}} p^{\mathsf{DC}}_{\tau,s}/(2T^{\mathsf{s}}+1) \tag{8}$$

$$p^{\mathsf{DC}}_{t,s} \geq (1-\delta) \sum_{\tau=t-T^{\mathsf{s}}}^{t+T^{\mathsf{s}}} p^{\mathsf{DC}}_{\tau,s}/(2T^{\mathsf{s}}+1) \tag{9}$$

$$0 \leq p^{\mathsf{q}}_{t,s} \leq \overline{P}^{\mathsf{Q}} \tag{10}$$

The parameters d^{IT}_{t} and $\hat{d}^{\mathsf{IT}}_{t}$ denote the total and time-sensitive IT loads at time t, respectively. The parameter P^{ramp} represents the load-ramping limit, while $\overline{P}^{\mathsf{DC}}$ specifies the upper limit on the DC load. The rated DC load is denoted by P^{R}. The binary time series $D^{\mathsf{dr}}_{t,s}$ indicates whether a demand response request is active at time t in scenario s. The parameter δ defines the allowable relative deviation from the local average, and $2T^{\mathsf{s}}$ denotes the width of the smoothing window in time steps. The parameter $\overline{P}^{\mathsf{Q}}$ specifies the maximum amount of unprocessed load that may remain in the queue.

The decision variables $p^{\mathsf{q}}_{t,s}$, $p^{\mathsf{IT}}_{t,s}$, and $p^{\mathsf{DC}}_{t,s}$ represent the unprocessed queue load, processed IT load, and total DC load, respectively. The variable $p^{\mathsf{Cool}}_{t,s}$ denotes the cooling-system load, while $p^{\mathsf{ch_{op}}}_{t,s}$ and $p^{\mathsf{dis_{op}}}_{t,s}$ represent the BESS charging and discharging power associated with DC operation, respectively.

Constraint (3) describes the evolution of the load queue. Constraint (4) ensures that the time-sensitive portion of the IT load is processed on time. Constraint (5) enforces the DC load balance, and constraint (6) limits changes in the DC load between consecutive time periods. Constraint (7) imposes the maximum DC load with and without an active demand response request. Constraints (8)–(9) limit short-term load fluctuations using a moving-average formulation. Finally, constraint (10) limits the amount of unprocessed load that can remain in the queue.

4.2.3 BESS operation for demand response

By operating the BESS, the DC can adjust its load by α of the rated load level P^{R} for a period of T^{dr} hours. This demand response capability is represented by

$$p^{\mathrm{ch_{dr}}} = \alpha P^{\mathrm{R}} \tag{11}$$
$$p^{\mathrm{dis_{dr}}} = \alpha P^{\mathrm{R}} \tag{12}$$
$$p_{t,s}^{\mathrm{ch_{op}}} - p_{t,s}^{\mathrm{dis_{op}}} + p^{\mathrm{ch_{dr}}} \leq P^{\mathrm{Batt}} \tag{13}$$
$$p_{t,s}^{\mathrm{ch_{op}}} - p_{t,s}^{\mathrm{dis_{op}}} - p^{\mathrm{dis_{dr}}} \geq -P^{\mathrm{Batt}} \tag{14}$$
$$E^{\mathrm{dr}} = \alpha P^{\mathrm{R}} T^{\mathrm{dr}} \tag{15}$$

The decision variables $p_{t,s}^{\mathrm{ch_{dr}}}$ and $p_{t,s}^{\mathrm{dis_{dr}}}$ denote the BESS charging and discharging power allocated to load flexibility, respectively. The variable E^{dr} represents the energy required to provide demand response for a duration of T^{dr} hours.

Constraints (11)–(12) determine the BESS rated power required for demand response. Constraints (13)–(14) limit the combined BESS power associated with DC operation and demand response. Constraint (15) determines the energy capacity required to provide demand response.

4.2.4 BESS for fast load fluctuation smoothing

The BESS power capacity required to smooth fast load fluctuations, denoted by P^{fs}, is modeled as a fraction β of the DC's rated load. This relationship is given by

$$P^{\mathrm{fs}} = \beta P^{\mathrm{R}} \tag{16}$$

The parameter β can be adjusted according to the fluctuation characteristics of the DC workload.

4.2.5 BESS state dynamic constraints

The BESS operating and state dynamic constraints are formulated as follows:

$$SoC_{t,s} - SoC_{t-1,s} = (\eta^{\mathrm{ch}} p_{t,s}^{\mathrm{ch_{op}}} - \eta^{\mathrm{dis}} p_{t,s}^{\mathrm{dis_{op}}}) * T^{\mathrm{L}} \tag{17}$$
$$0 \leq p_{t,s}^{\mathrm{ch_{op}}} \leq P^{\mathrm{Batt}} z^{\mathrm{op}} \tag{18}$$
$$0 \leq p_{t,s}^{\mathrm{dis_{op}}} \leq P^{\mathrm{Batt}} (1 - z^{\mathrm{op}}) \tag{19}$$
$$0.1E^{\mathrm{BESS}} + E^{\mathrm{dr}}(1 - D_{t,s}^{\mathrm{dr}}) \leq SoC_{t,s} \leq 0.9E^{\mathrm{BESS}} \tag{20}$$

The parameters η^{ch} and η^{dis} denote the BESS charging and discharging efficiencies, respectively. The variable $SoC_{t,s}$ represents the BESS state of charge (SoC).

Constraint (17) describes the evolution of the BESS SoC. Constraints (18)–(19) prevent simultaneous charging and discharging of the BESS. Constraint (20) imposes SoC limits while accounting for both DC operation and demand response. The SoC is maintained between 10% and 90% of the total energy capacity to prevent overcharging and overdischarging. This operating range also allows for minor SoC variations associated with fast load fluctuation smoothing.

Constraint (17) represents the BESS SoC dynamic equation. Constraints (13)–(14) ensure that the BESS does not charge and discharge simultaneously, while constraint (20) enforces the SoC limits associated with DC operation and demand response. The SoC remains between 10% and 90% of the total energy capacity to prevent overcharging and overdischarging. This range also accommodates minor SoC fluctuations caused by fast load fluctuation smoothing.

