Break-Even Economics of Nuclear Microreactor Fleets in Southeast Michigan

Marco Sarmiento[a,b*], Majdi I. Radaideh[b,c*]

[a] *Energy Engineering Program, University of the Philippines Diliman, Diliman, Quezon City 1101, Philippines*

[b] *Department of Nuclear Engineering and Radiological Sciences, University of Michigan, Ann Arbor, Michigan 48109, USA*

[c] *Department of Computer Science and Engineering, University of Michigan, Ann Arbor, Michigan 48109, USA*

Abstract

Nuclear microreactors are proposed for remote communities, industrial facilities, and data centers, but their economics remain commercially unproven. Without an operating fleet, first-of-a-kind (FOAK) to (n)th-of-a-kind (NOAK) cost progression relies on bottom-up estimates, scaling assumptions, and assumed learning, leaving overnight capital cost (OCC) among the least-established inputs in direct levelized cost of electricity (LCOE) comparisons. Rather than prescribing OCC, this work asks: what OCC must a microreactor fleet reach for its generation cost to equal the fossil generation it replaces? Fleets of 13-, 10-, and 5-MW_e Holos-Quad, Aalo-1, and eVinci units are dispatched against measured 2025 coal, natural-gas combined-cycle, and gas-peaking generation from DTE Electric in Southeast Michigan. Because capital cost does not affect hourly merit order, replacement cost is linear in a common capital multiplier (m), enabling break-even (m^*) from one dispatch. Ten metaheuristics algorithms screen mixed fleets over ten seeds using 1,000 evaluations per run. Selected economic candidates contain only Aalo-1 units: 185, 111, 35, and 124 for coal, combined-cycle gas, gas peaking, and combined cases. Using midpoint fossil costs from Lazard's 2026 LCOE+, corresponding NOAK break-even OCCs are 6,484, 1,063, 2,414, and 8,918 \$/kW. Replacement fraction alone does not determine the threshold: combined-cycle replaces 96.73% of target energy yet supports the lowest positive multiplier, whereas the combined fleet replaces 29.32% and supports the highest, demonstrating the importance of utilization and displaced-source cost. Candidate-specific reactor references are then evaluated using nonlinear model predictive control. All fleets satisfy hourly-linear tracking and remain below 1% under a separate Midcontinent Independent System Operator (MISO)-informed 5-min sensitivity, although the coal case reaches the 0.5 degrees/s drum-rate limit. These OCCs are economic break-even targets, not future-cost predictions; coal and combined-case Aalo-1 thresholds fall within generic microreactor cost ranges reported by the Nuclear Energy Institute.

Keywords: Fleet Dispatch; Levelized Cost of Electricity; Capital Cost; Nuclear Microreactors; Optimization; Break-even Analysis

1. Introduction

Electricity demand in the United States has increased sharply in recent years, driven in part by the expansion of data centers and new industrial loads [1]. At the same time, coal units are being retired or considered for replacement, including units in the DTE Electric system [2], which is the largest electric utility in Michigan, United States. Any replacement resource must supply electricity when needed while remaining economically competitive and supporting the transition away from fossil generation. Nuclear energy is one option because it provides high-capacity and low-carbon generation. However, conventional nuclear plants have large unit sizes, long construction schedules, and high initial capital requirements. These limitations have supported the development of small modular reactors (SMRs) [3] and, more recently, nuclear microreactors [4].

Microreactors are generally designed for smaller power output, factory fabrication, transportability, and reduced on-site construction. These features make them relevant to remote systems, industrial facilities, military installations, data centers, and hybrid energy systems. Their main economic limitation is capital cost. Bottom-up estimates remain high on a per-kW basis, and no commercial fleet has yet produced the learning needed to confirm the nth-of-a-kind (NOAK) values used in long-term studies [5, 6]. This creates a problem for the normal

*Corresponding Author: M. M. Sarmiento (mmsarmiento3@up.edu.ph), M. I. Radaideh (radaideh@umich.edu)

techno-economic approach. When viability is evaluated at an assumed overnight capital cost (OCC), the conclusion depends strongly on the input that is also one of the least certain.

Existing microreactor economic studies generally assume reactor cost as an input. Market studies compare a specified reactor cost against diesel or other high-cost alternatives and then identify where deployment may be competitive [3, 4]. Bottom-up cost work instead develops the capital and operating inputs needed for those comparisons [7]. The Microreactor Optimization Using Simulation and Economics (MOUSE) framework provides the component-level first-of-a-kind (FOAK) and NOAK estimates used in the present analysis [6, 8]. In each case, the reactor cost is treated as an input, and the resulting conclusions depend on that assumed value.

A separate body of work examines fossil replacement and coal-to-nuclear transitions. These studies screen sites, identify infrastructure that may be reused, and evaluate the cost and community effects of replacing retiring coal generation with nuclear capacity [9]. The reactor technology and its cost are normally fixed before the site is evaluated. The question asked here is different. The fossil-generation profile is fixed from measured operation, while the fleet size and the capital cost it can support are calculated as results.

Reactor flexibility is another part of the problem. Economic dispatch can assign a power trajectory to each unit, but it does not show that the reactor model can follow that trajectory under actuator limits. Microreactor load-following studies have developed sliding-mode, proportional–integral–derivative (PID), and nonlinear model predictive controllers (NMPC) for control-drum systems [10]. Deep reinforcement learning has also been applied to the same general control problem [11]. These studies test reactor behavior but do not determine the fleet-level cost threshold. Economic studies tend to assume the operating behavior instead. The two analyses are therefore connected here, but they are not nested inside the same optimization loop.

The analysis uses DTE Electric as the test system. DTE Electric serves approximately 2.3 million customers in southeastern Michigan and operates 12.4 GW of generating capacity, of which 8.1 GW is fossil fired [12]. The utility has committed to ending coal generation in 2032, and its Monroe plant alone supplied more than 30% of its 2025 power plant generation. It also reports that it is monitoring modular nuclear reactors among the emerging technologies that may support future clean generation. Hourly operation of every affected unit is publicly reported, which makes the fossil profile a measured input rather than a modeled one.

This work asks the inverse economic question. Instead of assuming an OCC and asking whether a microreactor fleet is competitive, it asks what OCC would make the fleet competitive with the fossil generation it replaces. The question can be answered directly because the units are dispatched in order of marginal generation cost, defined as fuel plus variable operation and maintenance. Capital cost does not change that order. Scaling the capital cost of every unit by a common multiplier therefore leaves the merit order, hourly generation, capacity factors, and displaced fossil energy unchanged. Only the annualized capital term changes, so the break-even multiplier follows from one dispatch.

The break-even OCC is a fleet result rather than a reactor value by itself. A small fleet replaces only the highest-cost fossil generation, but its early units remain close to FOAK cost. A larger fleet experiences greater cost reductions through learning, but it reaches lower-cost fossil generation and may operate at a lower capacity factor. These effects work against each other. The fleet with the highest supported capital cost therefore lies between a very small and a fully replacing fleet, and neither fleet size nor replacement percentage is fixed in advance.

There is also a non-capital floor. Fixed operation and maintenance, fuel, variable operation and maintenance, and the cycling charge associated with hourly startup and shutdown events are paid even when OCC is reduced to zero. When these costs are already above the displaced fossil-generation cost, no positive OCC can achieve break-even. This is not a failure of the inverse method. It indicates that capital cost is not the only cost component that would need to change. Consequently, the main contributions of this work are as follows:

- An inverse economic formulation is developed to solve for the required OCC rather than assume it.
- The levelized replacement cost is shown to be exactly linear in a common capital multiplier, so the break-even multiplier of an hourly dispatch follows from a single evaluation.
- Mixed fleets are screened with several optimization algorithms, while each pure design is also checked by complete integer enumeration.
- Fleet size, capacity factor, and fossil-generation replacement are reported as results rather than as imposed targets.
- The economically selected fleet is passed to a candidate-specific worst-day reactor-level screen that preserves the unit-level hourly dispatch and tests online load following under explicit control-drum limits.

After we review some related studies in Section 2, the remainder of the paper follows the two analysis levels. Section 3 defines the candidate reactor designs and presents the annual fleet-level methodology, including the load profiles, cost progression, merit-order dispatch, cycling charge, inverse cost formulation, and optimization screening. Section 4 gives the reduced-order reactor model, nonlinear model predictive control formulation, worst-day reference construction, and controller tuning. The economic and reactor-level results are then presented separately in Section 5 before their combined implications and limitations are discussed in Section 6. Section 7 concludes and sets out possible extensions.

2. Related Work

The present work draws from several areas outside microreactor economics and reactor control. These include power-system dispatch and planning, fossil-generation displacement, generation-portfolio optimization, and break-even cost analysis. This section reviews these areas and identifies how they relate to the inverse fleet-cost problem considered in this study.

2.1. Power-System Dispatch and Planning

Power-system optimization includes both planning and operational problems. Capacity-expansion models determine future investments in generation, storage, and transmission resources [13], while economic-dispatch models allocate generation among available units according to operating costs and system constraints. Unit-commitment formulations further incorporate startup and shutdown costs, minimum up- and down-times, ramp-rate limits, reserve requirements, and binary operating decisions [14]. Uncertainty in demand, renewable generation, fuel prices, and unit availability has also motivated stochastic unit-commitment approaches [15].

Because these optimization frameworks often rely on repeated evaluation of computationally intensive engineering models, considerable effort has been devoted to reducing their computational cost. Recent advances in reduced-order modeling and digital twins have improved the efficiency of engineering analyses. Variational digital twins [16], physics-informed digital twins with active learning [17, 18], and reduced-order reconstruction techniques based on simplified matching pursuits [19] enable efficient representation of complex physical systems for optimization and uncertainty quantification.

The present work adopts a different level of abstraction because the objective is to evaluate thousands of candidate microreactor fleets rather than reconstruct detailed power-system operation. The measured DTE Electric fossil-generation profile is therefore treated as the replacement target, and candidate microreactors are dispatched according to marginal generation cost. Startup and shutdown costs, binary commitment decisions, reserve constraints, and other unit-commitment features are intentionally omitted, providing the hourly generation and utilization required for fleet-level inverse economic analysis while maintaining computational efficiency.

2.2. Fossil-Generation Displacement and Replacement

A substantial body of work has examined the displacement of fossil generation by wind, solar, energy storage, and other low-carbon technologies. These studies generally compare the cost of the entering resource with the operating or levelized cost of the conventional generation it displaces, showing that the economic benefit depends on both the quantity and type of fossil generation replaced.

Merit-order models provide one approach for quantifying this effect. Jain and Shrimali combined measured utility schedules with a merit-order model to estimate the economic impact of renewable generation, demonstrating that the value of additional generation declines as progressively lower-cost thermal resources are displaced [20]. Similarly, the Coal Cost Crossover analysis compared continued coal operation with replacement by nearby renewable resources and energy storage across the United States [21]. Recent studies have also employed surrogate-assisted methods to evaluate integrated energy systems under uncertain operating conditions, emphasizing that economic value depends strongly on system operation rather than aggregate energy production alone [22].

The present work adopts the same displaced-generation perspective but considers fleets of nuclear microreactors. Coal, combined-cycle gas, and gas-peaking generation are treated as separate replacement targets because they differ in hourly operating profiles and generation costs. Consequently, the supported reactor capital cost depends on both the source and timing of the displaced generation. In the combined replacement case, microreactor generation first displaces the highest-cost fossil generation before progressively replacing lower-cost sources. As a result, replacing more energy does not necessarily increase economic value, and replacement fraction alone is insufficient to determine the supported capital cost.

2.3. Optimization of Generation Portfolios

Optimization has been widely applied to generation expansion, hybrid energy systems, reactor design, and operational planning using mathematical programming and metaheuristic algorithms. While mathematical programming is effective for structured optimization problems, evolutionary and swarm-based methods are well suited for nonlinear, mixed-variable, and simulation-driven applications.

In nuclear engineering, population-based optimization has been successfully applied to boiling water reactor fuel bundle optimization [23], physics-informed reinforcement learning for nuclear assembly design [24], multi-objective optimization of microreactor reactivity control systems [25], hybrid energy-system sizing [26], and multiphysics reactor optimization [27]. High-fidelity neutronic and thermal-hydraulic analyses have also been coupled to determine burnup-dependent operating conditions for heat-pipe microreactors [28, 29]. These studies demonstrate the flexibility of metaheuristic optimization for computationally intensive nuclear engineering problems.

Meaningful comparisons between optimization algorithms require consistent benchmarking because performance depends on factors such as evaluation budget, stopping criteria, random initialization, and constraint handling [30]. Accordingly, the present work evaluates multiple optimization algorithms under identical computational budgets and repeated random seeds. The optimization variables consist only of the integer fleet sizes of the three candidate microreactor designs, while reactor ratings, learning assumptions, cost parameters, and measured fossil-generation profiles remain fixed.

Complete enumeration is additionally performed for each pure-design fleet to verify the optimizer results. Unlike previous studies that primarily optimize reactor design, reactor control, or multiphysics performance, the proposed framework optimizes discrete microreactor fleet composition while directly solving for the maximum overnight capital cost supported by displaced fossil generation through inverse economic analysis.

2.4. Break-Even and Inverse Cost Analysis

Most techno-economic studies evaluate candidate energy technologies by assuming capital and operating costs before calculating system cost or economic competitiveness. Consequently, the conclusions depend directly on technology costs that often remain uncertain for emerging energy systems. A smaller body of work instead adopts an inverse perspective by determining the cost or performance required for a technology to become economically competitive.

Li et al. determined break-even capital costs for generation technologies using a capacity-expansion model, identifying the cost thresholds at which candidate technologies entered the optimal generation portfolio [31]. Similar inverse approaches have been developed for energy storage by estimating the capital cost supported by avoided operating costs and grid services [32], identifying cost and performance requirements for long-duration storage [33], and backcasting the cost targets needed for storage and firm low-carbon resources to become economically competitive [34, 35]. These studies demonstrate that inverse techno-economic analysis identifies technology targets rather than evaluating assumed designs.

Within nuclear engineering, however, economic studies have largely retained the conventional forward approach. Alonso et al. evaluated the competitiveness of small modular reactors by assuming reactor capital costs [36], while the MOUSE framework estimates FOAK and NOAK costs from component-level manufacturing models [6]. In parallel, recent advances in surrogate modeling [37], digital twins [16], physics-informed digital twins [17, 18], and reinforcement learning for autonomous microreactor control [38, 39, 40] have improved reactor analysis and operation, but they do not address the inverse economic question of determining the capital-cost thresholds required for competitive fleet deployment.

The present work applies inverse techno-economic analysis to dispatched fleets of nuclear microreactors. Measured fossil-generation profiles and source-specific generation costs are fixed, while fleet composition, annual utilization, displaced fossil generation, and supported capital cost are determined as optimization results. Because the common capital multiplier does not affect the marginal-cost dispatch order, hourly generation and displaced fossil energy remain unchanged for a fixed fleet, making the levelized replacement cost exactly linear in the multiplier. The break-even OCC can therefore be obtained from a single annual dispatch evaluation rather than repeated techno-economic simulations.

The resulting OCC represents an economic target rather than a prediction of future construction cost. Positive values identify the maximum capital cost supported under the assumed operating conditions, whereas negative values indicate that non-capital costs alone exceed the value of the displaced fossil generation. Although previous studies have addressed dispatch, fossil replacement, optimization, reactor economics, and AI-enabled reactor analysis separately, the authors did not identify prior work that combines measured source-specific fossil-generation profiles, unit-by-unit FOAK-to-NOAK cost progression, discrete mixed-fleet optimization, and a direct analytical solution for the supported reactor OCC. The proposed framework integrates these elements while separately verifying that the economically selected fleet configurations satisfy reactor-level dynamic control constraints.

3. Methodology

Figure 1 summarizes the analysis framework used in this study. A candidate fleet is first defined by the numbers of Aalo-1, Holos-Quad, and eVinci units. Each physical unit is assigned its position along the FOAK-to-NOAK cost progression, and the resulting fleet is dispatched against the selected 2025 DTE fossil-generation profile. The dispatch determines the annual microreactor generation, realized capacity factors, displaced fossil energy, and startup and shutdown events, while the corresponding capital and non-capital costs are calculated for the same fleet. These quantities are combined in the inverse break-even formulation to obtain the capital multiplier m^* supported by that candidate.

The mixed-fleet composition is searched using metaheuristic optimization methods, with each run limited to 1,000 fleet evaluations. The candidate with the largest m^* is retained as the selected economic fleet. Because the selected fleets in the present analysis contain only Aalo-1 units, a complete integer enumeration of the corresponding pure-design problem is also performed as a follow-up check of the selected unit count. The reactor-level calculation is performed only after the economic fleet has been selected. Its physical-unit hourly dispatch is retained to identify the candidate-specific most-variable day and construct the power references used for the NMPC load-following screen.