4.2.6 DC PCC voltage equations

The constraints governing voltage support are presented below. The DC is represented as an aggregated point, and the transmission line connecting the PCC to the root node of the DC distribution system is modeled explicitly.

$$p^{l}_{t,s} = p^{\mathsf{DC}}_{t,s} \tag{21}$$

$$q^{l}_{t,s} = q^{\mathsf{vs}}_{t,s} \tag{22}$$

$$V^{0}_{t} - \gamma F^{\mathsf{vs}}_{t,s} - V^{\mathsf{DC}}_{t,s} = r p^{l}_{t,s} + x q^{l}_{t,s} \tag{23}$$

$$\left(P^{\mathsf{Batt}} + P^{\mathsf{fs}}\right)^2 + \left(q^{\mathsf{vs}}_{t,s}\right)^2 \leq \left(P^{\mathsf{BESS}}\right)^2 \tag{24}$$

$$V^{\mathsf{ref}} - \Delta V \leq V^{\mathsf{DC}}_{t,s} \leq V^{\mathsf{ref}} + \Delta V \tag{25}$$

The parameter V^{0}_{t} denotes the normal PCC voltage, and V^{ref} represents the reference voltage at the root node of the DC distribution grid. The parameter ΔV specifies the maximum allowable deviation from the reference voltage. The parameters r and x denote the resistance and reactance, respectively, of the main transmission line connecting the PCC to the DC root node. The binary time series $F^{\mathsf{vs}}_{t,s}$ indicates whether a PCC voltage sag occurs at time t in scenario s. When $F^{\mathsf{vs}}_{t,s} = 1$, the PCC voltage decreases by γ.

The decision variables $p^{l}_{t,s}$ and $q^{l}_{t,s}$ represent the real and reactive power flowing through the main transmission line, respectively. The variable $q^{\mathsf{vs}}_{t,s}$ denotes the reactive power supplied by the BESS, and $V^{\mathsf{DC}}_{t,s}$ represents the voltage at the root node of the DC distribution grid.

Constraints (21)–(22) enforce the real- and reactive-power balances at the root node of the DC distribution grid, respectively. Constraint (23) defines the power-flow relationship for the main transmission line. Constraint (24) imposes the total BESS power limit, and constraint (25) specifies the allowable voltage range at the root node of the DC distribution grid.

4.2.7 DC cooling power equation

Let PUE denote the DC's power usage effectiveness, and assume that nonessential loads are negligible. Under these assumptions, the cooling-system power is approximated by

$$p^{\mathsf{Cool}}_{t,s} = (\mathsf{PUE} - 1) p^{\mathsf{IT}}_{t,s} \tag{26}$$

4.2.8 Objective function

Two optimization settings are considered. In the first, the demand response capacity α and the maximum allowable PCC voltage drop γ are specified as parameters. These quantities represent grid requirements that the DC must satisfy to become interconnected. The objective is then to minimize the total BESS cost while meeting the specified requirements.

In the second setting, α and γ are treated as decision variables. The resulting problem is formulated as a multi-objective optimization that minimizes BESS cost while maximizing demand response and voltage support capabilities. The Pareto front is then generated to allow the decision maker to evaluate the trade-offs among these objectives.

4.2.8.1 Single objective formulation

Let C^E and C^P denote the unit costs of BESS energy capacity and rated power, respectively. The cost-minimization objective is formulated as

$$\begin{aligned} \min \quad & C^E E^{\text{BESS}} + C^P P^{\text{BESS}} \\ & + \sum_{t\in\mathcal{T}} \sum_{s\in\mathcal{S}} |V^{\text{DC}}_{t,s} - V^{\text{ref}}| \end{aligned} \tag{27}$$

The first two terms represent the costs associated with BESS energy capacity and rated power, respectively. The final term penalizes deviations of the DC root-node voltage from its reference value.

The BESS size pair $(E^{\text{BESS}}, P^{\text{BESS}})$ is selected to be sufficiently large to satisfy the requirements of all scenarios. The BESS operating schedule, however, varies according to the demand response requests and main-grid fault signals associated with each scenario.

4.2.8.2 Multi-objective formulation

The weighted-sum method is commonly used to solve multi-objective optimization problems. Let w^{BESS}, w^{DR}, and w^{VS} denote the weighting factors assigned to BESS cost, demand response capacity, and voltage support capacity, respectively. The resulting scalarized objective is

$$\begin{aligned} \min \quad & w^{\text{BESS}}(C^E E^{\text{BESS}} + C^P P^{\text{BESS}}) - w^{\text{DR}}\alpha \\ & - w^{\text{VS}}\gamma + \sum_{t\in\mathcal{T}} \sum_{s\in\mathcal{S}} |V^{\text{DC}}_{t,s} - V^{\text{ref}}| \end{aligned} \tag{28}$$

The Pareto front can be obtained by repeatedly solving the resulting optimization problem with different combinations of weighting factors. DC designers can then choose a solution from the Pareto front that provides an appropriate balance among BESS cost, demand response capacity, and voltage support capacity.

4.3 Data Center Case Study

4.3.1 Use-case setting

The configuration of the DC system is depicted in Fig. 17. The root node is connected to the main grid through the PCC, and both the total DC load and the aggregate charging and discharging power of all BESS units are represented at this node. The BESS size obtained from the optimization may correspond to the combined capacity of multiple devices distributed throughout the DC system. However, the network details beyond the root node, including the internal DC system topology, are not explicitly modeled in this study.

The original DC load data used in this case study are obtained from the MIT Supercloud dataset (Samsi et al. 2021). The total IT load over a 48-hour period, from Sep. 1 at 00:00 to

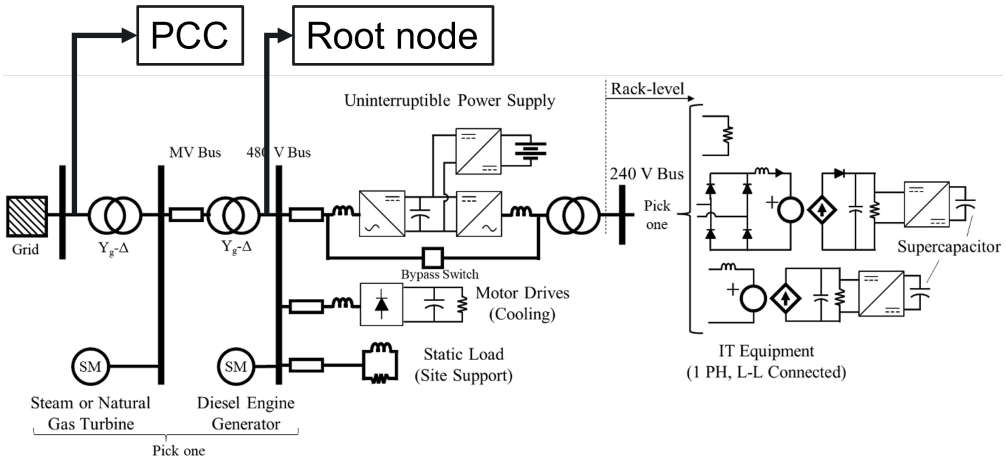


Figure 17: Data center grid connection and power supply architecture. The root node is connected to the main grid through the PCC. The total DC load and the aggregate charging and discharging power of all BESS units are represented at the root node.

Sep. 2 at 23:59, is extracted, resampled at a 1-minute resolution, and rescaled. The resulting load profile is presented in Fig. 18. We assume that 80% of the load is time-sensitive and therefore cannot be deferred. The PUE is set to 1.3, the rated load is set to 35 MW, and the upper limit on the total load is set to 30 MW.

The BESS cost is estimated using PNNL's Energy Storage Evaluation Tool[1]. The unit cost of energy capacity is set to $C^E = \$370/\text{kWh}$, while the unit cost of rated power is set to $C^P = \$105/\text{kW}$.

Most cases consider two representative operating scenarios. In the DR scenario, two demand response events are scheduled at 7:00 and 22:00 on Sep. 1. During each event, the DC reduces its load for one hour. A six-hour interval without an additional demand response event is assumed after each event to allow the BESS to recharge. In the PCC voltage sag scenario, voltage drops occur at 3:20 and 16:40 on Sep. 1, with each event lasting 20 minutes. The maximum allowable voltage deviation, ΔV, at the root node is set to 0.2 p.u. The event timings and durations used here are simplified assumptions for the case study. Although performance and sensitivity analyses are also conducted, more realistic demand response schedules and voltage sag durations would require detailed transmission system simulations and are therefore beyond the scope of this paper.

The following four cases are investigated. The resulting BESS sizing solutions and corresponding costs are summarized in Table 7.

- **Demand response only**: Only the DR scenario is considered, and the single-objective function (27) is optimized. When a demand response event is requested, the DC reduces its peak load by 20% of its rated load (i.e., $\alpha = 0.2$).

- **Voltage sag (VS) only**: Only the voltage sag scenario is considered, and the single-objective function (27) is optimized. The PCC voltage is assumed to decrease to 0.6 p.u.

- **Integrated DR and VS**: Both the DR and voltage sag scenarios are included, using the same parameter values as those in the corresponding single-scenario cases. The performance of

1. https://eset.pnnl.gov

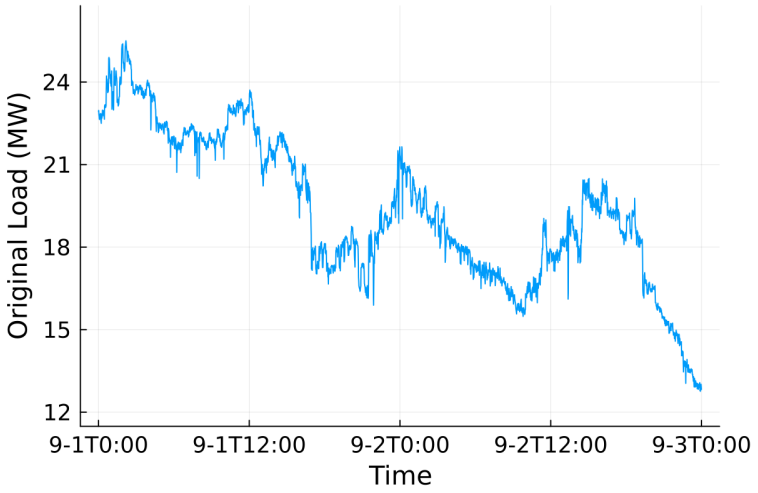


Figure 18: Data center IT load profile over a 48-hour period at 1-minute resolution.

the resulting BESS is evaluated by varying the interval between DR events in the DR scenario. In addition, the sensitivity of the BESS size is examined using randomly generated scenarios.

- **Multi-objective**: Both the demand response and voltage sag scenarios are considered, and the scalarized objective (28) is optimized for different combinations of weighting factors. The demand response capacity α is constrained to $[0.1, 0.3]$, while the allowable PCC voltage drop γ is constrained to $[0.2$ p.u., 0.4 p.u.$]$.

Table 7: BESS Size Solutions and Cost for Single-Objective Cases

	E^{BESS} (MWh)	P^{BESS} (MW)	Cost ($M)
DR Only	18.69	30.28	10.09
VS Only	9.94	78.54	11.93
DR + VS	18.69	80.82	15.40

4.3.2 Demand response scenario only

As reported in the first row of Table 7, considering only the demand response scenario yields an optimal BESS size of 30.28 MW/18.69 MWh and a total cost of $10.09 million. Relative to the voltage-support-only case, providing demand response requires a greater energy capacity. Because no main-grid fault occurs in this scenario, the BESS does not need to provide reactive power for root-node voltage support, resulting in a lower required power rating.

Figure 19 compares the DC's initial total load with its power draw when the BESS is deployed, while Fig. 20 presents the BESS energy-capacity allocation and SoC trajectory. As shown in Fig. 19, the BESS keeps the DC's total load below the 30 MW upper limit and enables it to decrease further to 23 MW when demand response is requested. The load profile with the BESS is also smoother than the profile without the BESS.

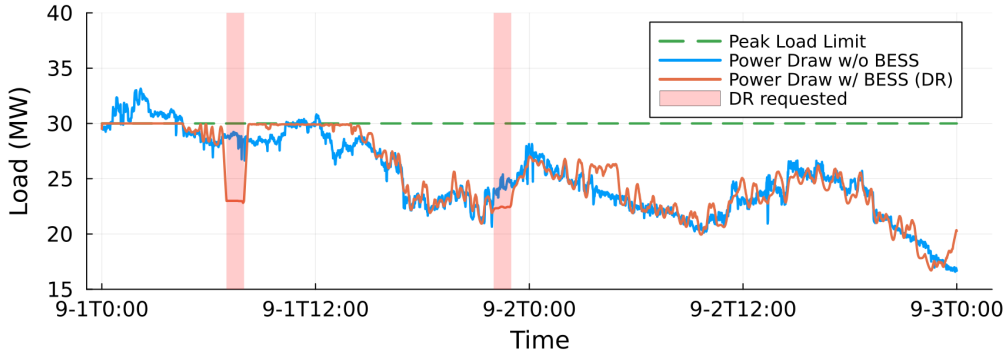


Figure 19: Data center net load in the demand-response-only case. The periods during which demand response events are requested are marked in pink.

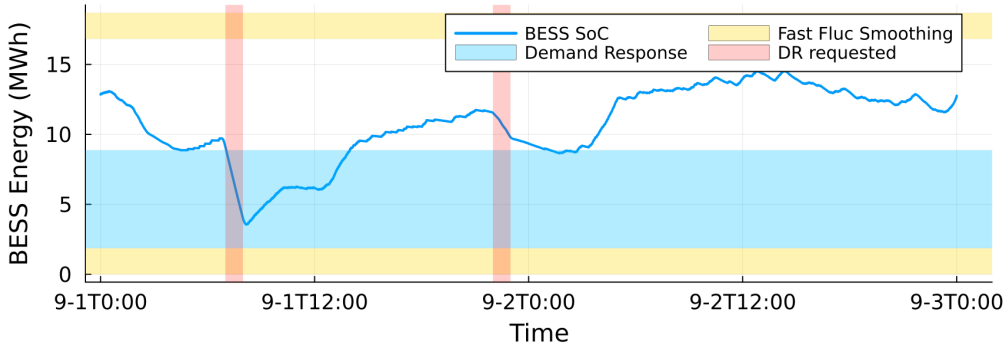


Figure 20: BESS energy-capacity allocation and SoC trajectory in the demand-response-only case. The energy capacity allocated for demand response is marked in blue. The yellow regions represent the upper and lower 10% of the total BESS energy capacity reserved to accommodate energy exchange for fast fluctuation smoothing.