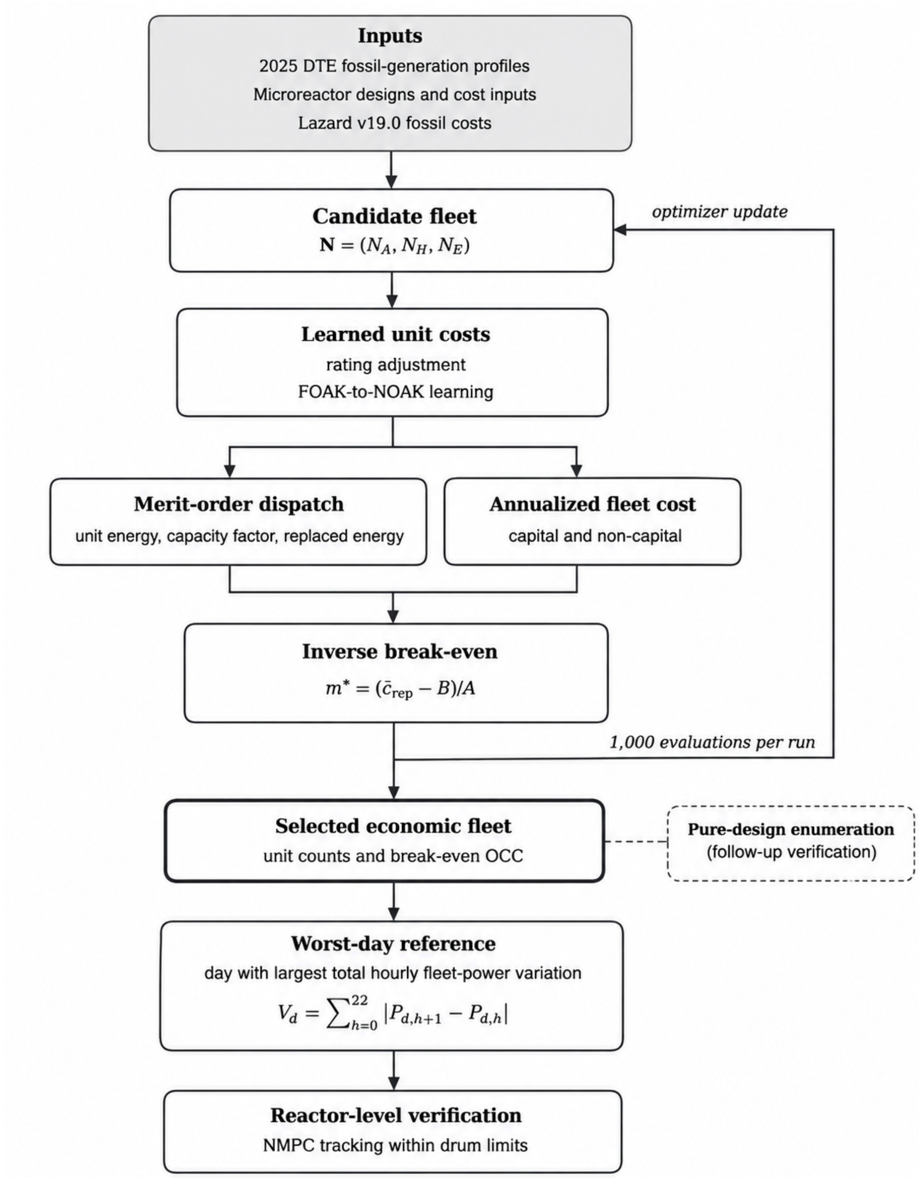

Figure 1: Analysis framework used in this study. Candidate fleets are evaluated using annual merit-order dispatch and inverse break-even cost analysis, after which the selected economic fleet is used to construct the candidate-specific worst-day reference for reactor-level verification.

3.1. DTE Electric Data

DTE Electric is used as the test system for the present analysis. Publicly available 2025 operating data are processed to construct hourly generation profiles for the fossil resources considered in the replacement cases. Coal, natural-gas combined-cycle, and gas-peaking generation are retained as separate source categories so that their different operating profiles and economic values can be represented explicitly.

3.1.1. DTE Electric generation fleet

DTE Electric owns 12,414 MW of generating capacity, all located in Michigan [12]. Table 1 gives the breakdown by source as of December 31, 2025. Fossil units account for 8,145 MW, or 65.6% of the total, and coal is the largest single category at 4,100 MW. The three fossil categories in the table are the basis for the replacement cases evaluated in this study.

Table 1: DTE Electric installed generating capacity by source in Michigan, United States, as of December 31, 2025.

Category	Source	Capacity (MW)	Share (%)
Fossil	Coal	4,100	33.0
	Natural-gas combined cycle and steam	2,092	16.9
	Natural-gas peaking	1,953	15.7
	Subtotal	8,145	65.6
Nuclear	Fermi 2	1,141	9.2
Renewable	Wind	1,597	12.9
	Solar	395	3.2
	Subtotal	1,992	16.0
Storage	Pumped hydro	1,122	9.0
	Battery	14	0.1
	Subtotal	1,136	9.2
Total		12,414	100.0

The reported values are summer net ratings rather than nameplate capacity, except for wind and solar.

3.1.2. DTE fossil-generation profiles

The economic analysis uses processed 2025 hourly generation profiles for coal, natural-gas combined cycle, and gas peaking in the DTE Electric system. The underlying operations data are taken from the U.S. Environmental Protection Agency Clean Air Markets Program Data (CAMPD) system [41]. The three profiles are shown superimposed in Fig. 2.

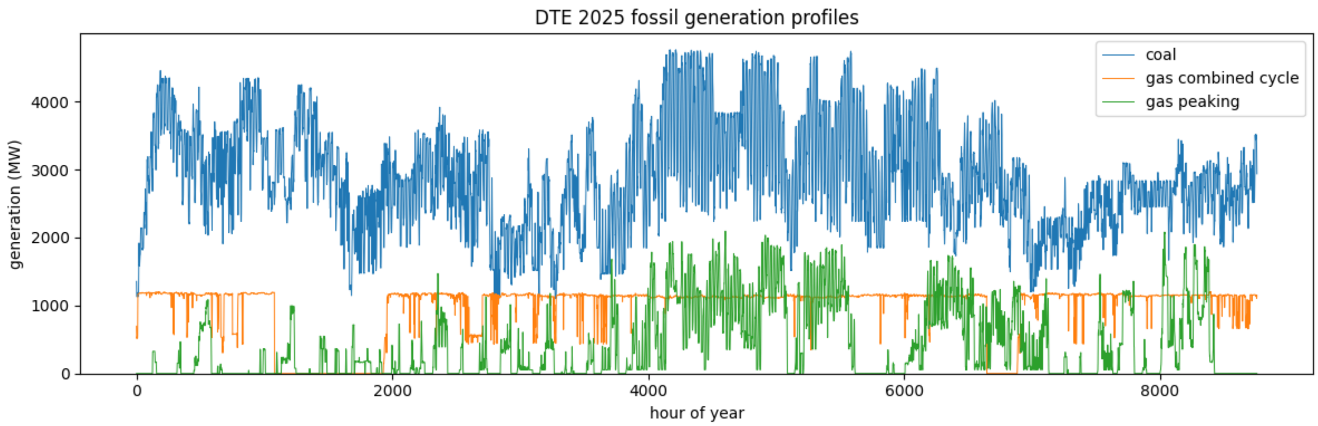

Figure 2: Processed 2025 hourly fossil-generation profiles for the DTE Electric system used as replacement targets in the economic analysis. The figure shows coal, natural-gas combined-cycle, and gas-peaking generation over the 8,760 hours of the study year using operations data from the U.S. Environmental Protection Agency Clean Air Markets Program Data (CAMPD) system [41]. The coal, combined-cycle, and gas-peaking cases use the corresponding individual profiles, while the combined coal-and-gas case uses the sum of the three source-specific profiles while retaining their separate fossil-cost layers.

Table 2 gives the annual energy and peak output of each profile. Coal supplies 67.6% of the fossil energy in the study year.

Table 2: Measured 2025 DTE fossil-generation profiles used in the economic analysis.

Source	Energy (TWh)	Peak (MW)
Coal	25.05	4,761
Natural-gas combined cycle	8.50	1,200
Gas peaking	3.49	2,094
Total fossil	37.05	7,823

The total fossil peak is lower than the sum of the individual peaks because the three sources do not reach their maxima in the same hour. Four replacement cases are evaluated: coal, combined-cycle gas, gas peaking, and the combined coal-and-gas profile. For case q, the target generation is

$$L_q(t) = \sum_{s \in \mathcal{S}_q} L_s(t), \tag{1}$$

where $\mathcal{S}_q$ contains the fossil sources included in that case. The combined case retains the source-specific layers instead of using a generic grid source.

The values in Table 3 are from Lazard's Levelized Cost of Energy+ (LCOE+) 2026 Version 19.0 report [42]. The midpoint fossil-generation costs are used in the main analysis, while low and high values are used for sensitivity. These values represent generic U.S. new-build cost benchmarks rather than plant-specific operating costs for the existing DTE fleet.

3.2. Economic Model

3.2.1. Candidate reactor designs

Three representative microreactor designs are included in this study based on publicly available vendor and technical information. Holos-Quad is a high-temperature gas-cooled microreactor (GCMR) developed by HolosGen [43, 44]. It is a modular and transportable reactor concept intended to provide compact power generation with

Table 3: Fossil-generation cost assumptions used in the inverse cost analysis.

Source	Low	Midpoint ($/MWh)	High
Coal	72.0	124.5	177.0
Natural-gas combined cycle	51.0	90.0	129.0
Gas peaking	144.0	210.0	276.0

passive safety and heat removal. Aalo-1 is a 30 MW_{th}/10 MW_e sodium-thermal reactor developed by Aalo Atomics. The current public design uses UO_2 ceramic fuel, graphite moderation, and liquid sodium coolant [45]. The Microreactor Applications Research Validation and Evaluation (MARVEL) project is an Idaho National Laboratory (INL)-led experimental microreactor program developed as a test bed for microreactor operation, control, and integration rather than as a commercial reactor design. Because detailed vendor dynamic parameters are not public, the reactor-level Aalo-1 branch is represented by a MARVEL/MOUSE liquid-metal thermal microreactor (LTMR) proxy assembled from public INL data [46, 6]. The proxy provides the reduced-order dynamic parameters used in this study and should not be interpreted as the current Aalo-1 fuel, coolant, or material specification. eVinci is a heat-pipe microreactor (HPMR) developed by Westinghouse [47]. It uses passive heat pipes to transfer thermal energy from the core without a conventional pumped primary-coolant system and is designed for factory fabrication, transportable deployment, and extended operation between refueling [48]. The reactor-level eVinci branch uses a public HPMR/generic heat-pipe microreactor (gHPMR) proxy that retains the heat-pipe thermal structure and available design characteristics without representing proprietary vendor data [29, 49].

The economic values for the GCMR, LTMR, and HPMR designs are derived from the corresponding models in the MOUSE framework [6]. These values are used as representative economic inputs and should not be interpreted as vendor quotations or exact cost estimates for Holos-Quad, Aalo-1, or eVinci. The rated electrical capacities used in the fleet model are shown in Table 4. The model rating is used for dispatch and annual energy, while the reference rating defines the published bottom-up cost basis before the adjustment of Section 3.2.2. Although this study considers three representative microreactor designs, the proposed framework is readily scalable to additional designs.

Table 4: Modeled unit ratings and corresponding MOUSE economic proxies.

Design	Economic proxy	Symbol	Reference rating (MW_e)	Model rating (MW_e)
Holos-Quad	GCMR	P_H	6.00	13
Aalo-1	LTMR	P_A	6.20	10
eVinci	HPMR	P_E	2.52	5

For a fleet with unit counts N_H, N_A, and N_E, the installed capacity is

$$P_{\text{cap}} = N_H P_H + N_A P_A + N_E P_E. \tag{2}$$

The economic model can contain many physical units of each design. The reactor-level layer preserves the hourly dispatch assigned to those physical units on the selected most-variable day. Units with identical hourly profiles are grouped only to avoid repeated controller solves, and the multiplicity of each profile is restored when the fleet response is reconstructed. This keeps the dynamic reference tied directly to the economic dispatch without requiring the reactor controller to be solved inside every annual fleet evaluation.

3.2.2. Microreactor cost basis and unit progression

The microreactor cost inputs are based on the bottom-up Idaho National Laboratory cost work [5, 6]. The associated component-level implementation is provided through the MOUSE framework [8]. The LTMR, GCMR, and HPMR values are used as representative economic proxies for Aalo-1, Holos-Quad, and eVinci, respectively. They are not vendor quotations, and the adjustment described below is intended to place the three proxies on a common fleet-model basis rather than reconstruct an exact vendor design.

The published MOUSE cost totals are reported as per-kW values using reference electrical ratings of 6.2 $\mathrm{MW_e}$ for LTMR, 6.0 $\mathrm{MW_e}$ for GCMR, and 2.52 $\mathrm{MW_e}$ for HPMR. Applying these per-kW values directly to the modeled ratings of 10, 13, and 5 $\mathrm{MW_e}$ would scale every account with unit output, including accounts that are represented as direct dollar costs and are not tied to reactor rating. The reference per-kW value is therefore converted back to a total before the unit rating is changed.

The published reference cost is separated into a rating-insensitive portion and a portion that scales with reactor rating. Let c_{ref} denote the published cost in \$/kW or \$/kW-yr, P_{ref} the published electrical rating, P the modeled rating, and C_{fix} the rating-insensitive amount in \$ for capital costs or \$/yr for fixed annual costs. Both ratings are expressed in kW in Eq. (3) to match the per-kW cost basis, although reactor ratings are quoted in $\mathrm{MW_e}$ elsewhere in the paper. The adjusted cost is

$$c_{\text{adj}} = \left(c_{\text{ref}} - \frac{C_{\text{fix}}}{P_{\text{ref}}}\right) + \frac{C_{\text{fix}}}{P}. \tag{3}$$

The first term retains the scalable portion of the published cost on its original per-kW basis, while the second term distributes the rating-insensitive portion over the modeled reactor rating.

The adjustment changes the relative capital basis because the ratings do not move by the same proportion. Every modeled unit is larger than its reference, so the fixed dollars are spread over more output and every per-kW cost falls. Aalo-1 moves the least and therefore gains the least: its NOAK OCC falls by 15.9%, against 22.7% for Holos-Quad and 22.8% for eVinci. The resulting values are listed in Appendix B. This treatment is representative: it does not reproduce the component geometry, mass, material, or account-specific scaling exponents of a new bottom-up calculation.

The source cost model reports both overnight capital cost (OCC) and total capital investment (TCI). Across the FOAK and NOAK values considered here, TCI is only approximately 2.00–2.32% higher than the corresponding OCC. The present analysis therefore uses OCC as the capital-cost basis for both the annualized capital charge and the inverse cost calculation. The small difference between OCC and TCI is neglected relative to the broader uncertainty in early-stage microreactor cost estimates. This also keeps the calculated break-even value directly interpretable as an OCC threshold.

After the rating adjustment, cost component z of design j and cumulative unit k is interpolated geometrically between the FOAK and NOAK endpoints[50],

$$z_{j,k} = z_{j,\text{FOAK}} \left(\frac{z_{j,\text{NOAK}}}{z_{j,\text{FOAK}}}\right)^{\log_2(k)/\log_2(N_L)}, \qquad 1 \le k < N_L, \tag{4}$$

with $z_{j,k} = z_{j,\text{NOAK}}$ for $k \ge N_L$ and $N_L = 100$. The cost preprocessing separates each design into learning and non-learning components before Eq. (4) is applied. For capital, only the learning portion of overnight capital cost (OCC) is interpolated between the FOAK and NOAK endpoints; the design-insensitive capital portion is held at its FOAK value. O&M is treated in the same form. Fixed O&M (FOM) is split into learning and non-learning portions, with the O&M staff account, embedded security allocation, and regulatory fee kept outside the learning curve. Fuel and variable O&M (VOM) are combined into the variable-cost category; the current input basis has no separate non-learning variable component. This keeps the explicitly design-insensitive accounts from being reduced by the build-sequence learning curve while retaining a compact six-component cost representation. The geometric path is still a simplified build-sequence assumption, since actual industry learning would depend on production volume, account type, supply-chain repetition, design stability, and project delivery [50].

The main case uses one unit per site, $K = 1$. In the sensitivity analysis, the embedded security account is shared among K co-located units; no other capital, staffing, fuel, or maintenance component is shared.

For unit u, the annualized reference capital cost is calculated using the capital recovery factor (CRF),

$$C_{\mathrm{cap},u} = \mathrm{CRF}(r, L)\, \mathrm{OCC}_u\, P_u, \tag{5}$$

where OCC_u is in \$/kW, P_u is the rated electrical power in kW, and the dimensionless CRF converts the initial capital cost into an equivalent annual cost in \$/yr over the reactor lifetime,

$$\mathrm{CRF}(r, L) = \frac{r(1+r)^L}{(1+r)^L - 1}. \tag{6}$$

The analysis uses $r = 7\%$ and $L = 60$ years. The annual non-capital cost is

$$C_{\mathrm{noncap},u} = \mathrm{FOM}_u\, P_u + \left(c_u^{\mathrm{fuel}} + c_u^{\mathrm{VOM}}\right) E_u, \tag{7}$$

where FOM_u is in \$/kW-yr, c_u^{fuel} and c_u^{VOM} are in \$/MWh, and E_u is the annual dispatched energy in MWh, so that FOM cost is carried by installed capacity while fuel and variable O&M are carried by dispatched energy. Both $C_{\mathrm{cap},u}$ and $C_{\mathrm{noncap},u}$ are annual costs in \$/yr.