According to constraint (15), 7 MWh of the BESS capacity is reserved for peak-load reduction during demand response events, as indicated by the blue region in Fig. 20. At the beginning of Sep. 1, the DC load is high, and the BESS discharges to maintain the 30 MW load limit. However, the energy reserved for demand response remains unused because of the SoC restriction in constraint (20). This restriction ensures that sufficient energy remains available whenever a demand response event is requested. The yellow regions in Fig. 20 denote the upper and lower 10% of the total BESS energy capacity. These portions are excluded from regular operation for safety and to accommodate the energy exchange associated with fast load fluctuation smoothing.

4.3.3 Voltage sag scenario

When only the voltage sag scenario is considered, the optimal BESS size is 78.54 MW/9.94 MWh, with a total cost of $11.93 million, as shown in the second row of Table 7. Because demand response is not included, the required energy capacity is substantially lower than that in the demand-response-only case. In contrast, the required power rating is considerably higher because the BESS must provide voltage support during voltage sag events.

Figure 21 shows the voltages at the PCC and the DC root node. The period from Sep. 1 at

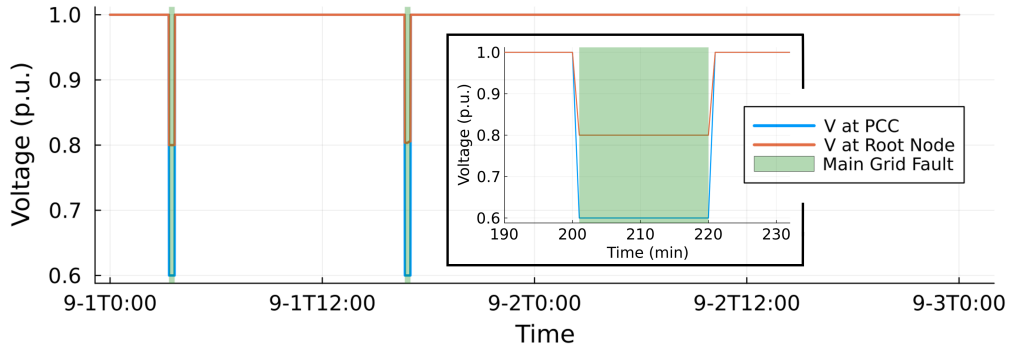


Figure 21: PCC and data center root-node voltages in the voltage-sag-only case. The interval from Sep. 1 at 3:10 am to 3:50 am is enlarged.

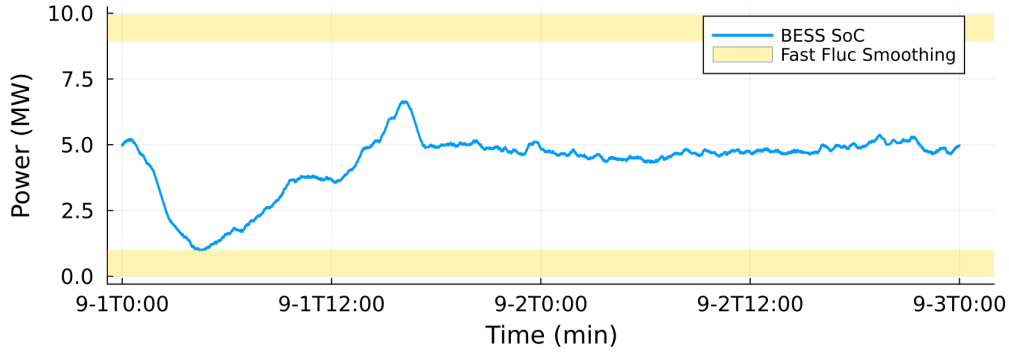


Figure 22: BESS SoC trajectory in the voltage-sag-only case. The yellow regions represent the upper and lower 10% of the total BESS energy capacity reserved to accommodate energy exchange for fast fluctuation smoothing.

3:10 am to 3:50 am is enlarged to emphasize the system response to the fault. During the event, the PCC voltage decreases to 0.6 p.u., whereas the root-node voltage remains above 0.8 p.u., thereby satisfying the maximum allowable deviation of 0.2 p.u.

The BESS SoC trajectory is shown in Fig. 22. At the beginning of the simulation, the BESS discharges to enforce the 30 MW upper limit on the total load. The SoC then remains close to 50% for the rest of the simulation horizon. Since no demand response events are included, the BESS undergoes no further substantial charging or discharging in this case.

Usually, the minimum required ride through time before transfer to backup power is sub-second to several seconds, which is much shorter than the assumed 20 minutes in our simulation. Since the time resolution in our problem is one minute, we run another test reducing the voltage time to one step. The voltages at the PCC and the DC root node is shown in Fig. 23. In addition, the resulting BESS size is 9.96 MWh and 78.45 MW, which is almost the same when voltage sag lasts for 20 minutes.

4.3.4 Integrating demand response and voltage sag scenarios

When both demand response and voltage sag scenarios are considered, the BESS must provide load flexibility as well as root-node voltage support. Consequently, both its rated power and energy capacity must increase. The resulting optimal BESS size is 80.82 MW/18.69 MWh,

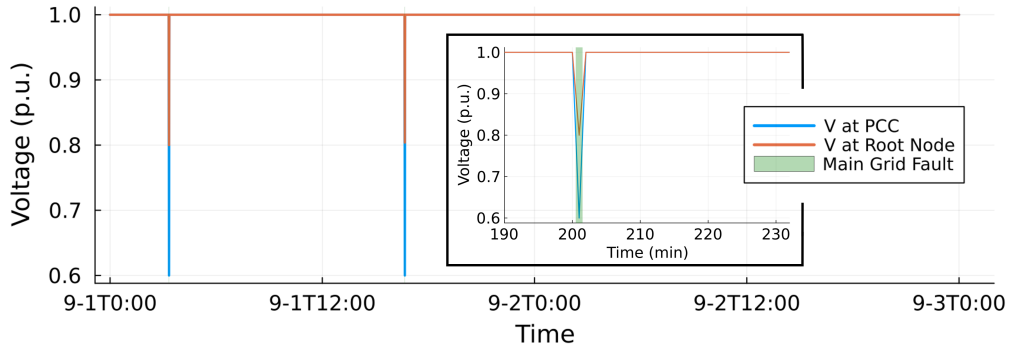


Figure 23: PCC and data center root-node voltages in the voltage-sag-only case with voltage sag last for one time step. The interval from Sep. 1 at 3:10 am to 3:50 am is enlarged.

with a total cost of $15.4 million. The resulting net-load profile, PCC and root-node voltage trajectories, and BESS SoC evolution are similar to those obtained in the corresponding single-scenario cases. The figures are omitted for brevity and to avoid redundancy.

The 80.82 MW/18.69 MWh BESS is further evaluated under extreme scenarios in which the intervals between DR request events are shortened. Specifically, after the first DR event at 7:00, the second event is assumed to occur after 1, 2, …, or 5 hours. The design problem assumes a six-hour recharge period before the next DR request. Thus, a 20% peak-load reduction is guaranteed only for the first DR event. During the second event, the BESS reduces the peak load to the greatest extent possible given its remaining SoC and the limited time available for recharging.

The resulting DC net-load and BESS SoC trajectories are shown in Fig. 24. When the available recharge time is less than 3 hours, the peak-load reduction achieved during the second DR event decreases below 20%. In every scenario, the BESS fully depletes the energy reserved for DR once the second event is requested. Furthermore, the original DC load reaches the maximum allowable value of 30 MW between approximately 10:00 and 12:00. The BESS therefore cannot recharge during this interval, which further restricts its ability to support the second DR event. These results demonstrate that guaranteeing load flexibility is particularly difficult when DR events occur frequently during periods of high DC load.

The sensitivity of the BESS size to DR and voltage sag events is evaluated by solving the design problem for 15 randomly generated scenario configurations. Each configuration contains two scenarios, with two events in each scenario. Thus, each configuration includes four randomly generated events—two DR events and two VS events—distributed across the two scenarios. All DR events are assumed to last 1 hour, whereas each VS event lasts 20 minutes.

The randomly generated scenarios, including the event types and occurrence times for each configuration, together with the corresponding design results, are presented in Fig. 25. For most configurations, the optimal BESS size is approximately 80 MW/18.7 MWh, indicating that this size is sufficient to support the DC under most conditions. Configurations 6 and 12 are the only exceptions. In both cases, one DR event occurs at the beginning of Sep. 1, when the original load is high. The BESS must therefore discharge additional energy to achieve the required peak-load reduction while simultaneously providing demand response support.

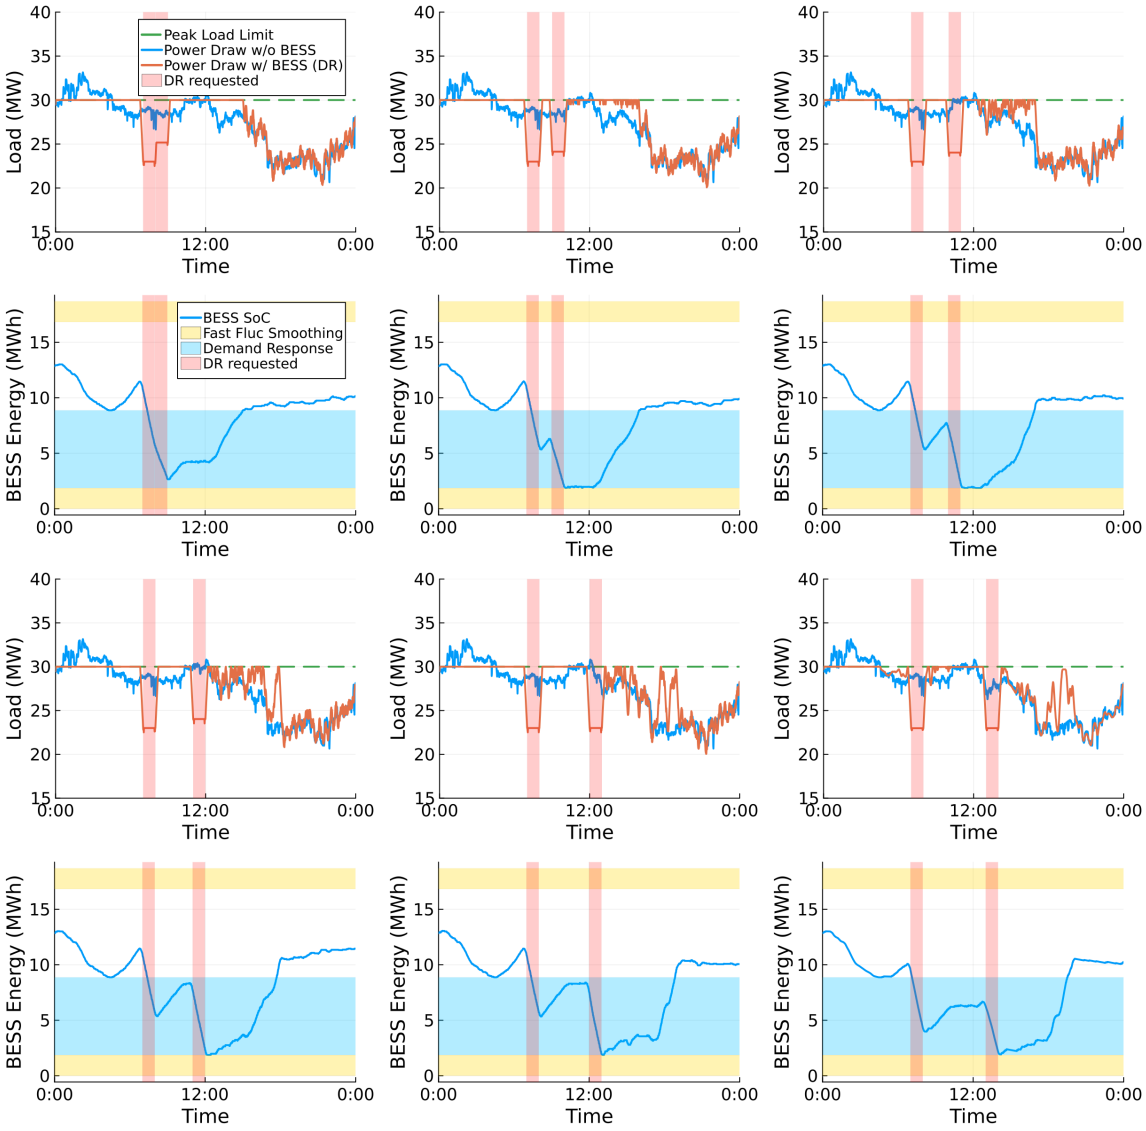


Figure 24: Data center net-load and BESS SoC trajectories under extreme demand response scenarios. During the second DR event, the peak-load reduction is less than 20% because the BESS has insufficient time to fully recharge.