OCC is used as the single capital-cost basis in both the annualized capital charge and the reported break-even threshold. The learning model therefore carries the learning and non-learning portions of OCC directly, together with the corresponding fixed and variable operating-cost components.

3.2.3. Merit-order dispatch

Every candidate fleet is expanded into individual learned units. The marginal generation cost of unit u is

$$c_u^{\mathrm{marg}} = c_u^{\mathrm{fuel}} + c_u^{\mathrm{VOM}}, \tag{8}$$

where c_u^{fuel} and c_u^{VOM} are in \$/MWh, and the reactors are sorted from the lowest to the highest value of Eq. (8). Capital cost and FOM do not determine the hourly dispatch order. Each reactor may be dispatched between zero and rated output, so the output of unit u at hour t is

$$P_{u,t} = \min\left[\max\left(L_t - \sum_{v<u} P_{v,t}\ ,\ 0\right),\ P_u\right], \tag{9}$$

where $v < u$ denotes reactors ahead of u in merit order, L_t is the hourly target in MW, and P_u is rated electrical power in $\mathrm{MW_e}$. The total fleet output is therefore

$$P_t^{\mathrm{fleet}} = \min\left(L_t,\ \sum_u P_u\right). \tag{10}$$

A zero hourly endpoint is treated as an off state for the economic commitment schedule. The annual energy of unit u is $E_u = \sum_t P_{u,t} \Delta t$ with $\Delta t = 1$ h. Startup and shutdown events are counted from consecutive hourly outputs. A startup occurs when $P_{u,t-1} = 0$ and $P_{u,t} > 0$, while a shutdown occurs when $P_{u,t-1} > 0$ and $P_{u,t} = 0$. The annual sequence is treated cyclically by connecting the final hour of the year to the first.

A generic nuclear startup cost of 107.86 \$/MW-event in 2019 dollars is taken from the Argonne reactor-size study [51]. The same value is assumed for shutdown, and both are escalated to the 2026 cost basis using the Consumer Price Index for All Urban Consumers (CPI-U) [52], giving $c_{\text{start}} = c_{\text{stop}} = 140.89$ \$/MW-event. The annual cycling cost is

$$C_{\text{cyc}} = \sum_u P_u \left(c_{\text{start}} N_u^{\text{start}} + c_{\text{stop}} N_u^{\text{stop}} \right), \tag{11}$$

where P_u is the rated electrical power of unit u in MW, N_u^{start} and N_u^{stop} are the annual numbers of startup and shutdown events, and c_{start} and c_{stop} are the corresponding cycling-cost coefficients in \$/MW-event. The cycling cost therefore increases with both unit rating and the number of zero-to-positive and positive-to-zero transitions in the hourly dispatch.

After the fleet output is calculated, its electricity is credited against fossil sources in descending order of midpoint levelized cost of electricity (LCOE). In the combined case the order is gas peaking, coal, and gas combined cycle. This is an economic replacement allocation; it does not reconstruct the actual commitment or retirement order of specific DTE units.

3.2.4. Inverse break-even formulation

For one fixed fleet, define the capital coefficient, the non-capital floor, and the displaced-cost target,

$$A = \frac{C_{\text{cap}}}{E_{\text{MR}}}, \tag{12}$$

$$B = \frac{C_{\text{noncap}} + C_{\text{cyc}}}{E_{\text{MR}}}, \tag{13}$$

$$\bar{c}_{\text{rep}} = \frac{\sum_s c_s E_s^{\text{rep}}}{E_{\text{MR}}}, \tag{14}$$

where $C_{\text{cap}} = \sum_u C_{\text{cap},u}$ and $C_{\text{noncap}} = \sum_u C_{\text{noncap},u}$ are the fleet capital and non-cycling non-capital costs in \$/yr, C_{cyc} is the annual startup/shutdown cost of Eq. (11), E_{MR} is the annual microreactor generation in MWh, and E_s^{rep} is the replaced generation of fossil source s in MWh at cost c_s in \$/MWh. The capital coefficient A, the non-capital floor B, and the target $\bar{c}_{\text{rep}}$ are therefore all in \$/MWh. A common dimensionless capital multiplier m is applied to the reference capital path. Since m does not appear in the marginal-cost merit order, the dispatch, cycling events, realized capacity factors, and displaced fossil quantities do not change with m. The cycling charge therefore remains part of the fixed non-capital floor, and the microreactor cost over its served energy is exactly linear,

$$\text{LCOE}_{\text{MR}}(m) = Am + B. \tag{15}$$

with LCOE_{MR} in \$/MWh. At break-even, the microreactor LCOE equals the average cost in \$/MWh of the displaced fossil generation $\bar{c}_{\text{rep}}$.

$$\text{LCOE}_{\text{MR}}(m^*) = \bar{c}_{\text{rep}} = Am^* + B \tag{16}$$

The break-even multiplier m^* is

$$m^* = \frac{\bar{c}_{\text{rep}} - B}{A}, \qquad A > 0, \tag{17}$$

The coefficient A is positive by construction for any nonempty fleet because the annualized reference capital cost and generated energy are both positive. The break-even OCC of design j is

$$\mathrm{OCC}_j^{\mathrm{BE}} = m^* \, \mathrm{OCC}_j^{\mathrm{ref}}. \tag{18}$$

Here $\mathrm{OCC}_j^{\mathrm{ref}}$ is the reference overnight capital cost in \$/kW and m^* is dimensionless, so $\mathrm{OCC}_j^{\mathrm{BE}}$ is also in \$/kW. The LCOE at break-even is therefore not a second independent model output; it is a consistency relation that places the result on the same \$/MWh axis as the Lazard ranges. A negative m^* is retained: it means $B > \bar{c}_{\mathrm{rep}}$, so the fleet costs more than the displaced fossil generation before any capital charge is added. A negative OCC is not physically implementable, but its magnitude shows how far the case is below zero-capital break-even.

3.2.5. Metaheuristic screening

The mixed-fleet decision vector is

$$\mathbf{N} = (N_A, N_H, N_E) \in \mathbb{Z}^3. \tag{19}$$

The three entries are nonnegative integer unit counts. The optimizers minimize

$$F(\mathbf{N}) = -m^*(\mathbf{N}). \tag{20}$$

The decision vector contains reactor counts only. The unit ratings, cost inputs, and learning path of each design are fixed in Section 3.2.2, and the hourly dispatch of Section 3.2.3 follows from the counts, so the fleet composition is the only free choice in the annual model. The number of units per site is held at $K = 1$ rather than optimized, because m^* is non-decreasing in K under the security-sharing model of Section 3.2.6 and an unconstrained search would return the upper bound of K in every case. The effect of K is reported as a sensitivity instead. Each entry of $\mathbf{N}$ is bounded using the number of units required for that design alone to cover the peak of the corresponding target profile,

$$N_{\max,j} = \left\lceil \frac{\max_t L_q(t)}{P_j} \right\rceil + 1, \tag{21}$$

where the additional unit provides a small margin above peak coverage. The bound is calculated separately for each replacement case and reactor design. Designs with lower unit ratings therefore receive larger count bounds, while cases with larger target peaks allow more units. These bounds define the limits of the mixed-fleet search rather than requiring the selected fleet to cover the full target peak.

The multiplier is used rather than the break-even OCC because it is common to all three designs. The break-even OCC of Eq. (18) is defined against a single design-specific reference cost, so a mixed fleet has no single OCC to scale, while m^* remains a single scalar for any composition. Maximizing m^* therefore selects the fleet that supports the highest capital cost relative to its own reference path. The screening is implemented with the NeuroEvolution Optimization with Reinforcement Learning (NEORL) optimization framework [53, 54]. Ten methods within NEORL are adopted and compared: differential evolution (DE) [55], particle swarm optimization (PSO) [56], grey wolf optimizer (GWO) [57], whale optimization algorithm (WOA) [58], Harris hawks optimization (HHO) [59], simulated annealing (SA) [60], tabu search (TS) [61], the prioritized replay evolutionary and swarm algorithm (PESA) and its modern variant (PESA2) [62], and the ensemble of differential evolution variants (EDEV) [63]. Each method is repeated for seeds 1 through 10, and the objective wrapper limits each run to 1,000 fleet evaluations.

The unique fleet candidates evaluated across all algorithm and seed combinations are compared by m^* within each replacement case. The candidate with the largest m^* is retained for the downstream sensitivity and reactor-reference calculations. When different optimizer and seed combinations return the same fleet composition and the same m^*, they represent the same economic solution. Algorithm-level mean, spread, and frequency of finding the selected composition are reported separately to describe search consistency.

3.2.6. Security-cost sharing sensitivity

The economic analysis can select fleets containing tens to hundreds of microreactor units. For fleets of this size, placing every reactor at a separate site is unlikely to represent how a multi-unit deployment would be organized. Several units may instead be co-located and share costs that are associated with the site rather than with an individual reactor. In the adopted cost model, the security account is the main cost that can be identified directly as shareable between co-located units. A units-per-site sensitivity is therefore included to determine how sharing this cost changes the break-even capital threshold.

The main case uses one reactor per site, or $K = 1$, so that every reactor carries a full security allocation. Larger values of K represent co-located sites in which one security allocation is shared among multiple reactors.

K is not an optimization variable. The multiplier m^* is non-decreasing in K, and no term in the model penalizes a larger site. An optimizer would therefore return the upper bound of K in every case, and that bound would be set by the search range rather than by the economics. K is instead swept over a fixed range and reported as a sensitivity.

Only the embedded security allocation is shared. The remaining fixed O&M accounts are tied to the individual reactor or already carry an internal pooling assumption. Sharing them again at the site level would count the same economy twice. The sensitivity therefore represents only the portion of the co-location benefit associated with the explicitly shareable security account. Other possible site-level savings are not modeled.

The sweep is applied to the fleet already selected in Section 3.2.5. The unit counts, dispatch, and displaced energy do not change with K. If up to K units share one security allocation, the number of site allocations is

$$N_{\text{site}} = \left\lceil \frac{N}{K} \right\rceil, \tag{22}$$

and the retained security share per unit is

$$f_{\text{site}} = \frac{N_{\text{site}}}{N}. \tag{23}$$

Here N is the total reactor count of the fleet and K is the number of co-located reactors sharing one allocation, so N_{site} is a count and f_{site} is a dimensionless fraction between $1/K$ and 1. The sensitivity uses $K = 1, 2, 5, 10, 15, 20,$ and 40.

4. Description of Reactor Models

This section presents the reactor-level portion of the analysis. Section 4.1 defines the point-kinetics, reactivity-feedback, xenon, and thermal models used for the three reactor designs. Section 4.2 gives the nonlinear model predictive control formulation and actuator constraints. Section 4.3 describes how the hourly economic dispatch is converted into candidate-specific reactor references, and Section 4.4 presents the controller tuning and worst-day replay used for the dynamic screen.

The total reactivity of Eq. (28) has two parts. The control system supplies an external reactivity, and the fuel, moderator, coolant, and xenon states supply feedback. This section describes the external term, since it is the only actuated quantity in the reactor-level model.

A conventional plant moves control rods linearly through the core. The designs considered here instead rotate cylindrical drums seated in the outer region of the core. Each drum carries a boron carbide absorber over part of its circumference and reflector material over the remainder. Rotating a drum changes how close the absorber sits to the fueled region, which changes the absorption cross section seen by the core. Holos-Quad uses eight drums, and Aalo-1 and eVinci use twelve. The physical drums can be actuated individually, while the reduced-order controller represents the selected active drums using a common bank angle.

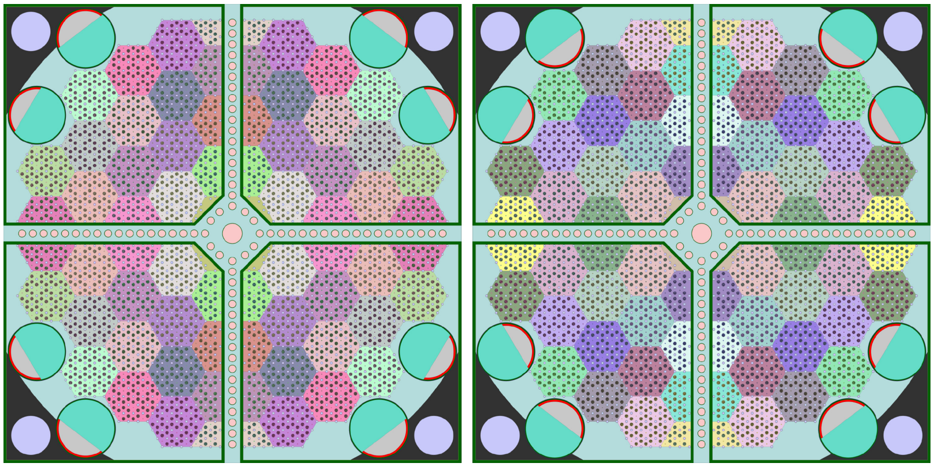

Figure 3: Control-drum arrangement of the Holos-Quad microreactor, reproduced from [10]. The two configurations illustrate the limiting orientations of the rotating control drums surrounding the fueled region. Rotation changes the absorber orientation relative to the core and therefore changes the external reactivity available for power control. The reduced-order reactor model represents this behavior through the nonlinear drum-worth relation of Eq. (29).

Figure 3 shows the two limiting positions. At full outward rotation the absorber faces away from the fueled region and reactivity is at its maximum. At full inward rotation the absorber faces the fueled region and reactivity is at its minimum. The angle convention places $\theta = 0^\circ$ at the fully inserted position.

First-order perturbation theory relates the drum reactivity to the absorption perturbation [64],

$$\rho = \frac{\delta\Sigma_a}{\Sigma_a + DB^2}, \tag{24}$$

where $\delta\Sigma_a$ is the perturbation to the macroscopic absorption cross section, Σ_a the unperturbed absorption cross section, D the diffusion coefficient, and B the geometric buckling. The drum reactivity is therefore proportional to the change in absorption that rotation produces. Integrating over the rotation gives the raised-cosine worth of Eq. (29).

Two drum properties set the maneuvering capability of a design. The integral drum worth is the reactivity difference between the fully withdrawn and fully inserted positions, summed over the active drums. It bounds how far the reactor can be moved from its initial power. The values used here are listed in Table A.1. The maximum drum rate $\dot{\theta}_{\max}$ is the fastest rotation the actuator can produce. It bounds how quickly that reactivity can be inserted and therefore sets the achievable ramp rate. The analysis uses $\dot{\theta}_{\max} = 0.5$ degrees/s for all three designs, applied through Eqs. (39) and (40).

4.1. Reduced-order Reactor Model

The reactor-level model follows the point-kinetics and thermal-feedback structure used in previous microreactor load-following work [10, 64]. Its purpose is to provide a fast plant model whose dominant time constants are physically representative, so that the maneuvering capability of the three designs can be compared on a common basis.

The state vector is

$$\mathbf{x} = \begin{bmatrix} \bar{n} & \bar{C}_1 & \bar{C}_2 & \cdots & \bar{C}_6 & X & I & T_f & T_m & T_c \end{bmatrix}^T, \tag{25}$$

which collects twelve states: the normalized reactor power $\bar{n} = n/n_0$, where n is the neutron density and n_0 is the reference neutron density at rated power; the six normalized delayed-neutron precursor concentrations $\bar{C}_i = C_i/C_{i,0}$, where $C_{i,0}$ is the corresponding equilibrium precursor concentration at the rated-power reference state; the ^{135}Xe and ^{135}I concentrations X and I; and the fuel, moderator, and coolant temperatures T_f, T_m, and T_c. Power is carried in per-unit form so that a single controller structure and reference trajectory apply to all three designs regardless of rating.

The six-group point-kinetics equations are [65]

$$\frac{d\bar{n}}{dt} = \frac{\rho - \beta}{\Lambda}\bar{n} + \sum_{i=1}^{6} \frac{\beta_i}{\Lambda}\bar{C}_i, \tag{26}$$

$$\frac{d\bar{C}_i}{dt} = \lambda_i \left(\bar{n} - \bar{C}_i\right), \qquad i = 1, \ldots, 6, \tag{27}$$

where Λ is the prompt-neutron generation time, β the total effective delayed-neutron fraction, β_i and λ_i the fraction and decay constant of precursor group i, and ρ the total reactivity. Equation (26) balances prompt production against the delayed neutrons returned by precursor decay, while Eq. (27) relaxes each group toward the instantaneous power at its own rate.

The total reactivity contains control-drum, temperature, and xenon feedback,

$$\rho = \rho_d(\theta) + \alpha_f(T_f - T_{f0}) + \alpha_m(T_m - T_{m0}) + \alpha_c(T_c - T_{c0}) - \frac{\sigma_X}{\Sigma_f}(X - X_0). \tag{28}$$

where α_f, α_m, and α_c are the fuel, moderator, and coolant reactivity coefficients, T_{f0}, T_{m0}, and T_{c0} are the corresponding reference temperatures at rated power, σ_X is the microscopic absorption cross section of ^{135}Xe, Σ_f is the macroscopic fission cross section, and X_0 is the equilibrium xenon concentration at the rated-power reference state.