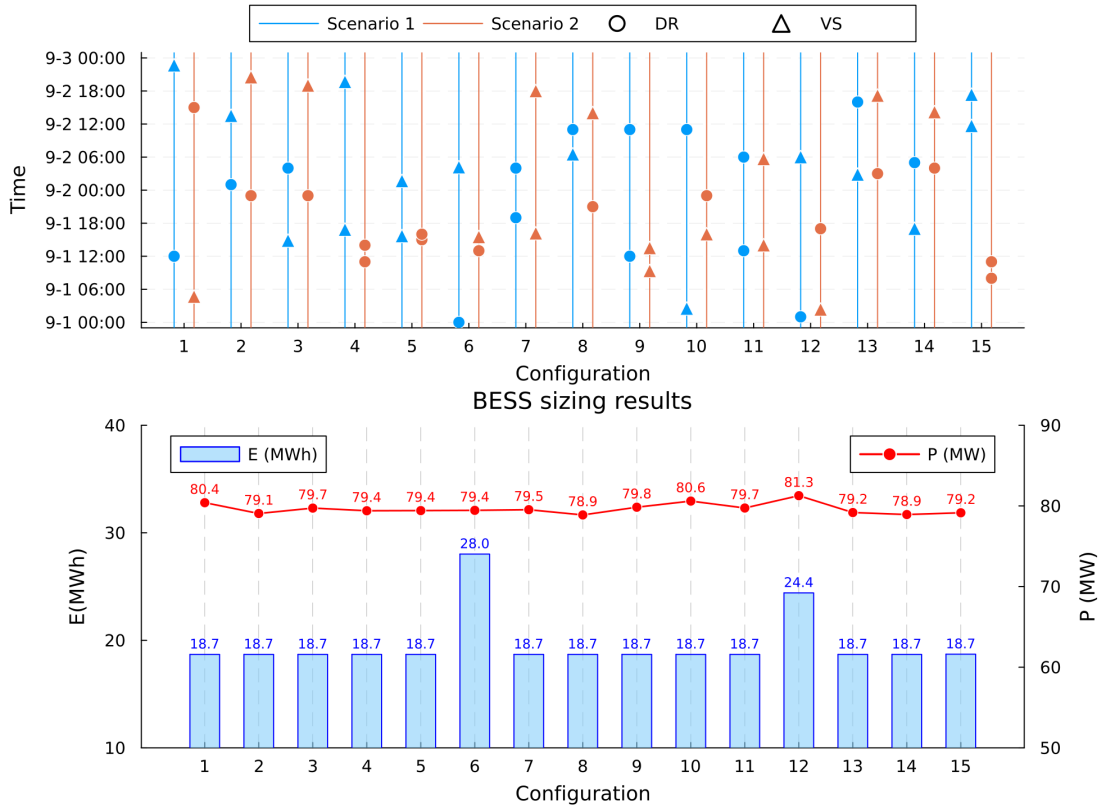


Figure 25: Randomly generated scenario configurations and corresponding BESS sizes from the design results. The upper subplot illustrates the occurrence times and types of the two events in each scenario for all configurations. The lower subplot shows the BESS rated-power and energy-capacity results for each configuration.

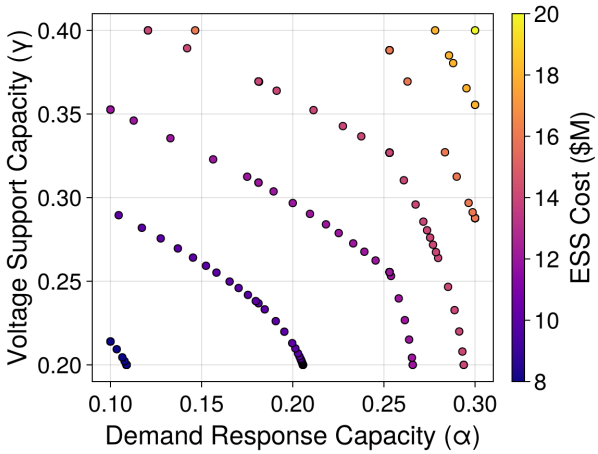


Figure 26: Pareto front for the multiple objectives. The x-axis represents demand response capacity, the y-axis represents voltage support capacity, and the color of the dots indicates the total BESS cost.

4.3.5 Multi-objective optimization with both demand response and voltage sag scenarios

In the final case, the demand response capacity α and allowable PCC voltage drop γ are treated as decision variables. An upper bound on the total BESS cost is included to facilitate the generation of the Pareto front. By varying this cost limit and the weighting factors w^{DR} and w^{VS} in the scalarized objective function (28), the Pareto front is obtained, as shown in Fig. 26.

Under a limited BESS budget, both the achievable demand response capacity and the voltage support capability are restricted. As the BESS cost increases, the DC can reduce more net load during demand response events and withstand deeper PCC voltage sags during main-grid fault contingencies. When the capability levels become high ($\alpha > 0.25$ and $\gamma > 0.35$), further improvements lead to a steep increase in cost. The resulting Pareto front presents a range of trade-off solutions from which designers can select the option that best satisfies their requirements.

5.0 Conclusion

Multi-criteria, multi-scenario infrastructure planning requires nuanced human evaluation rather than a single universal algorithmic ranking. The pyMOODS framework addresses this by bridging multi-objective optimization outputs with decision-making workflows through reusable data contracts, backend analytics, and linked visual views. By combining preference-based rankings with parameter-free generalizer-specializer analysis, the system exposes the structural trade-offs and balanced baselines of candidate designs, ensuring the final judgment remains with the decision-maker. While formative case studies demonstrate the tool's technical feasibility and domain relevance, future work will focus on quantifying data reduction, testing sensitivity to candidate-set changes, and conducting controlled user studies to validate usability and cognitive load improvements. Additionally, future iterations will address visual scalability limitations—such as overplotting in parallel coordinates and dense overlapping in Euler diagrams—by introducing interactive tools for rearranging axes and comparing orderings.

Complementing this decision-support layer, CAMEO accelerates the development and validation of multi-objective optimization formulations for power and energy systems. By automating the translation of domain expertise into scalable workflows, CAMEO efficiently generates the computational studies required for large-scale planning. When integrated with pyMOODS, CAMEO closes the analytical loop, empowering interdisciplinary teams to make informed, data-driven decisions regarding optimal designs and hyper-parameters.

As part of the project the work also established a co-design capabilities to develop battery energy storage (BESS) optimization formulation to provides practical guidance for data center planners to support the power grids by providing flexibility in operations through onsite storage sizing. The scenario-based model demonstrates that demand response is the primary driver of BESS energy-capacity requirements, whereas providing reactive power for voltage sags necessitates a considerably higher power rating storage. Providing these grid-support services simultaneously forces the BESS to satisfy combined technical constraints, resulting in higher overall storage costs. Furthermore, stress tests reveal that frequent demand response requests during periods of high data center load severely limit available recharge time, reducing the guaranteed load reduction. Future enhancements to this data center framework will investigate more detailed event representations and system-level dynamics to more accurately evaluate fluctuation-smoothing and voltage ride-through performance.

6.0 References

Alsallakh, Bilal, Luana Micallef, Wolfgang Aigner, Helwig Hauser, Silvia Miksch, and Peter Rodgers. 2014. Visualizing sets and set-typed data: state-of-the-art and future challenges. In *Eurographics conference on visualization (eurovis).*

Arendt, Dustin L., Milan Jain, Rounak Meyur, Wei Wang, Jennifer N. Pham, Palak Mattoo, Thiagarajan Ramachandran, Sumit Purohit, Veronica A. Adetola, and Himanshu Sharma. 2027. Multi-criteria multi-scenario decision-making: no solution is perfect, but some are special. In *Proceedings of the hawaii international conference on system sciences (hicss).*

Atkins, Jonathan E., Erik G. Boman, and Bruce Hendrickson. 1998. A spectral algorithm for seriation and the consecutive ones problem. *SIAM Journal on Computing* 28 (1): 297–310. https://doi.org/10.1137/S0097539795285771.

Azizi, Arian, Samaneh Morovati, Amin Zamani, Joshua Piruzza, David Guo, Zhuoning Liu, and Pasi Taimela. 2026. Strengthening data center operations using grid-forming battery energy storage as a line-interactive uninterruptible power supply. *International Journal of Electrical Power & Energy Systems* 175:111638.

Chambers, John M. 2018. *Graphical methods for data analysis.* Chapman & Hall/CRC.

Chen, Min, Ciwei Gao, Meng Song, Songsong Chen, Dezhi Li, and Qiang Liu. 2020. Internet data centers participating in demand response: a comprehensive review. *Renewable and Sustainable Energy Reviews* 117:109466.

Cho, Seolhee, Can Li, and Ignacio E. Grossmann. 2022. Recent advances and challenges in optimization models for expansion planning of power systems and reliability optimization. *Computers & Chemical Engineering* 165:107924. https://doi.org/10.1016/j.compchemeng.2022.107924.

Cinelli, Marco, Miłosz Kadziński, Michael Gonzalez, and Roman Słowiński. 2020. How to support the application of multiple criteria decision analysis? let us start with a comprehensive taxonomy. *Omega* 96:102261. https://doi.org/10.1016/j.omega.2020.102261.

Cupelli, Lisette, Nikhil Barve, and Antonello Monti. 2017. Optimal sizing of data center battery energy storage system for provision of frequency containment reserve. In *Iecon 2017-43rd annual conference of the ieee industrial electronics society,* 7185–7190. IEEE.

Cupelli, Lisette, Thomas Schütz, Pooyan Jahangiri, Marcus Fuchs, Antonello Monti, and Dirk Müller. 2018. Data center control strategy for participation in demand response programs. *IEEE Transactions on Industrial Informatics* 14 (11): 5087–5099.

Dabbagh, Mehiar, Bechir Hamdaoui, Ammar Rayes, and Mohsen Guizani. 2017. Shaving data center power demand peaks through energy storage and workload shifting control. *IEEE Transactions on Cloud Computing* 7 (4): 1095–1108.

Dvorkin, Yury, Ricardo Fernández-Blanco, Yishen Wang, Bolun Xu, Daniel S. Kirschen, Hrvoje Pandžić, Jean-Paul Watson, and Cesar A. Silva-Monroy. 2018. Co-planning of investments in transmission and merchant energy storage. *IEEE Transactions on Power Systems* 33 (1): 245–256. https://doi.org/10.1109/TPWRS.2017.2705187.

Filipič, Bogdan, and Tea Tušar. 2018. A taxonomy of methods for visualizing pareto front approximations. In *Proceedings of the genetic and evolutionary computation conference,* 649–656. GECCO '18: Genetic and Evolutionary Computation Conference. Kyoto Japan: ACM, July 2, 2018. ISBN: 978-1-4503-5618-3, accessed July 16, 2024. https://doi.org/10.1145/3205455.3205607. https://dl.acm.org/doi/10.1145/3205455.3205607.

Gratzl, Samuel, Alexander Lex, Nils Gehlenborg, Hanspeter Pfister, and Marc Streit. 2013. LineUp: Visual Analysis of Multi-Attribute Rankings [in en]. *IEEE Transactions on Visualization and Computer Graphics* 19, no. 12 (December): 2277–2286. ISSN: 1077-2626, accessed July 16, 2024. https://doi.org/10.1109/TVCG.2013.173. http://ieeexplore.ieee.org/document/6634146/.

Hoffman, Patrick, Georges Grinstein, and David Pinkney. 1999. Dimensional anchors: a graphic primitive for multidimensional multivariate information visualizations. In *Proceedings of the 1999 workshop on new paradigms in information visualization and manipulation in conjunction with the eighth acm international conference on information and knowledge management,* 9–16.

Inc., Plotly Technologies. 2015. *Collaborative data science.* Montreal, QC. https://plot.ly.

Inselberg, Alfred. 2008. *Parallel coordinates: visual multidimensional geometry and its applications.* Springer.

Jain, Milan, Soumya Kundu, Arnab Bhattacharya, Sen Huang, Vikas Chandan, Nikitha Radhakrishnan, Veronica Adetola, and Draguna Vrabie. 2021. Occupancy-driven stochastic decision framework for ranking commercial building loads. In *2021 american control conference (acc),* 4171–4177. IEEE.

Kohonen, Teuvo. 2012. *Self-organizing maps.* Vol. 30. Springer Science & Business Media.

Korte, Bernhard, and Jens Vygen. 2008. *Combinatorial optimization: theory and algorithms.* Springer.

Krishnan, Venkat, Jonathan Ho, Benjamin F. Hobbs, Andrew L. Liu, James D. McCalley, Mohammad Shahidehpour, and Qipeng P. Zheng. 2016. Co-optimization of electricity transmission and generation resources for planning and policy analysis: review of concepts and modeling approaches. *Energy Systems* 7 (2): 297–332. https://doi.org/10.1007/s12667-015-0158-4.

Kumar, Abhishek, Bikash Sah, Arvind R. Singh, Yan Deng, Xiangning He, Praveen Kumar, and R. C. Bansal. 2017. A review of multi criteria decision making (MCDM) towards sustainable renewable energy development. *Renewable and Sustainable Energy Reviews* 69:596–609. https://doi.org/10.1016/j.rser.2016.11.191.