The integral drum worth is represented by

$$\rho_d(\theta) = \frac{\rho_{\max}}{2}\left[1 - \cos\left(\frac{\pi\theta}{180}\right)\right] - \rho_{\text{ss}}. \tag{29}$$

where θ is the drum rotation angle in degrees, $\rho_{\max}$ the total integral worth available at $\theta = 180°$, and the factor $\pi/180$ converts the angle to radians. Rotating a drum changes how much of its boron absorber faces the fueled region, and first-order perturbation theory gives the resulting worth this raised-cosine shape [64]. Here ρ_{ss} is the steady-state drum reactivity required to hold the reactor critical at the rated-power reference drum angle θ_0. The equilibrium xenon concentration X_0 appearing in Eq. (28) is evaluated at the same rated-power reference state.

$$\frac{dI(t)}{dt} = \gamma_I \Sigma_f v n_0 \bar{n}(t) - \lambda_I I(t), \tag{30}$$

$$\frac{dX(t)}{dt} = \gamma_X \Sigma_f v n_0 \bar{n}(t) - \lambda_X X(t) + \lambda_I I(t) - \sigma_X v n_0 \bar{n}(t) X(t), \tag{31}$$

where $I(t)$ and $X(t)$ are the iodine-135 and xenon-135 concentrations, γ_I and γ_X are their respective fission yields, Σ_f is the macroscopic fission cross section, v is the average thermal-neutron speed, n_0 is the reference neutron density, $\bar{n}(t)$ is the normalized neutron density, λ_I and λ_X are the iodine and xenon decay constants, and σ_X is the microscopic xenon absorption cross section.

A lumped three-temperature model represents fuel, moderator, and coolant feedback,

$$m_f c_f \frac{dT_f}{dt} = f_f P_0 \bar{n} - K_{fm}(T_f - T_m), \tag{32}$$

$$m_m c_m \frac{dT_m}{dt} = (1 - f_f) P_0 \bar{n} + K_{fm}(T_f - T_m) - K_{mc}(T_m - T_c), \tag{33}$$

$$m_c c_c \frac{dT_c}{dt} = K_{mc}(T_m - T_c) - Q_{\text{rem}}(T_c), \tag{34}$$

where m and c with the corresponding subscript are the mass and specific heat of each node, K_{fm} and K_{mc} are the fuel-to-moderator and moderator-to-coolant heat-transfer coefficients, f_f is the fraction of fission power deposited in the fuel, P_0 is the rated thermal power of one unit, and Q_{rem} is the heat removed from the coolant node. Heat flows through the three nodes in series, and the products mc together with the coefficients with the time increments determines how strongly thermal feedback resists a commanded ramp. Holos-Quad and Aalo-1 use a convective coolant-removal coefficient. The eVinci proxy uses a heat-pipe removal coefficient to represent its passive heat-transfer path [48, 66]. For Holos-Quad and Aalo-1, heat removal is represented by

$$Q_{\text{rem}} = 2\dot{m} c_c (T_c - T_{\text{in}}), \tag{35}$$

which places the nominal coolant node at the mean of the inlet and outlet temperatures. For the eVinci-like heat-pipe branch,

$$Q_{\text{rem}} = K_{\text{hp}}(T_c - T_{\text{sink}}), \tag{36}$$

where T_c is the heat-pipe coupling node and T_{sink} is the condenser-side sink temperature.

These parameter sets are reduced-order comparison models and are not vendor-validated licensing models but are sufficient to capture power maneuvering for our load following studies. The numerical parameters used in the reactor models are listed in Appendix A. The Aalo-1 and eVinci parameter sets are comparative proxy values and should not be interpreted as vendor-validated design data. The authors developed these parameter sets to the best of their ability using the publicly available data.

4.2. Nonlinear model predictive control

Model predictive control is an optimal control strategy for constrained systems. At each sampling instant it predicts the reactor response over a finite horizon, solves for the input sequence that minimizes a cost over that horizon, applies only the first drum position, and repeats the calculation from the new state [67]. The drum position and rate limits therefore enter the controller as constraints rather than as saturation applied after the command is calculated.

Nonlinear model predictive control uses the nonlinear reactor model directly in the prediction step [10]. This is useful here because the point-kinetics equation contains the product $\rho\bar{n}$, the xenon equation contains a neutron-density-dependent burnout term, and the drum worth is nonlinear in angle. The full reactor state is propagated internally by the prediction model, while the controlled output is the normalized reactor power $\bar{n}$. The manipulated control input is the angular position of the active control-drum bank, θ. The NMPC therefore searches over a sequence of allowable future drum positions and uses the nonlinear reactor model to determine the corresponding power response over the prediction horizon.

The objective used at each controller update is

$$J = \sum_{k=0}^{N_p-1} Q\left(\bar{n}_{k+1} - \bar{n}_{k+1}^{\text{ref}}\right)^2 + \sum_{j=0}^{N_c-1} R\,\theta_j^2 + \sum_{j=1}^{N_c-1} R_d\left(\theta_j - \theta_{j-1}\right)^2. \tag{37}$$

where J is the scalar cost assigned to a candidate control sequence, N_p is the prediction horizon, and N_c is the control horizon. The first term measures the predicted tracking error between reactor power and the requested reference. The weight Q therefore determines how strongly the controller prioritizes power tracking. The second term penalizes large control-drum angles through R, which limits the tendency of the optimizer to use unnecessarily large drum positions to reduce tracking error. The third term penalizes changes between consecutive drum positions through R_d and therefore discourages rapid changes in the planned control sequence. Together, the three terms define the tradeoff between tracking accuracy, drum position, and drum movement. The interpolated reference at the current controller step is held over the prediction horizon.

The drum position and rate constraints are

$$0 \leq \theta_j^{\text{pred}} \leq 180^\circ, \tag{38}$$

$$-\dot{\theta}_{\max} T_s \leq \theta_0^{\text{pred}} - \theta_{\text{prev}} \leq \dot{\theta}_{\max} T_s, \tag{39}$$

$$-\dot{\theta}_{\max} T_s \leq \theta_j^{\text{pred}} - \theta_{j-1}^{\text{pred}} \leq \dot{\theta}_{\max} T_s, \qquad j = 1, \ldots, N_c - 1. \tag{40}$$

The first constraint defines the physical input space of the control drums, from the fully inserted to the fully withdrawn position. The second limits the first predicted move relative to the drum position applied during the previous controller step, while the third applies the same rate limit between subsequent positions within the control horizon. These constraints represent the finite actuator speed and prevent the optimizer from producing unrealistically fast changes in external reactivity. With a maximum drum rate of 0.5 degrees/s and a controller interval of $T_s = 60$ s, the drum position can change by at most 30° during one controller update. A requested power change that would require a faster drum motion must therefore be followed with some tracking error rather than by violating the actuator limit.

The constrained problem is solved by sequential least-squares programming. The prediction and control horizons are $N_p = 5$ and $N_c = 2$, and the controller interval is $T_s = 60$ s, corresponding to 5-min and 2-min prediction and control horizons. The nonlinear state equations are propagated with an implicit backward differentiation formula (BDF) integrator over each interval. With $N_c = 2$, two future drum positions are optimized at each controller update, while the reactor response is predicted over five controller intervals. Only the first optimized drum position is applied before the NMPC problem is solved again from the updated reactor state. This keeps the same general receding-horizon structure as the reference microreactor NMPC implementation while allowing the economic schedule to be replayed over a full day [10].

4.3. Economic dispatch reference and worst-day selection

The reactor-level reference is constructed only after the economic optimization has selected a fleet. This keeps the dynamic calculation tied to the operating schedule that produced the economic result rather than applying the same prescribed transient to every candidate fleet. Let $P_{d,h}$ denote the total dispatched fleet power on day d at hour h. The variability of each day is calculated as

$$V_d = \sum_{h=0}^{22} |P_{d,h+1} - P_{d,h}|, \tag{41}$$

where V_d represents the total movement of the fleet during the day. The metric sums the magnitude of each consecutive hourly change, so repeated increases and decreases in output are retained rather than considering only the largest single ramp. The day with the largest V_d is selected for the reactor-level screen. Because the hourly fleet output depends on both the selected reactor counts and their merit-order dispatch, the resulting most-variable day is specific to each economic replacement case.

After the most-variable day is selected, the day index is omitted for clarity. The annual economic calculation dispatches individual physical reactors, and this unit-level information is retained when the reference is transferred to the reactor model. The fleet output is therefore not redistributed equally among the reactors at reactor time resolution. Instead, each physical unit follows the hourly endpoints assigned to it by the economic dispatch. Between hours h and $h+1$, the reference of unit u is reconstructed by linear interpolation,

$$P_u(t) = (1-\alpha)P_{u,h} + \alpha P_{u,h+1}, \qquad \alpha = \frac{t-t_h}{3600}, \tag{42}$$

where $P_{u,h}$ and $P_{u,h+1}$ are the dispatch values of unit u at the two hourly endpoints. Within each hour, α increases from zero to one, giving a continuous power reference between the two economic-dispatch points rather than an instantaneous change at the hour boundary. The reference contains 25 hourly endpoints: the 24 points of the selected day together with the first endpoint of the following day. The additional point completes the interpolation of the final hour.

The same interpolation is preserved at the fleet level. Let P_h denote the total fleet output at hour h of the selected day. Summing the individual unit references gives

$$\sum_u P_u(t) = (1-\alpha)P_h + \alpha P_{h+1}, \tag{43}$$

so the sum of the physical-unit references remains equal to the interpolated hourly fleet schedule. This is important because the economic dispatch may contain units at full power, units at partial power, and units that are off during the same hour. Redistributing the aggregate fleet output would remove this merit-order structure and could create reactor trajectories that were not present in the economic calculation.

The hourly schedule also contains transitions between online and off states. An interval is treated as online load following only when both hourly endpoints are positive. A zero-to-positive interval defines a startup boundary, a positive-to-zero interval defines a shutdown boundary, and a zero-to-zero interval remains off. These classifications are retained from the economic schedule, while NMPC is applied only to positive-to-positive intervals. The reactor calculation therefore tests whether an already-online unit can follow the power changes assigned by the economic dispatch; it does not model reactor startup or shutdown. Startup and shutdown remain part of the economic calculation through the cycling-cost treatment.

Several physical reactors may receive exactly the same sequence of hourly dispatch values. Units of the same design with identical 24-point profiles are therefore grouped so that the same NMPC calculation does not have to be repeated. One representative trajectory is simulated for each unique profile, and its response is multiplied by the number of physical units represented by that profile when the fleet response is reconstructed. This grouping reduces computational effort only and does not change the unit dispatch or the total fleet reference. The hourly-linear reference is the primary reactor-level check because it is constructed directly from the temporal resolution available in the annual economic dispatch. A second MISO-informed reference is included only as a higher-resolution load-following sensitivity. Public 5-min MISO Real-Time Cleared Offers data provide a normalized within-hour shape that is rescaled around each hourly economic-dispatch fleet value and then interpolated to the 60-s NMPC interval. The MISO-informed trajectory is therefore used as a surrogate for additional intra-hour variation rather than as a reconstruction of DTE operation. The matching and reconstruction details are given in Appendix D [68].

4.4. Evolution-strategy controller tuning and worst-day replay

The NMPC weights are tuned offline using the evolution strategy (ES) implementation in NEORL [53, 54]. ES provides a derivative-free search over bounded continuous variables and has also been applied through NEORL to microreactor control-drum optimization [25]. The active drum counts are fixed at 8 for Holos-Quad, 8 for the Aalo/LTMR branch, and 12 for the eVinci/HPMR branch. For a selected active subset, the available integral worth is scaled from the full-bank value in Table A.1 in proportion to the number of active drums, and the selected drums are represented by the common bank angle θ. The tuning vector for an active design i is

$$z_i = [\log_{10} Q_i, \log_{10} R_i, \log_{10} R_{d,i}], \tag{44}$$

with

$$2 \le \log_{10} Q_i \le 10, \tag{45}$$

$$-8 \le \log_{10} R_i \le 1, \tag{46}$$

$$-4 \le \log_{10} R_{d,i} \le 1. \tag{47}$$

The logarithmic variables allow all three controller weights to be searched over several orders of magnitude. During NMPC operation, Q_i, R_i, and $R_{d,i}$ are fixed while the controller solves for the future drum sequence. During tuning, the three logarithmic weights are the decision variables and each candidate is evaluated from its closed-loop tracking response.

For each reference basis, a continuous 1-h online window is selected from the most demanding positive-to-positive portion of each candidate-specific day. The hourly-linear case uses the interval containing the largest hourly fleet change, while the MISO-informed sensitivity uses the 1-h interval with the largest cumulative power movement in the reconstructed 5-min reference. Startup and shutdown boundaries are excluded. If a selected fleet contains more than one reactor design, each design is tuned separately using only the cases in which that design has an online load-following interval. Tracking error is measured using the mean absolute percentage error (MAPE), and the tuning objective is

$$F_i(z_i) = \max_{c \in \mathcal{C}_i} \mathrm{MAPE}_{i,c}(z_i), \tag{48}$$

where $\mathcal{C}_i$ is the set of applicable replacement cases for design i. The 0.5 degrees/s drum-rate limit remains enforced within each online NMPC solution, and a worst-case tuning MAPE below 1% is used as the tracking criterion.

The ES search uses $\lambda = 60$ individuals, $\mu = 30$ survivors, a crossover probability of 0.6, a mutation probability of 0.3, blend crossover with $\alpha = 0.5$, and 10 generations. Eight parallel workers evaluate the population. These settings give 660 outer objective evaluations for each active reactor design. The hourly-linear case uses the bounds of Eqs. (45)–(47). The two reference bases are tuned independently so that the higher-resolution sensitivity does not alter the controller used for the primary hourly-linear check.

The selected controller weights are then held fixed within their respective reference basis while all changing online intervals are replayed to reconstruct the full candidate-specific day. In a mixed fleet, each physical unit uses the controller tuned for its own design, and units with identical profiles remain grouped only for computational efficiency. Dynamic tracking error is evaluated only on the time points belonging to simulated positive-to-positive intervals. The reported comparison uses fleet MAPE, maximum absolute fleet tracking error, and maximum drum rate for both reference bases. The resulting calculations are online load-following capability screens and do not model reactor startup or shutdown.

5. Results

This section presents the annual economic results and the candidate-specific worst-day references obtained from the selected fleets. The economic results come from the merit-order dispatch, cycling-cost, and inverse break-even model of Section 3, evaluated on the 2025 DTE profiles with midpoint fossil costs unless a range is stated. The worst-day references are constructed from the same unit-level dispatch using Section 4.3. Interpretation of the economic results and their connection to the reactor-level workflow is given in Section 6.

5.1. Break-even results: selected fleets and per-design thresholds

The mixed-fleet screening and the exact single-design enumeration were run for each of the four replacement cases. In every case, the selected fleet from the NEORL screening contains only Aalo-1 units, and the same Aalo-1 unit count is obtained from the single-design enumeration. Table 5 reports the complete single-design results for all three reactor designs. The Aalo-1 rows therefore also correspond to the selected mixed-fleet results, while the Holos-Quad and eVinci rows show the best pure-design result for the same replacement case. The table gives the unit count, installed capacity, break-even multiplier m^*, break-even OCC at the FOAK and NOAK reference costs, replaced-energy fraction, and realized fleet capacity factor (CF).

Table 5: Break-even results for all three designs at midpoint fossil-generation costs from the exact single-design enumeration. The Aalo-1 rows also correspond to the selected mixed-fleet results. Negative values indicate that the non-capital cost floor is above the displaced fossil cost; they are not feasible negative construction costs.

Design	Case	Units	Capacity (MW)	m^*	FOAK OCC* ($/kW)	NOAK OCC* ($/kW)	Replaced (%)	CF (%)
Aalo-1	Coal	185	1850	0.629555	14,087	6,484	63.83	98.66
	Gas combined cycle	111	1110	0.103166	2,309	1,063	96.73	84.58
	Gas peaking	35	350	0.234395	5,245	2,414	40.49	46.15
	Coal and gas	124	1240	0.865866	19,375	8,918	29.32	100.00
Holos-Quad	Coal	158	2054	0.521247	12,043	5,061	70.11	97.61
	Gas combined cycle	87	1131	−0.045977	−1,062	−446	98.28	84.33
	Gas peaking	31	403	0.055520	1,283	539	45.27	44.82
	Coal and gas	136	1768	0.736341	17,013	7,149	41.76	99.89
eVinci	Coal	954	4770	−0.623762	−29,574	−12,516	100.00	59.95
	Gas combined cycle	241	1205	−0.841257	−39,886	−16,880	100.00	80.54
	Gas peaking	210	1050	−0.435339	−20,641	−8,735	85.40	32.45
	Coal and gas	196	980	−0.587891	−27,874	−11,796	23.17	100.00

The design comparison is consistent across the four replacement cases. Aalo-1 gives the highest positive break-even multiplier in every case and is the design retained by the mixed-fleet screening. Holos-Quad remains positive for coal, gas peaking, and the combined coal-and-gas case, but its combined-cycle threshold is negative. eVinci has a negative threshold in all four cases. A negative m^* means that the non-capital cost floor, including fixed O&M, fuel, variable O&M, and the cycling charge, is already above the average cost of the fossil generation displaced by that fleet. In those cases, reducing OCC alone cannot produce break even.