Liu, Shaohui, Sungho Shin, and Deepjyoti Deka. 2026. Watts vs. Bytes: Turning data centers into grid assets via storage—compute co-optimization. Submitted to arXiv, *arXiv preprint arXiv:2605.16190,* https://arxiv.org/abs/2605.16190.

Mamun, A, I Narayanan, D Wang, A Sivasubramaniam, and HK Fathy. 2016. Multi-objective optimization of demand response in a datacenter with lithium-ion battery storage. *Journal of Energy Storage* 7:258–269.

Mattoo, Palak, Jennifer N. Pham, Milan Jain, Dustin Arendt, Phoenix Avila, Thiagarajan Ramachandran, Veronica Adetola, Ji Young Yun, and John Wenskovitch. 2026. Pymoods: an open-source framework for multi-objective decision support in large-scale infrastructure planning. *SoftwareX.*

McInnes, Leland, John Healy, and James Melville. 2020. *UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction,* arXiv:1802.03426, September. Accessed November 6, 2025. https://doi.org/10.48550/arXiv.1802.03426.

Meyur, Rounak, Sam Donald, Tonya Martin, Thiagarajan Ramachandran, and Sumit Purohit. 2025. Cameo: a co-design architecture for multi-objective energy system optimization. In *Proceedings of the sc'25 workshops of the international conference for high performance computing, networking, storage and analysis,* 2201–2212.

Miettinen, Kaisa. 1999. *Nonlinear multiobjective optimization.* Springer.

MUI. 2014. *Mui: react component library.* https://mui.com/.

Munzner, Tamara. 2025. Visualization analysis and design. In *Proceedings of the special interest group on computer graphics and interactive techniques conference courses,* 1–2.

Osika, Zuzanna, Jazmin Zatarain Salazar, Diederik M. Roijers, Frans A. Oliehoek, and Pradeep K. Murukannaiah. 2023. What lies beyond the Pareto front? a survey on decision-support methods for multi-objective optimization. In *Proceedings of the thirty-second international joint conference on artificial intelligence,* 6741–6749. https://doi.org/10.24963/ijcai.2023/755.

Pham, Jennifer, Milan Jain, and Palak Mattoo. 2026. *pyMOODS: Python for multi-objective optimization and decision support.* Zenodo software archive, version 0.0.4. https://doi.org/10.5281/zenodo.22882257. https://doi.org/10.5281/zenodo.22882257.

Samsi, Siddharth, Matthew L Weiss, David Bestor, Baolin Li, Michael Jones, Albert Reuther, Daniel Edelman, William Arcand, Chansup Byun, John Holodnack, et al. 2021. The MIT supercloud dataset. In *2021 ieee high performance extreme computing conference (hpec),* 1–8. IEEE.

Sharma, Himanshu, Wei Wang, Bowen Huang, Buxin She, and Thiagarajan Ramachandaran. 2025. Control co-design under uncertainty for offshore wind farms: optimizing grid integration, energy storage, and market participation, April 11, 2025. https://doi.org/10.48550/arXiv.2504.08555. arXiv: 2504.08555. https://arxiv.org/abs/2504.08555.

Shehabi, Arman, Andrew Newkirk, Sarah Smith, A. Hubbard, N. Lei, M. Siddik, et al. 2024. *2024 United States data center energy usage report.* Technical report LBNL-2001637. Lawrence Berkeley National Laboratory. https://doi.org/10.71468/P1WC7Q. https://escholarship.org/uc/item/32d6m0d1.

Shi, Yuanyuan, Bolun Xu, Baosen Zhang, and Di Wang. 2016. Leveraging energy storage to optimize data center electricity cost in emerging power markets. In *Proceedings of the seventh international conference on future energy systems,* 1–13.

Smedberg, Henrik, and Sunith Bandaru. 2023. Interactive knowledge discovery and knowledge visualization for decision support in multi-objective optimization. *European Journal of Operational Research* 306 (3): 1311–1329.

Snider, Shane. 2026. Why AI data center projects face years of delays after approval. Accessed May 12, 2026. https://www.datacenterknowledge.com/energy-power-supply/why-ai-data-center-projects-face-years-of-delays-after-approval.

Tadelis, Steven. 2013. *Game theory: an introduction.* Princeton university press.

Takci, Mehmet Türker, Meysam Qadrdan, Jon Summers, and Jonas Gustafsson. 2025. Data centres as a source of flexibility for power systems. *Energy Reports* 13:3661–3671.

Thompson, Christopher C, PE Konstantinos Oikonomou, Amir H Etemadi, and Volker J Sorger. 2016. Optimization of data center battery storage investments for microgrid cost savings, emissions reduction, and reliability enhancement. *IEEE Transactions on Industry Applications* 52 (3): 2053–2060.

Tian, Junfeng, Binyan Wang, and Shijun Wang. 2023. The paradigms and main content of land use conflict research. *Resources Science* 45 (3): 465–479.

Triantaphyllou, Evangelos, and Khalid Baig. 2005. The impact of aggregating benefit and cost criteria in four mcda methods. *IEEE Transactions on Engineering Management* 52 (2): 213–226.

Tušar, Tea, and Bogdan Filipič. 2015. Visualization of pareto front approximations in evolutionary multiobjective optimization: a critical review and the prosection method. *IEEE Transactions on Evolutionary Computation* 19, no. 2 (April): 225–245. ISSN: 1089-778X, 1089-778X, 1941-0026, accessed November 6, 2024. https://doi.org/10.1109/TEVC.2014.2313407. https://ieeexplore.ieee.org/document/6777535/.

U.S. Environmental Protection Agency. 2026. *Clean air act resources for data centers.* https://www.epa.gov/stationary-sources-air-pollution/clean-air-act-resources-data-centers. Accessed September 22, 2026.

Wall, Emily, Subhajit Das, Ravish Chawla, Bharath Kalidindi, Eli T Brown, and Alex Endert. 2017. Podium: ranking data using mixed-initiative visual analytics. *IEEE Transactions on Visualization and Computer Graphics* 24 (1): 288–297.

Wayne, Wang, Bhatnagar Archit, Miao Tongyuan, Kalamkar Saniya, Cui Wenqi, Incer Inigo, and Chen Ang. 2026. Synthesizing voltage ride-through controllers for data centers. *arXiv preprint arXiv:2608.07289.*

Weng, Di, Ran Chen, Zikun Deng, Feiran Wu, Jingmin Chen, and Yingcai Wu. 2018. SRVis: towards better spatial integration in ranking visualization. *IEEE Transactions on Visualization and Computer Graphics* 25 (1): 459–469.

Xie, Yiheng, Wenqi Cui, and Adam Wierman. 2025. Enhancing data center low-voltage ride-through. *arXiv preprint arXiv:2510.03867.*

Pacific Northwest National Laboratory

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7675)

www.pnnl.gov