Among the selected Aalo-1 fleets, the combined coal-and-gas case gives the largest multiplier ($m^* = 0.866$) at an approximately 100% capacity factor while replacing 29.32% of the target energy. The gas combined-cycle case gives the smallest positive multiplier ($m^* = 0.103$) while replacing 96.73% of its target energy. This shows that a larger replacement fraction does not necessarily support a larger capital threshold. The cost of the displaced generation and the utilization of the fleet are more important than the replacement fraction by itself.

Figure 4 shows the original, replaced, and remaining generation for the selected Aalo-1 fleet of each case. In the coal and combined cases the fleet forms a nearly constant block below the load; in the combined-cycle case the fleet capacity sits near the upper range of the load; and in the gas-peaking case the fleet band lies above the intermittent load for much of the year.

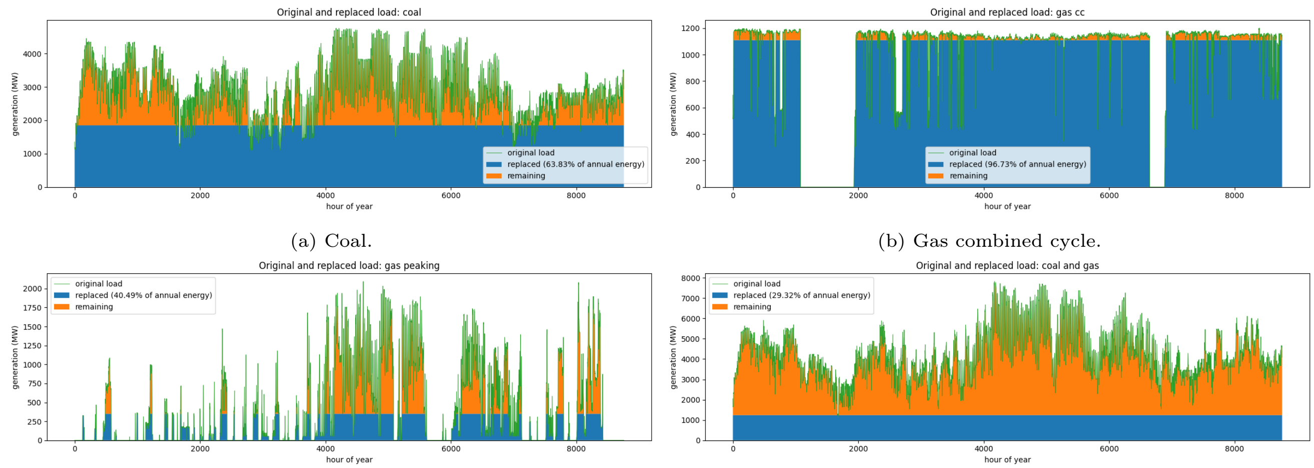

(a) Coal. (b) Gas combined cycle.

(c) Gas peaking. (d) Combined coal and gas.

Figure 4: Hourly fossil-generation replacement by the selected Aalo-1 fleet for each replacement case at midpoint fossil-generation costs. The original fossil-generation profile, the portion replaced by microreactor generation, and the remaining fossil generation are shown over the 8,760-h study year. The selected fleets contain (a) 185 Aalo-1 units for coal, replacing 63.83% of annual target energy; (b) 111 units for gas combined cycle, replacing 96.73%; (c) 35 units for gas peaking, replacing 40.49%; and (d) 124 units for the combined coal-and-gas case, replacing 29.32%. The different profiles show that the amount of energy replaced depends on both the installed fleet capacity and the hourly shape of the fossil-generation target.

5.2. Sensitivity to fossil cost and comparison against Lazard v19.0

The selected Aalo-1 fleets were re-evaluated at the low, midpoint, and high fossil costs of Table 3, holding the fleet count fixed and changing only the displaced fossil cost. Table 6 reports the resulting NOAK break-even OCC. At the low costs, the coal and combined cases remain positive; at the high costs, all four cases are positive.

Table 6: Aalo-1 NOAK OCC* sensitivity to fossil-generation cost ($/kW).

Case	Low	Midpoint	High
Coal	630	6,484	12,339
Gas combined cycle	−2,475	1,063	4,600
Gas peaking	−348	2,414	5,177
Coal and gas	2,780	8,918	15,057

The microreactor fleet LCOE is compared with the Lazard 2026 LCOE+ Version 19.0 conventional-generation ranges in Figure 5 [42]. Two features of the comparison are relevant.

The break-even targets are low relative to the Lazard new-build U.S. nuclear range. The coal and combined-cycle markers fall below the lower bound of the Lazard U.S. nuclear range, so the fleet would have to reach a cost below that of new conventional nuclear to break even in those cases. The gas-peaking marker is the only one that falls inside the nuclear range, because gas peaking is the most expensive displaced source in the study.

The distance between the endpoint bands and the break-even markers differs sharply by case. The coal and combined coal-and-gas rows are much closer to break even than the two gas rows. These endpoint bars are illustrative all-FOAK and all-NOAK bounds; the inverse thresholds themselves use the learned build sequence of Eq. (4). The MOUSE cost basis is also presented by its authors as a design-stage estimate with scope for further reduction [6].

5.3. Comparison with generic microreactor OCC estimates

The calculated break-even OCC values are economic thresholds rather than predictions of future reactor costs. To place these thresholds in an external context, they are compared with the generic microreactor cost assumptions

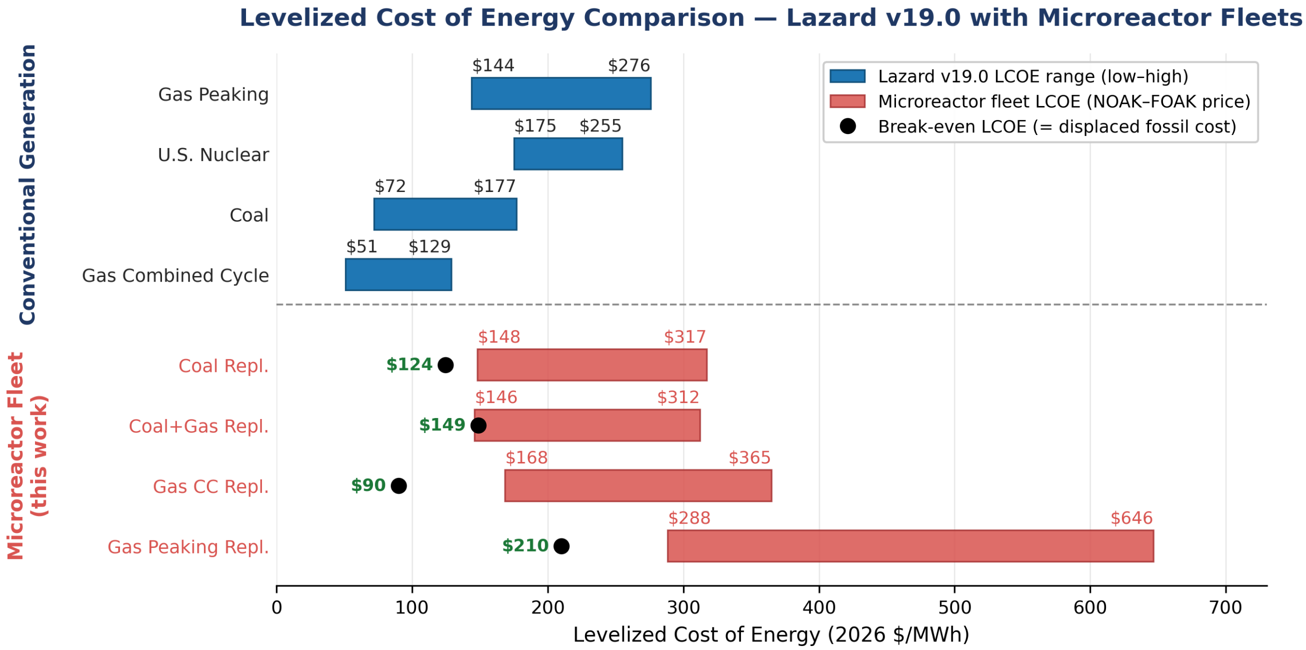

Figure 5: Levelized cost of electricity for the selected Aalo-1 fleets compared against the Lazard 2026 LCOE+ Version 19.0 conventional-generation ranges [42]. The conventional-generation bars give Lazard's unsubsidized new-build LCOE ranges. The microreactor bars give the fleet LCOE when every unit is evaluated at the NOAK and FOAK reference OCC endpoints. The marker on each microreactor row is the break-even LCOE, which is equal to the average cost of the fossil generation displaced by that fleet. The endpoint bars provide an LCOE-scale comparison and should not be interpreted as the capital multiplier of the learned build sequence used in the inverse-cost calculation.

published by the Nuclear Energy Institute (NEI) [69]. The NEI study used a first-unit overnight capital cost range of 10,000–20,000 $\$/kW_e$. Its learning scenarios produced 50th-reactor costs ranging from 3,996 $\$/kW_e$ under a high learning rate to 14,973 $\$/kW_e$ under a low learning rate. The first-unit range is treated here as a generic FOAK benchmark, while the 50th-reactor range is treated as a mature-unit or NOAK proxy. These estimates are generic and are not quotations for Aalo-1, Holos-Quad, or eVinci.

Only the replacement scenarios whose calculated break-even OCC falls within the corresponding generic range are shown in Table 7.

Table 7: Calculated break-even OCC scenarios that fall within the generic NEI first-unit and 50th-reactor cost estimates [69]. Values are shown in $/kW.

Generic cost benchmark	Published range ($/kW)	Replacement scenarios within the range ($/kW)
First unit—FOAK proxy	10,000–20,000	Coal: Aalo-1 (185 units), 14,087; Holos-Quad (158 units), 12,043 Coal and gas: Aalo-1 (124 units), 19,375; Holos-Quad (136 units), 17,013
50th reactor—NOAK proxy	3,996–14,973	Coal: Aalo-1 (185 units), 6,484; Holos-Quad (158 units), 5,061 Coal and gas: Aalo-1 (124 units), 8,918; Holos-Quad (136 units), 7,149

The NEI first-unit and 50th-reactor estimates are generic microreactor benchmarks rather than vendor quotations.

The coal and combined coal-and-gas cases place both Aalo-1 and Holos-Quad within the generic first-unit and mature-unit ranges. These overlaps do not indicate that either design will necessarily achieve the calculated cost. Rather, they show that the capital-cost levels supported by these replacement cases are within the range of generic microreactor cost assumptions reported by NEI [69].

The remaining midpoint scenarios are outside the generic ranges. The Aalo-1 gas combined-cycle and gas-peaking thresholds are below the lower bounds of both ranges. The Holos-Quad gas-peaking threshold is below both lower bounds, the Holos-Quad combined-cycle threshold is negative, and all eVinci cases have negative midpoint break-even values. The eVinci values used in this study represent a generic heat-pipe microreactor cost archetype and reduced-order reactor proxy; they are not Westinghouse cost quotations or vendor-validated eVinci parameters.

A design-specific comparison was also made using the Army G-4 mobile nuclear power plant study, which reports cost estimates for a 13.266 MW_e Holos Quad [70]. The study reports 2017 best-estimate OCC values of 9,488 \$/kW for FOAK and 5,535 \$/kW for NOAK. After adjustment to June 2026 dollars using the Consumer Price Index for All Urban Consumers (CPI-U), these values become approximately 12,926 and 7,541 \$/kW, respectively [52].

The rating-adjusted Holos-Quad inputs that are used in this study are 23,105 \$/kW for FOAK and 9,709 \$/kW for NOAK, both above the corresponding adjusted Army estimates. The adjusted FOAK value is above the 12,043 \$/kW coal threshold but below the 17,013 \$/kW combined-case threshold. The adjusted NOAK value remains above the 7,149 \$/kW combined threshold and the 5,061 \$/kW coal threshold.

5.4. Units-per-site sensitivity

The sensitivity is applied to the selected Aalo-1 fleets of Table 5. The unit counts are held fixed and only the number of reactors sharing one security allocation is changed, so the dispatch, replaced energy, and capacity factor are identical at every K. The reported quantity is the change in the break-even threshold alone.

Security is the only account in the cost basis that is explicitly attributable to a site rather than to a reactor, so it is the only cost for which co-location can be represented without introducing an assumption outside the reference. It is not a small account. At the mature fixed O&M of the model it is 101.1 of 320.2 \$/kW-yr, or about 32% of the total, and about 79% of the fixed O&M that does not decline through learning. Sharing it is therefore the one available measure of how much co-location could reduce the required capital cost.

Figure 6 shows all four cases over $K = 1, 2, 5, 10, 15, 20, 40$. The combined coal-and-gas case rises from $m^* = 0.866$ at $K = 1$ to 0.984 at $K = 40$, the coal case from 0.630 to 0.753, gas combined cycle from 0.103 to 0.220, and gas peaking from 0.234 to 0.333. On the NOAK reference basis the break-even OCC rises from 8,918 to 10,133 \$/kW, 6,484 to 7,754 \$/kW, 1,063 to 2,267 \$/kW, and 2,414 to 3,431 \$/kW.

The gain saturates early. Five reactors per site capture 82% of the total improvement available out to forty, and ten capture 92%. The percentages are within one point of each other across all four cases. Beyond ten units the ceiling of Eq. (22) dominates and further consolidation removes few remaining allocations. Most of the economic benefit of co-location is therefore available at a small site, and a plant of tens of units per site is not required to obtain it.

This is consistent with the direction taken by the Aalo-1 vendor, whose Aalo Pod product groups five Aalo-1 reactors into a single module [45]. At $K = 5$ the coal break-even OCC rises to 7,528 \$/kW and the combined case to 9,920 \$/kW. The result here does not establish that five is the optimal grouping, since the model shares only one account and imposes no siting or licensing limit. It does show that a module of that size sits in the range where most of the modeled security-sharing benefit has already been captured.

The absolute gain is similar in all four cases, between about 1,017 and 1,270 \$/kW from $K = 1$ to $K = 40$. The shared cost is a fixed annual amount per reactor, so the improvement in the multiplier is the ratio of the security saving to the fleet capital charge, and both scale with fleet size. The remaining spread follows fleet size through learning: the 185-unit coal fleet has the lowest average capital cost per unit and the largest gain, and the 35-unit gas-peaking fleet the smallest.

The relative gain differs sharply. The combined-cycle threshold rises by more than a factor of two, while the combined coal-and-gas threshold rises by about 14%. A fixed absolute improvement matters most where the threshold is smallest. Security sharing does not change the ordering of the four cases.

Only the embedded security account is shared, so these values represent the portion of the co-location benefit associated with that account rather than the full economic effect of co-location.

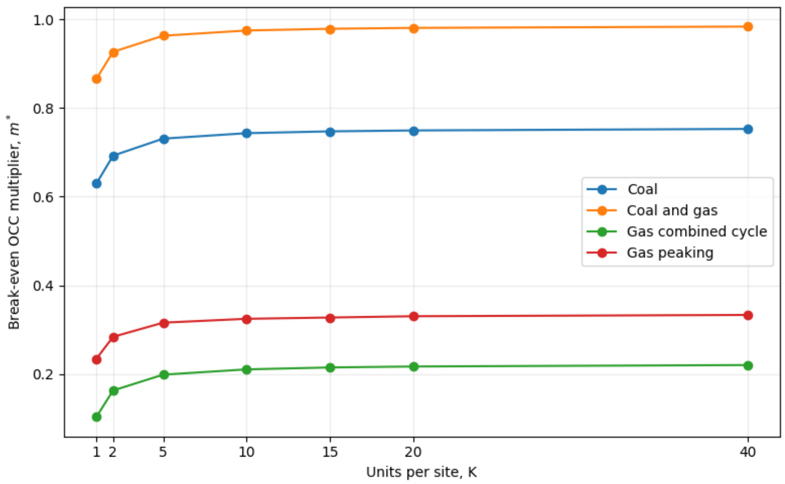

Figure 6: Sensitivity of the break-even capital multiplier m^* to the number of microreactor units sharing one security allocation. The selected fleet composition, hourly dispatch, replaced fossil energy, and realized capacity factor are held fixed while K is varied over 1, 2, 5, 10, 15, 20, and 40 units per site. At $K = 1$, each reactor carries a full security allocation; increasing K reduces only the embedded security cost according to Eqs. (22) and (23). No other capital, staffing, fuel, or maintenance account is shared. The early flattening of each curve shows that most of the modeled security-sharing benefit is obtained at relatively small values of K.

5.5. NEORL screening behavior

The mixed-fleet NEORL calculation uses ten optimizers, ten seeds, and a common budget of 1,000 fleet evaluations per run. The same fleet bounds, objective function, and evaluation budget are used for every method. Several optimizers find the same selected Aalo-1 fleet in each replacement case. DE, GWO, HHO, PESA2, and WOA find the selected fleet in all four cases, while EDEV finds it in three. Across the ten methods, the selected composition is found by six algorithms for coal, seven for gas combined cycle, six for gas peaking, and five for the combined coal-and-gas case.

When different optimizer and seed combinations return the same fleet counts and the same m^*, they represent the same economic solution. Other runs return either nearby Aalo-1 counts or mixed fleets with a lower m^*. PESA, PSO, SA, and TS show larger variation across seeds under the same evaluation budget. The complete mean, standard deviation, and best result for every algorithm are reported in Appendix C.

The NEORL calculation searches the full three-design fleet space, where the numbers of Aalo-1, Holos-Quad, and eVinci units are all allowed to vary. The selected fleet in each replacement case contains only Aalo-1 units. Once this pure-design result is obtained, the remaining problem becomes one-dimensional because the Holos-Quad and eVinci counts are zero. A complete integer enumeration of the Aalo-1 count is therefore performed as a direct check of the fleet size returned by NEORL. The enumeration gives the same Aalo-1 count in all four replacement cases. It is used only to verify the pure-design solution identified by the mixed-fleet search and is not used to determine the reactor design in advance.

5.6. Reactor-level load-following results

All four selected economic fleets contain only Aalo-1 units, so both reactor-level screens use the eight-drum Aalo/LTMR branch. The hourly-linear and MISO-informed reference bases are tuned independently. For the hourly-linear replay, ES gives

$$Q = 1.2081 \times 10^4, \qquad R = 1.36 \times 10^{-4}, \qquad R_d = 2.6058, \tag{49}$$

with a worst-case 1-h tuning MAPE of 0.004557%. For the MISO-informed sensitivity, ES gives

$$Q = 2.1876 \times 10^3, \qquad R = 1.0 \times 10^{-6}, \qquad R_d = 2.598 \times 10^{-3}, \tag{50}$$

with a worst-case 1-h tuning MAPE of 0.496988%. These weights are then held fixed within their respective reference basis during the complete selected-day replay.

Table 8 compares the hourly-linear and MISO-informed results. The hourly-linear case is the direct verification of the hourly economic-dispatch resolution. The MISO-informed case is a higher-resolution sensitivity and should not be interpreted as a measured DTE trajectory.

Table 8: Reactor-level NMPC tracking results for the selected economic fleets using the hourly-linear economic-dispatch reference and the MISO-informed 5-min sensitivity. Tracking error is evaluated only over simulated positive-to-positive online intervals.

Reference	Case	Fleet MAPE (%)	Max. error (MW)	Max. drum rate (deg/s)
Hourly-linear	Coal	0.000037	0.147151	0.004336
	Gas combined cycle	0.000037	0.147246	0.004336
	Gas peaking	0.000389	0.019814	0.002758
	Coal and gas	0.000007	0.001453	0.001166
MISO-informed	Coal	0.000173	0.540229	0.500000
	Gas combined cycle	0.000106	0.048955	0.088465
	Gas peaking	0.002691	0.938246	0.078547
	Coal and gas	0.000014	0.219354	0.013265

All four hourly-linear cases remain well below the 1% tracking criterion and use only a small fraction of the imposed drum-rate limit. Adding the MISO-informed 5-min variation increases both the tracking error and the required drum motion, as expected from the sharper intra-hour changes, but the fleet MAPE remains below 0.003% in every case. The coal sensitivity reaches the imposed 0.5 degrees/s drum-rate constraint, while the other three cases remain below it. The result therefore does not show that the MISO-informed reference is a more accurate representation of DTE operation; rather, it shows that the favorable tracking result is retained when the same hourly power basis is subjected to substantially greater intra-hour variation.

Figure 7 shows the coal case as a representative comparison. The hourly-linear reference follows the resolution of the economic dispatch, while the MISO-informed case contains the additional 5-min structure. The remaining case plots and the MISO matching diagnostics are reported in Appendix D.

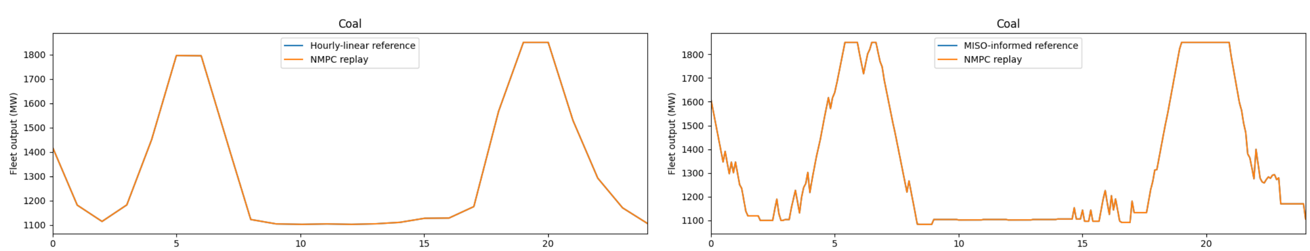

(a) Hourly-linear reference. (b) MISO-informed 5-min sensitivity.

Figure 7: Representative coal-case worst-day NMPC tracking. The hourly-linear case is constructed directly from the hourly economic dispatch, while the MISO-informed case preserves the same hourly power basis and adds higher-resolution within-hour variation. The MISO-informed trajectory is used as a load-following sensitivity and not as a reconstruction of measured DTE operation.

6. Discussion

This section interprets the economic and reactor-level results together. The discussion first considers the break-even capital thresholds and the roles of fossil-generation cost, cycling cost, fleet utilization, and reactor design. The optimization and security-sharing results are then considered in terms of their effect on the selected fleets. The reactor-level results are used to assess whether the corresponding economic dispatch can be followed within the imposed control-drum limits. The main assumptions and limitations of both analysis levels are discussed last.

The screening and the design-specific enumeration both select Aalo-1 exclusively. Table 5 shows the same ordering in every case: Aalo-1 has the highest positive break-even threshold, Holos-Quad is positive for coal, gas peaking, and the combined case, and eVinci is negative throughout. Under the modeled cost inputs, Aalo-1 has the lower FOAK capital basis and lower fixed O&M and fuel costs. Holos-Quad has a slightly lower NOAK capital endpoint, but that advantage is not enough to offset its higher non-capital costs and separate learning sequence. Adding a Holos-Quad or eVinci unit therefore does not provide a sufficient gain in displaced value to overcome the corresponding capital and non-capital costs in the mixed-fleet screening. The eVinci result is worth isolating: its fixed and fuel-plus-variable cost floor already exceeds the midpoint fossil cost, so even a zero capital multiplier does not reach break even. This is a consequence of the selected cost archetype, study power rating, dispatch order, and DTE profiles, and it should not be interpreted as a general comparison of reactor vendors.

The four cases show that a high replacement fraction does not translate into a high capital threshold. The gas combined-cycle fleet replaces 96.73% of its target energy yet yields the lowest positive multiplier ($m^* = 0.103$), because combined-cycle gas has the lowest midpoint fossil cost and therefore the smallest margin above the microreactor non-capital floor. The combined coal-and-gas fleet replaces 29.32% of its target and yields the highest multiplier ($m^* = 0.866$), because it runs at a 100% capacity factor and, by the descending-cost credit order of Section 3.2.3, displaces the most expensive fossil layers first. The gas-peaking case sits between them: it faces the highest fossil cost target ($210/MWh) but its intermittent load caps utilization at 46.15%, so the capital charge is spread over little energy. The controlling variables are therefore the cost of the displaced source and the utilization the fleet can achieve against the load shape, not the fraction of the source replaced. This is the same behavior seen in the replaced-load panels of Figure 4, where the near-baseload cases fill a flat block under the load while the peaking case is stranded above an intermittent profile.

The inverse OCC is a threshold, not a prediction. The $6,484/kW Aalo-1 NOAK coal value means the fleet would have to reach approximately this OCC or lower for its levelized replacement cost to equal the average cost of the coal it displaces; it does not imply that Aalo-1 will reach this value. The Lazard comparison of Figure 5 places this on the same $/MWh axis as the wider generation stack. As emphasized by Eq. (16), the break-even LCOE is a consistency relation rather than a second independent output: it equals the displaced fossil cost by construction and only serves to place the result on the Lazard axis. Negative thresholds carry information too; they indicate that capital reduction alone is insufficient under the modeled fuel and O&M assumptions, so the non-capital cost structure would also have to change.

The generic NEI comparison provides context for the magnitude of the calculated break-even OCC values [69]. The Aalo-1 and Holos-Quad coal and combined coal-and-gas thresholds fall within the reported generic first-unit and mature-unit ranges. The remaining cases are below the lower bounds of the corresponding ranges or have negative break-even OCC values. This overlap does not indicate that either design will achieve the calculated cost; it only shows that some of the required capital-cost levels are consistent with a previously published generic microreactor estimate.

The Holos-specific comparison provides a more direct example of how the inverse-cost thresholds can be interpreted. The inflation-adjusted FOAK and NOAK estimates from the Army G-4 study are both substantially below the rating-adjusted Holos-Quad cost inputs used in the present model. This suggests that the MOUSE-based GCMR proxy used here represents a more conservative capital-cost basis for Holos-Quad than the earlier Holos-specific estimate.

A lower cost than the reference input does not, however, automatically mean that every replacement case reaches break even. Under the non-capital cost, dispatch, and fossil-cost assumptions used here, the adjusted FOAK value of 12,926 $/kW remains above the 12,043 $/kW coal threshold but below the 17,013 $/kW combined-case threshold. The adjusted NOAK value of 7,541 $/kW remains above the 7,149 $/kW combined threshold and 5,061 $/kW coal threshold. The comparison remains approximate because the Army study and the present analysis do not use identical cost boundaries, deployment assumptions, or non-capital inputs.

The optimizer comparison shows that several methods identify the same selected Aalo-1 composition, but not every method does so under the common 1,000-evaluation budget. DE, GWO, HHO, PESA2, and WOA find the selected fleet in all four replacement cases, while EDEV does so in three. Across the ten methods, the selected composition is found by six algorithms for coal, seven for gas combined cycle, six for gas peaking, and five for the combined coal-and-gas case. The pure-design enumeration is then used as a follow-up check of the Aalo-1 unit count after the mixed-fleet search returns a pure-design solution. This provides a separate check on the selected counts without implying that the complete three-design mixed-fleet space was exhaustively enumerated.

The selected fleets contain tens to hundreds of units. These are not near-term deployment recommendations; they show the number of small units required to replace a meaningful part of a utility fossil profile at the selected ratings. A commercial project at this scale would require a more detailed treatment of sites, outages, refueling, shared balance of plant, staffing, licensing, construction schedule, and supply-chain capacity. The units-per-site sensitivity illustrates one multi-unit benefit: sharing only the embedded security account raises the NOAK threshold by about 1,017–1,270 $/kW between $K = 1$ and $K = 40$, with about 81–82% of that gain already captured at five units per site. This should not be extrapolated to every fixed cost because different accounts scale per site, module, reactor, or operating crew.

The dynamic layer is applied after the economic fleet selection and does not change m^*. The selected day and physical-unit endpoints come from the same annual dispatch. Because all four selected fleets contain Aalo-1 units, the reactor-level calculation uses the same reduced-order Aalo/LTMR branch for every case. The primary hourly-linear screen preserves the temporal resolution of the annual economic dispatch. A separate MISO-informed sensitivity adds 5-min within-hour variation while keeping the hourly fleet-power basis fixed, and the controller is tuned independently for the two reference bases.

The hourly-linear and higher-resolution results give the same qualitative conclusion. All four hourly-linear cases track with very small error and remain well below the 0.5 degrees/s drum-rate constraint. The MISO-informed trajectories are more demanding: the largest fleet MAPE is 0.002691% in the gas-peaking case, and the coal case reaches the 0.5 degrees/s rate limit. The higher-resolution calculation is not a validation against measured DTE sub-hourly operation because the MISO profiles are used only as temporal shape surrogates. Its purpose is narrower: it shows that the favorable load-following result is not limited to the smooth hourly interpolation and is retained when additional 5-min variability is imposed.

Several *limitations* bound our analysis and results. The annual dispatch has no binary commitment variable, minimum up- or down-time, reserve requirement, refueling schedule, forced outage, or transmission constraint. Startup and shutdown transitions carry an economic cycling charge, but the zero/nonzero hourly boundaries are still algebraic commitment decisions rather than simulated reactor startup or shutdown sequences. A zero hourly endpoint therefore should not be interpreted as evidence that routine shutdown and restart are physically feasible on that schedule. Cost inputs are public archetype estimates rather than vendor quotations, adjusted to 2026 prices without account-specific labor, equipment, material, or fuel escalation indexes, and the FOAK-to-NOAK path is an interpolation between two endpoint estimates rather than an observed commercial learning history. The generic NEI comparison is also subject to differences in dollar year, plant scope, reactor maturity, production sequence, and included cost accounts; overlap between the ranges is therefore contextual rather than a like-for-like validation. On the optimization side, the ten metaheuristics share a common evaluation budget but not identical internal search behavior. The selected Aalo-1 unit count is checked independently by exact integer enumeration after the mixed-fleet search returns a pure-design solution, while the full three-design mixed-fleet space remains screened heuristically.

The reactor-level analysis carries its own bounds. Its parameter sets are reduced-order comparison models rather than vendor-validated safety models. The selected-day replay evaluates online positive-to-positive intervals and does not model startup, shutdown, protection-system action, or accident behavior. The public DTE/CAMPD data used in the economic model are hourly, so the hourly-linear reference cannot reconstruct observed intra-hour DTE operation. The MISO-informed sensitivity partly addresses this temporal-resolution limitation by imposing 5-min variability while preserving the hourly economic-dispatch power basis, but the selected MISO profiles are not measurements of the corresponding DTE units and should not be interpreted as a validation dataset. The screen also evaluates the most variable 24-hour day rather than every hour of the year, so the resulting test remains an online load-following capability check rather than proof of full annual operating feasibility. Finally, the study computes cost targets and does not propose a specific method for reducing microreactor OCC or O&M; the break-even thresholds indicate what must be achieved, not how.

7. Conclusions and Future Work

This study developed an inverse break-even cost analysis for microreactor fleets replacing fossil generation in the DTE Electric system. The economic model calculates the capital multiplier at which the levelized microreactor replacement cost equals the cost of the fossil generation actually displaced. Startup and shutdown transitions are included as a generic nuclear cycling charge, while the dispatch remains an hourly merit-order model.

Using midpoint 2026 fossil-generation costs, Aalo-1 gives the highest positive NOAK threshold in all four cases. The calculated values are 6,484 \$/kW for coal, 1,063 \$/kW for combined-cycle gas, 2,414 \$/kW for gas peaking, and 8,918 \$/kW for the combined coal-and-gas profile. The selected fleets contain 185, 111, 35, and 124 Aalo-1 units, respectively. Holos-Quad remains positive for coal, gas peaking, and the combined case, while the eVinci non-capital floor is above the displaced fossil cost in every midpoint case.

The result depends strongly on fossil cost and utilization. The combined case supports the highest threshold because its selected fleet operates almost continuously and displaces the most expensive fossil layers first. The combined-cycle case supports the lowest positive threshold even though most of its annual target energy is replaced.

The Aalo-1 and Holos-Quad coal and combined-case thresholds fall within the corresponding generic first-unit and mature-unit microreactor cost ranges considered in the external comparison. This overlap provides context for the magnitude of the calculated thresholds but does not establish that the reactor designs will achieve those costs.

The NEORL mixed-fleet screen and exact single-design enumeration identify the same best Aalo-1 counts. The optimizer comparison also shows that some methods are more consistent than others under the same call limit.

The reactor-level workflow constructs candidate-specific selected-day references from each economic fleet. The hourly-linear replay provides the direct check of the hourly economic-dispatch resolution, and all four cases remain well below the 1% tracking criterion and the 0.5 degrees/s drum-rate limit. A separate MISO-informed 5-min sensitivity introduces additional intra-hour variation while preserving the hourly economic-dispatch power basis. Tracking remains below 0.003% fleet MAPE in all four sensitivity cases, although the coal case reaches the imposed drum-rate limit. The MISO-informed calculation is therefore interpreted as a higher-resolution load-following capability test rather than as a reconstruction of measured DTE operation.

The controller effort itself does not enter the economic objective. Startup and shutdown transitions are already represented through the generic cycling charge of Eq. (11), but the economic model does not include design-specific startup trajectories, minimum up- and down-times, or other unit-commitment constraints. Future work can replace the generic event charge with design-specific cycling costs and explicit commitment constraints for replacement cases that contain frequent zero-to-positive and positive-to-zero transitions.

Finally, the framework is not specific to DTE. It requires an hourly generation profile by fossil source, a cost for each source, and a set of reactor cost archetypes. Applying it to systems with a higher displaced cost, such as island or remote grids served by diesel, would test whether the negative floors reported here are a property of the reactor archetypes or of the comparatively low displaced cost of the DTE fossil fleet.

Data Availability

The authors have all the datasets and codes to reproduce all the results in this work currently in a private GitHub repository. To ensure the confidentiality of this research, the authors will make this repository public during an advanced stage of the review process, which will be listed under our research group's public GitHub page: https://github.com/aims-umich.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors acknowledge the developers of the NEORL optimization framework and the public data and technical-report resources provided by the U.S. Environmental Protection Agency, Idaho National Laboratory, and Lazard. M. M. Sarmiento acknowledges the University of the Philippines System Office of International Linkages for support through the Continuous Operational and Outcomes-based Partnership for Excellence in Research (COOPERATE) program, which funded his exchange research visit to the University of Michigan.

Appendix A. Reactor-level model parameters

The parameters in this appendix are the numerical values used in the reactor-level model. The Holos-Quad branch follows the adopted reduced-order load-following basis [10, 43]. The Aalo-1 branch uses the public 30 $\mathrm{MW_{th}}$/10 $\mathrm{MW_e}$ rating and a MARVEL/MOUSE LTMR dynamic proxy [45, 46, 6, 71]. The eVinci branch uses the 5 $\mathrm{MW_e}$/15 $\mathrm{MW_{th}}$ HPMR/gHPMR proxy lineage [48, 29, 66, 49]. These parameter sets are used for comparative load-following analysis and should not be interpreted as vendor licensing models.

For the Aalo/LTMR and eVinci/HPMR branches, Σ_f and V_f are defined on matching fuel bases. The Aalo value is a one-group thermal fuel-material estimate based on U–ZrH fuel containing 30 wt% uranium at 19.75% enrichment and a fuel density of 6.872 g/cm^3; the corresponding LTMR fuel inventory gives $V_f = 1.68003 \times 10^5$ cm^3 [46, 6]. The eVinci-like value is homogenized over the tristructural isotropic (TRISO) fuel compact using the uranium oxycarbide (UCO) content of the compact at 19.75% enrichment, and its matching compact volume is 1.14299×10^6 cm^3 [48, 29, 49]. Both calculations use a 585.1-b thermal ^{235}U fission cross section as the one-group

Table A.1: Neutronic, poison, and control-drum parameters used in the reduced-order reactor models.

Parameter	Symbol or unit	Holos-Quad	Aalo-1	eVinci
Rated electrical power	P_e (MW$_e$)	13	10	5
Reference thermal power	P_r (MW$_{th}$)	22	30	15
Full-bank drum count	$N_{d,\mathrm{tot}}$	8	12	12
Nominal drum angle	θ_0 (deg)	77.8	84.04	84.5
Full-bank maximum worth	$\rho_{\max,\mathrm{tot}}$	0.047	0.13727	0.14574
Prompt-neutron generation time	Λ (s)	1.68×10^{-3}	1.799×10^{-5}	1.815×10^{-4}
Total delayed-neutron fraction	β	0.00480153	0.00753000	0.00664849
Macroscopic fission cross section	Σ_f (cm^{-1})	0.3358	0.610397	0.148697
Fission-region normalization volume	V_f (cm^3)	8.000×10^4	1.68003×10^5	1.14299×10^6
Xenon absorption cross section	σ_X (cm^2)		2.65×10^{-22}	
Iodine yield	γ_I		0.061	
Direct xenon yield	γ_X		0.002	
Iodine decay constant	λ_I (s^{-1})		2.87×10^{-5}	
Xenon decay constant	λ_X (s^{-1})		2.09×10^{-5}	

modeling value and neglect the thermal ^{238}U fission contribution [72]. The Holos-Quad value and 80,000 cm^3 normalization are retained as the effective basis of the adopted reduced-order model rather than interpreted as a geometric full-core volume.

Table A.2: Six-group delayed-neutron parameters used in the reactor models. Vector entries correspond to groups 1 through 6.

Parameter	Holos-Quad	Aalo-1	eVinci
$\boldsymbol{\beta}_i$	1.42481×10^{-4}, 9.24281×10^{-4}, 7.79956×10^{-4}, 2.06583×10^{-3}, 6.71175×10^{-4}, 2.17806×10^{-4}	2.40×10^{-4}, 1.31×10^{-3}, 1.31×10^{-3}, 2.90×10^{-3}, 1.25×10^{-3}, 5.20×10^{-4}	2.3238×10^{-4}, 1.1992×10^{-3}, 1.1459×10^{-3}, 2.5692×10^{-3}, 1.0591×10^{-3}, 4.4271×10^{-4}
$\boldsymbol{\lambda}_i$ (s^{-1})	1.272×10^{-2}, 3.174×10^{-2}, 1.160×10^{-1}, 3.110×10^{-1}, 1.400, 3.870	1.334×10^{-2}, 3.270×10^{-2}, 1.2086×10^{-1}, 3.0365×10^{-1}, 8.5312×10^{-1}, 2.86517	1.3337×10^{-2}, 3.2734×10^{-2}, 1.2079×10^{-1}, 3.0291×10^{-1}, 8.5005×10^{-1}, 2.8549

Table A.3: Effective lumped thermal parameters used in the three-temperature reactor models.

Parameter	Unit	Holos-Quad	Aalo-1	eVinci
Fuel specific heat, c_f	J/(kg K)	977	374.238	1147.979
Moderator specific heat, c_m	J/(kg K)	1697	659.255	1905.802
Coolant/heat-pipe-node specific heat, c_c	J/(kg K)	5188.6	876.153	1001
Effective fuel mass, m_f	kg	2002	1154.541	2678.072
Effective moderator mass, m_m	kg	11573	316	14001.010
Effective coolant/heat-pipe-node mass, m_c	kg	500	48.671	368.569
Fuel heat fraction, f_f	–	0.96	0.97	0.97
Reference fuel temperature, T_{f0}	K	1105	872	1235
Reference moderator temperature, T_{m0}	K	1087	857	1200
Reference coolant/heat-pipe temperature, T_{c0}	K	985	748	1100
Coolant inlet / heat-sink temperature	K	864	703	523.15
Coolant outlet temperature	K	1106	793	–
Coolant mass flow rate, $\dot{m}$	kg/s	17.5	380.451	0
Fuel feedback coefficient, α_f	K^{-1}	-2.875×10^{-5}	-5.28×10^{-5}	-4.27×10^{-5}
Moderator/solid feedback coefficient, α_m	K^{-1}	-3.696×10^{-5}	$+1.22 \times 10^{-5}$	-1.42×10^{-5}
Coolant feedback coefficient, α_c	K^{-1}	0	$+1.60 \times 10^{-6}$	0

For the LTMR proxy used to represent the Aalo-1 branch and the HPMR proxy used to represent the eVinci branch, Σ_f and V_f are defined on matching fuel bases. The LTMR proxy uses a one-group thermal fuel-material estimate based on U–ZrH fuel, while the HPMR proxy uses the corresponding homogenized TRISO fuel basis. The sodium–potassium (NaK) heat capacity is evaluated at the 748-K coolant-node temperature using the INL NaK correlation [73]. The Aalo moderator/solid coefficient is an effective nonfuel-solid feedback mapped onto the lumped moderator node. For the eVinci/HPMR branch, T_{c0} = 1100 K is the heat-pipe coupling node and 523.15 K is the condenser-side sink, so a conventional coolant outlet and bulk coolant flow are not used [49, 29]. The inter-node conductances are calculated internally from the listed full-power states and are not independent inputs.

Appendix B. Microreactor cost inputs

Table B.1 lists the six cost components used in the learning model. The learning (L) components are interpolated between the first-of-a-kind (FOAK) and nth-of-a-kind (NOAK) endpoints using Eq. (4); the non-learning (NL)

components are held constant with build number. Capital is represented through the learning and non-learning portions of OCC. Fixed O&M contains both learning and non-learning portions, while the variable-cost basis used here has a zero non-learning component. All values are in 2026 dollars at the modeled ratings. The design capacity factors are 0.9878 for Aalo-1, 0.9890 for Holos-Quad, and 0.9849 for eVinci.

Table B.1: Microreactor learning and non-learning cost inputs in 2026 dollars. Values are escalated from the 2024 INL bottom-up basis [6, 8]. L and NL denote learning and non-learning components, respectively.

Design	Basis	OCC_L ($/kW)	OCC_{NL}	FOM_L ($/kW-yr)	FOM_{NL}	Var._L ($/MWh)	Var._{NL}
Aalo-1	FOAK	16,834.9	5,542.0	756.1	127.4	25.80	0
Aalo-1	NOAK	4,757.8	5,542.0	192.8	127.4	23.81	0
Holos-Quad	FOAK	18,849.0	4,255.8	1,086.6	96.3	47.10	0
Holos-Quad	NOAK	5,452.8	4,255.8	257.4	96.3	24.58	0
eVinci	FOAK	38,325.4	9,087.4	1,902.5	254.8	293.78	0
eVinci	NOAK	10,978.2	9,087.4	434.3	254.8	156.68	0

The two OCC columns sum to the total OCC used in Eq. (5); for example, the Aalo-1 FOAK and NOAK totals are 22,377 and 10,300 $/kW.

Appendix C. NEORL screening summary

This appendix summarizes the mixed-fleet NEORL screening results used in Section 5.5. Each of the ten optimization methods is run for ten seeds using the same case-specific fleet bounds and a common budget of 1,000 fleet evaluations per run. The mean and standard deviation are calculated from the best m^* returned by each of the ten seed runs. The "Best m^*" column gives the largest value found by that optimizer across the ten seeds, together with the corresponding Aalo-1, Holos-Quad, and eVinci unit counts. The replacement fraction and fleet capacity factor correspond to that best fleet.

The NEORL calculation searches the full three-design fleet space. The pure-design enumeration is not used to generate the values in this appendix; it is applied only after the mixed-fleet screening returns a pure Aalo-1 selected fleet, as a follow-up check of the selected Aalo-1 unit count. Identical fleet counts and m^* values found by different algorithms or seeds represent the same economic solution.

Table C.1: NEORL screening summary for the coal replacement case. Each optimizer was evaluated over ten seeds with a common budget of 1,000 fleet evaluations per run.

Algorithm	Mean m^*	Std. m^*	Best m^*	N_A	N_H	N_E	Replaced (%)	CF (%)
DE	0.629132	0.000357	0.629555	185	0	0	63.83	98.66
EDEV	0.629219	0.000501	0.629555	185	0	0	63.83	98.66
GWO	0.618724	0.034250	0.629555	185	0	0	63.83	98.66
HHO	0.569131	0.066379	0.629555	185	0	0	63.83	98.66
PESA	0.381017	0.052594	0.466588	160	14	5	62.44	98.82
PESA2	0.622829	0.021270	0.629555	185	0	0	63.83	98.66
PSO	0.242192	0.143095	0.502393	133	18	0	54.44	99.55
SA	0.146284	0.118306	0.314687	230	19	35	86.67	91.06
TS	0.438035	0.140037	0.629260	179	0	0	61.89	98.88
WOA	0.597063	0.052318	0.629555	185	0	0	63.83	98.66

Appendix D. Higher-resolution reference construction and tracking diagnostics

The MISO-informed calculation is included only as a higher-resolution sensitivity to the hourly economic dispatch. The MISO profile supplies normalized within-hour variation, while each hourly economic-dispatch fleet value remains the power basis of the reconstructed reference. For the single-source cases, the selected resource minimizes

Table C.2: NEORL screening summary for the coal and gas replacement case. Each optimizer was evaluated over ten seeds with a common budget of 1,000 fleet evaluations per run.

Algorithm	Mean m^*	Std. m^*	Best m^*	N_A	N_H	N_E	Replaced (%)	CF (%)
DE	0.859696	0.009333	0.865866	124	0	0	29.32	100.00
EDEV	0.861490	0.005685	0.865838	126	0	0	29.79	100.00
GWO	0.865866	0.000000	0.865866	124	0	0	29.32	100.00
HHO	0.839961	0.054612	0.865866	124	0	0	29.32	100.00
PESA	0.482707	0.078952	0.584980	4	172	11	54.80	99.42
PESA2	0.856911	0.022796	0.865866	124	0	0	29.32	100.00
PSO	0.262328	0.164721	0.663495	135	1	10	33.40	99.98
SA	0.181328	0.150522	0.542360	170	149	3	81.27	94.11
TS	0.532032	0.188377	0.848749	180	0	0	42.51	99.88
WOA	0.852913	0.040959	0.865866	124	0	0	29.32	100.00

Table C.3: NEORL screening summary for the gas combined cycle replacement case. Each optimizer was evaluated over ten seeds with a common budget of 1,000 fleet evaluations per run.

Algorithm	Mean m^*	Std. m^*	Best m^*	N_A	N_H	N_E	Replaced (%)	CF (%)
DE	0.102981	0.000559	0.103166	111	0	0	96.73	84.58
EDEV	0.101613	0.003142	0.103166	111	0	0	96.73	84.58
GWO	0.103166	0.000000	0.103166	111	0	0	96.73	84.58
HHO	0.088252	0.047163	0.103166	111	0	0	96.73	84.58
PESA	0.009300	0.029297	0.038888	112	3	2	99.84	82.89
PESA2	0.103166	0.000000	0.103166	111	0	0	96.73	84.58
PSO	-0.110346	0.057602	-0.002021	121	3	5	100.00	76.18
SA	-0.117219	0.072536	-0.032910	79	27	0	98.88	84.10
TS	0.033278	0.067911	0.103166	111	0	0	96.73	84.58
WOA	0.103166	0.000000	0.103166	111	0	0	96.73	84.58

Table C.4: NEORL screening summary for the gas peaking replacement case. Each optimizer was evaluated over ten seeds with a common budget of 1,000 fleet evaluations per run.

Algorithm	Mean m^*	Std. m^*	Best m^*	N_A	N_H	N_E	Replaced (%)	CF (%)
DE	0.233892	0.001060	0.234395	35	0	0	40.49	46.15
EDEV	0.232127	0.002907	0.234395	35	0	0	40.49	46.15
GWO	0.234395	0.000000	0.234395	35	0	0	40.49	46.15
HHO	0.180733	0.086405	0.234395	35	0	0	40.49	46.15
PESA	0.030829	0.057788	0.108786	42	4	1	51.32	42.92
PESA2	0.234395	0.000000	0.234395	35	0	0	40.49	46.15
PSO	-0.093685	0.092672	0.083698	66	9	0	72.17	37.05
SA	-0.170599	0.038001	-0.121157	133	6	34	97.89	24.75
TS	0.035571	0.140021	0.197130	36	1	0	42.62	45.59
WOA	0.216507	0.056565	0.234395	35	0	0	40.49	46.15

the normalized 24-h shape root mean square error (RMSE) among complete full-day candidates. For the combined case, the active source layers on the selected day are matched separately and combined by daily energy share.

Table D.1: MISO shape-matching diagnostics used to construct the higher-resolution 5-min sensitivity reference.

Case	Candidates	Shape RMSE	Shape correlation	Selected MISO unit(s)
Coal	172	0.061132	0.984294	A3789
Gas combined cycle	167	0.024449	0.998253	A19218
Gas peaking	176	0.092625	0.998929	A3644
Coal and gas	169	0.130053	0.932433	coal:A5739; gas peaking:A5260

Figure D.1 illustrates the construction for the coal case. The 5-min MISO-informed points add within-hour structure around the hourly economic-dispatch power basis, and the NMPC reference is then interpolated at the 60-s controller interval.

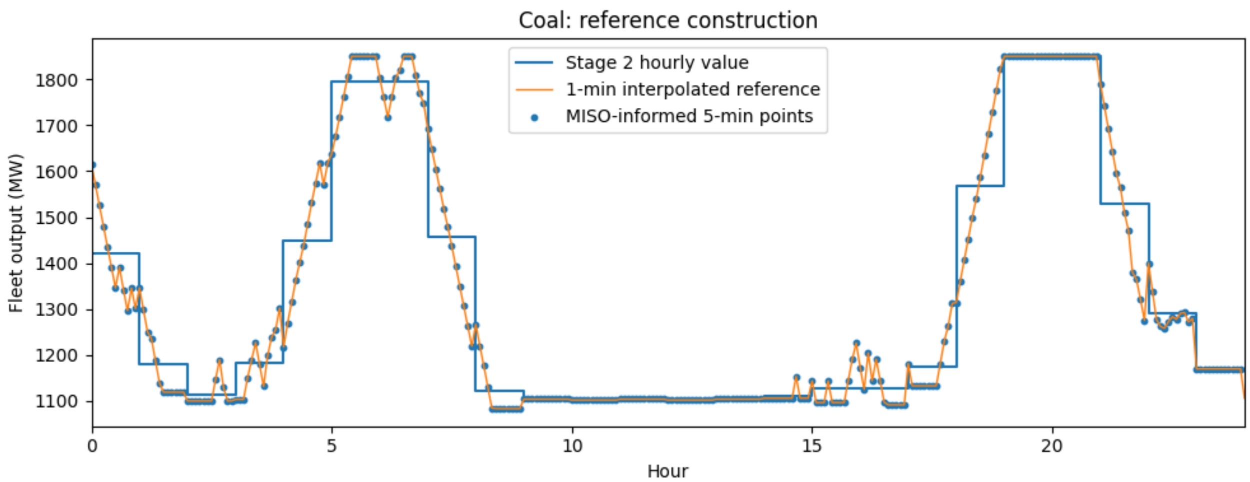

Figure D.1: Coal-case construction of the higher-resolution sensitivity reference. The MISO-informed 5-min points are used only to impose within-hour variation while the hourly economic-dispatch fleet values remain the reference power basis.

References

[1] U.S. Energy Information Administration, Fossil generation could rise with faster-than-expected growth in data center power demand, Today in Energy, accessed 2026-07-21 (Mar. 2026).
URL https://www.eia.gov/todayinenergy/detail.php?id=67344

[2] DTE Energy, Michigan public service commission approves dte's landmark clean energy plan, accessed 2026-04-02 (2023).
URL https://www.dteenergy.com/us/en/newsroom/2023/Michigan-Public-Service-Commission-approves-DTE-
html

[3] G. Black, D. Shropshire, K. Araújo, A. Heek, Prospects for nuclear microreactors: A review of the technology, economics, and regulatory considerations, Nuclear Technology 209 (2022) 1–20. doi:10.1080/00295450.2022.2118626.

[4] E. Morones-García, J.-L. François, Nuclear microreactors: Status review and potential applications; the mexican case, Nuclear Engineering and Design 420 (2024) 113021. doi:https://doi.org/10.1016/j.nucengdes.2024.113021.
URL https://www.sciencedirect.com/science/article/pii/S0029549324001213

[5] B. N. Hanna, R. G. Gonzaga de Oliveira, K. A. Mohammad Al Dawood, M. D. Patterson, A. Abou Jaoude, S. Garcia, B. Lindley, Technoeconomic evaluation of microreactor using detailed bottom-up estimate (rev.1), Tech. rep., Idaho National Laboratory (INL), Idaho Falls, ID (United States) (11 2024). doi:10.2172/2447366.
URL https://www.osti.gov/biblio/2447366

[6] B. Hanna, K. Al-Dawood, P. Seurin, R. Abdelnasser, A. Abou-Jaoude, A bottom-up cost estimation tool for nuclear microreactors, Tech. Rep. INL/RPT-25-87273-Rev000, Idaho National Laboratory, Idaho Falls, Idaho, United States, oSTI ID: 2588465 (Sep. 2025).
URL https://www.osti.gov/biblio/2588465

[7] A. Abou-Jaoude, L. Lin, C. Bolisetti, E. K. Worsham, L. M. Larsen, A. S. Epiney, Literature review of advanced reactor cost estimates, Tech. rep., Idaho National Laboratory (INL), Idaho Falls, ID (United States) (06 2023). doi:10.2172/1986466.
URL https://www.osti.gov/biblio/1986466

[8] IdahoLabResearch, Mouse: Microreactor cost estimation tool, GitHub repository, accessed 2026-02-06 (2025).
URL https://github.com/IdahoLabResearch/MOUSE

[9] J. Hansen, W. D. Jenson, A. Wrobel, K. Biegel, N. Stauff, T. K. Kim, R. Belles, F. Omitaomu, Investigating benefits and challenges of converting retiring coal plants into nuclear plants, Tech. Rep. INL/RPT-22-67964, Idaho National

Laboratory (2022). doi:10.2172/1886660.
URL https://www.osti.gov/biblio/1886660

[10] K. K. Abdulraheem, S. Choi, Q. Shen, B. Kochunas, M. I. Radaideh, A load following reactivity control system for nuclear microreactors, Progress in Nuclear Energy 184 (2025) 105676. doi:https://doi.org/10.1016/j.pnucene.2025.105676.
URL https://www.sciencedirect.com/science/article/pii/S0149197025000745

[11] L. Tunkle, K. Abdulraheem, L. Lin, M. I. Radaideh, Nuclear microreactor transient and load-following control with deep reinforcement learning, Energy Conversion and Management: X 27 (2025) 101090. doi:https://doi.org/10.1016/j.ecmx.2025.101090.
URL https://www.sciencedirect.com/science/article/pii/S2590174525002223

[12] DTE Energy Company, DTE Electric Company, Annual report on Form 10-K for the fiscal year ended december 31, 2025, SEC Filing Commission File Nos. 1-11607 and 1-2198, U.S. Securities and Exchange Commission, Washington, DC, accession No. 0000936340-26-000054; filed February 17, 2026 (Feb. 2026).
URL https://www.sec.gov/Archives/edgar/data/936340/000093634026000054/

[13] S. Cho, C. Li, I. E. Grossmann, Recent advances and challenges in optimization models for expansion planning of power systems and reliability optimization, Computers & Chemical Engineering 165 (2022) 107924. doi:10.1016/j.compchemeng.2022.107924.

[14] M. Carrión, J. M. Arroyo, A computationally efficient mixed-integer linear formulation for the thermal unit commitment problem, IEEE Transactions on Power Systems 21 (3) (2006) 1371–1378. doi:10.1109/TPWRS.2006.876672.

[15] Q. P. Zheng, J. Wang, A. L. Liu, Stochastic optimization for unit commitment—a review, IEEE Transactions on Power Systems 30 (4) (2015) 1913–1924. doi:10.1109/TPWRS.2014.2355204.

[16] L. A. Burnett, U. M. Nabila, M. I. Radaideh, Variational digital twins, Energy and AI (2026) 100756.

[17] U. M. Nabila, P. Seurin, L. Lin, M. I. Radaideh, Physics-informed digital twin development of thermal energy distribution systems with active learning, Applied Energy 416 (2026) 127903.

[18] U. M. Nabila, P. Seurin, L. Lin, M. I. Radaideh, Active learning for uncertainty quantification of a physics-informed digital twin of a thermal energy distribution system, in: Proceedings of the International Conference on the Physics of Reactors (PHYSOR 2026), Torino, Italy, 2026. doi:10.5281/zenodo.20804162.

[19] D. Price, M. I. Radaideh, B. Kochunas, Simplified matching pursuits applied to 3d nuclear reactor temperature distribution construction, Applied Mathematical Modelling 131 (2024) 134–158.

[20] S. Jain, G. Shrimali, Impact of renewable electricity on utility finances: Assessing merit order effect for an indian utility, Energy Policy 168 (2022) 113092. doi:10.1016/j.enpol.2022.113092.

[21] M. Solomon, E. Gimon, M. O'Boyle, U. Paliwal, A. Phadke, Coal cost crossover 3.0: Local renewables plus storage create new opportunities for customer savings and community reinvestment, Tech. rep., Energy Innovation Policy & Technology LLC, San Francisco, California (Jan. 2023).
URL https://energyinnovation.org/report/the-coal-cost-crossover-3-0/

[22] S. Sene, L. Lin, J. Kim, M. I. Radaideh, Surrogate-driven variance-based sensitivity analysis of thermal storage tanks in integrated energy systems, in: Proceedings of the International Topical Meeting on Nuclear Plant Instrumentation, Control, and Human-Machine Interface Technologies (NPIC&HMIT 2025), Chicago, IL, USA, 2025, pp. 1446–1456.

[23] M. I. Radaideh, B. Forget, K. Shirvan, Large-scale design optimisation of boiling water reactor bundles with neuroevolution, Annals of Nuclear Energy 160 (2021) 108355.

[24] M. I. Radaideh, I. Wolverton, J. Joseph, J. J. Tusar, U. Otgonbaatar, N. Roy, B. Forget, K. Shirvan, Physics-informed reinforcement learning optimization of nuclear assembly design, Nuclear Engineering and Design 372 (2021) 110966.

[25] D. Price, M. I. Radaideh, B. Kochunas, Multiobjective optimization of nuclear microreactor reactivity control system operation with swarm and evolutionary algorithms, Nuclear Engineering and Design 393 (2022) 111776.

[26] W. Li, G. Zhang, X. Yang, Z. Tao, H. Xu, Sizing a hybrid renewable energy system by a coevolutionary multiobjective optimization algorithm, Complexity 2021 (2021) 8822765. doi:10.1155/2021/8822765.

[27] A. Ş. Uyar, B. Türkay, A. Keleş, A novel differential evolution application to short-term electrical power generation scheduling, International Journal of Electrical Power & Energy Systems 33 (6) (2011) 1236–1242. doi:10.1016/j.ijepes.2011.01.036.

[28] D. Price, M. Radaideh, N. Roskoff, B. Kochunas, Serpent-openfoam coupling of an evinci motivated microreactor design for burnup-dependent critical control drum positions, in: International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, CNS, 2023.

[29] D. Price, N. Roskoff, M. I. Radaideh, B. Kochunas, Multiphysics modeling of heat pipe microreactor with critical control drum position search, Nuclear Science and Engineering (2024) 1–20.

[30] G. Chicco, A. Mazza, Metaheuristic optimization of power and energy systems: Underlying principles and main issues of the "rush to heuristics", Energies 13 (19) (2020) 5097. doi:10.3390/en13195097.

[31] B. Li, J. Thomas, A. R. de Queiroz, J. F. DeCarolis, Open source energy system modeling using break-even costs to inform state-level policy: A north carolina case study, Environmental Science & Technology 54 (2) (2020) 665–676. doi:10.1021/acs.est.9b04184.

[32] P. Denholm, The value of energy storage for grid applications, Tech. Rep. NREL/TP-6A20-58465, National Renewable Energy Laboratory, Golden, Colorado (2013). doi:10.2172/1079719.
URL https://www.nrel.gov/docs/fy13osti/58465.pdf

[33] N. A. Sepulveda, J. D. Jenkins, A. Edington, D. S. Mallapragada, R. K. Lester, The design space for long-duration energy storage in decarbonized power systems, Nature Energy 6 (5) (2021) 506–516. doi:10.1038/s41560-021-00796-8.

[34] M. S. Ziegler, J. M. Mueller, G. D. Pereira, J. Song, M. Ferrara, Y.-M. Chiang, J. E. Trancik, Storage requirements and costs of shaping renewable energy toward grid decarbonization, Joule 3 (9) (2019) 2134–2153. doi:10.1016/j.joule.2019.06.012.

[35] N. A. Sepulveda, J. D. Jenkins, F. J. de Sisternes, R. K. Lester, The role of firm low-carbon electricity resources in deep decarbonization of power generation, Joule 2 (11) (2018) 2403–2420. doi:10.1016/j.joule.2018.08.006.

[36] G. Alonso, S. Bilbao, E. del Valle, Economic competitiveness of small modular reactors versus coal and combined cycle plants, Energy 116 (1) (2016) 867–879. doi:10.1016/j.energy.2016.10.030.

[37] D. Price, M. Radaideh, B. Kochunas, Active learning with deep neural network ensembles for heat pipe microreactor modeling in abnormal operation, in: Conference: 18th International Probabilistic Safety Assessment and Analysis. PSA, 2023, p. 2023.

[38] L. Tunkle, K. Abdulraheem, M. Amans, M. I. Radaideh, Load-following control in microreactors with deep reinforcement learning, in: Proc. International Conference on Mathematics & Computational Methods Applied to Nuclear Science and Engineering (M&C 2025). Denver, Colorado, 2025.

[39] L. Tunkle, K. Abdulraheem, L. Lin, M. I. Radaideh, Enhancing autonomous control of microreactors using multi-agent reinforcement learning, in: Proceedings of the International Topical Meeting on Nuclear Plant Instrumentation, Control, and Human-Machine Interface Technologies (NPIC&HMIT 2025), Chicago, IL, USA, 2025, pp. 40–49. doi:10.13182/NPICHMIT25-46807.

[40] M. I. Radaideh, L. Tunkle, D. Price, K. Abdulraheem, L. Lin, M. Elias, Multistep criticality search and power shaping in nuclear microreactors with deep reinforcement learning, Nuclear Science and Engineering (2025) 1–13.

[41] U.S. Environmental Protection Agency, Clean air markets program data (campd), hourly generation data; accessed July 2026 (2026).
URL https://campd.epa.gov/

[42] Lazard, Lazard's levelized cost of energy+, version 19.0, Tech. rep., Lazard (July 2026).

[43] N. Stauff, C. H. Lee, C. Filippone, Core design of the holos-quad microreactor, Tech. rep., Argonne National Laboratory (ANL), Argonne, IL (United States) (03 2022). doi:10.2172/1863405.
URL https://www.osti.gov/biblio/1863405

[44] HolosGen, Technology, accessed 2026-01-05 (2026).
URL https://www.holosgen.com/technology/

[45] Aalo Atomics, Aalo pod: 50 MWe nuclear power plant for AI data centers, https://www.aalo.com/technology, accessed: 2026-07-27 (2026).

[46] Idaho National Laboratory, MARVEL 90% final design report, Tech. Rep. INL/RPT-23-74280, Revision 0, Idaho National Laboratory, Idaho Falls, Idaho, United States (sep 2023).

[47] Westinghouse Electric Company, evinci microreactor, Company technology page, accessed 2026-01-06 (2026).
URL https://westinghousenuclear.com/innovation/evinci-microreactor/

[48] D. Price, N. Roskoff, M. I. Radaideh, B. Kochunas, Thermal modeling of an evinci™-like heat pipe microreactor using openfoam, Nuclear Engineering and Design 415 (2023) 112709. doi:https://doi.org/10.1016/j.nucengdes.2023.112709.
URL https://www.sciencedirect.com/science/article/pii/S0029549323005587

[49] J. Ortensi, M. K. Jaradat, J. Hansel, S. Terlizzi, The monolithic heat pipe microreactor reference plant model, Tech. Rep. INL/RPT-24-77914, Idaho National Laboratory (apr 2024).
URL https://virtualtestbed.inl.gov/microreactors/gHPMR/gHPMR_reactor_description.html

[50] E. S. Rubin, I. M. Azevedo, P. Jaramillo, S. Yeh, A review of learning rates for electricity supply technologies, Energy Policy 86 (2015) 198–218. doi:https://doi.org/10.1016/j.enpol.2015.06.011.
URL https://www.sciencedirect.com/science/article/pii/S0301421515002293

[51] N. E. Stauff, W. N. Mann, K. Biegel, T. Levin, J. D. Rader, A. Cuadra, S. H. Kim, Reactor power size impacts on nuclear competitiveness in a carbon-constrained future, Tech. Rep. ANL/NSE-21/55, Argonne National Laboratory (sep 2021). doi:10.2172/1865640.
URL https://www.osti.gov/biblio/1865640

[52] U.S. Bureau of Labor Statistics, Consumer price index historical tables for u.s. city average, consumer Price Index for All Urban Consumers (CPI-U), all items, U.S. city average.
URL https://www.bls.gov/regions/mid-atlantic/data/consumerpriceindexhistorical_us_table.htm

[53] M. I. Radaideh, K. Du, P. Seurin, D. Seyler, X. Gu, H. Wang, K. Shirvan, Neorl: Neuroevolution optimization with reinforcement learning, arXiv preprint arXiv:2112.07057.

[54] M. I. Radaideh, K. Du, P. Seurin, D. Seyler, X. Gu, H. Wang, K. Shirvan, Neorl: Neuroevolution optimization with reinforcement learning—applications to carbon-free energy systems, Nuclear Engineering and Design 412 (2023) 112423. doi:https://doi.org/10.1016/j.nucengdes.2023.112423.
URL https://www.sciencedirect.com/science/article/pii/S0029549323002728

[55] R. Storn, K. Price, Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces, Journal of Global Optimization 11 (1997) 341–359. doi:10.1023/A:1008202821328.

[56] J. Kennedy, R. Eberhart, Particle swarm optimization, in: Proceedings of ICNN'95 - International Conference on Neural Networks, Vol. 4, 1995, pp. 1942–1948 vol.4. doi:10.1109/ICNN.1995.488968.

[57] S. Mirjalili, S. M. Mirjalili, A. Lewis, Grey wolf optimizer, Advances in Engineering Software 69 (2014) 46–61. doi:https://doi.org/10.1016/j.advengsoft.2013.12.007.
URL https://www.sciencedirect.com/science/article/pii/S0965997813001853

[58] S. Mirjalili, A. Lewis, The whale optimization algorithm, Advances in Engineering Software 95 (2016) 51–67. doi:https://doi.org/10.1016/j.advengsoft.2016.01.008.
URL https://www.sciencedirect.com/science/article/pii/S0965997816300163

[59] A. A. Heidari, S. Mirjalili, H. Faris, I. Aljarah, M. Mafarja, H. Chen, Harris hawks optimization: Algorithm and applications, Future Generation Computer Systems 97 (2019) 849–872. doi:https://doi.org/10.1016/j.future.2019.02.028.
URL https://www.sciencedirect.com/science/article/pii/S0167739X18313530

[60] S. Kirkpatrick, C. Gelatt, M. Vecchi, Optimization by simulated annealing, Science (New York, N.Y.) 220 (1983) 671–80. doi:10.1126/science.220.4598.671.

[61] F. Glover, M. Laguna, Tabu search i, ORSA Journal on Computing 1. doi:10.1287/ijoc.1.3.190.

[62] M. Radaideh, K. Shirvan, Pesa: Prioritized experience replay for parallel hybrid evolutionary and swarm algorithms - application to nuclear fuel, Nuclear Engineering and Technology 54. doi:10.1016/j.net.2022.05.001.

[63] G. Wu, X. Shen, H. Li, H. Chen, A. Lin, P. N. Suganthan, Ensemble of differential evolution variants, Information Sciences 423 (2018) 172–186. doi:10.1016/j.ins.2017.09.053.

[64] B. Kochunas, K. Barr, S. Kinast, S. Choi, Global and local reactivity assessments for passive control systems of multi-module htgr special purpose reactors, Tech. Rep. NE/8887/2020-003-00NURAM-2020-005-00, University of Michigan (Sep. 2020).

[65] J. J. Duderstadt, L. J. Hamilton, Nuclear Reactor Analysis, Wiley, 1976.

[66] N. Choi, M. K. Mohammad Jaradat, V. Laboure, J. A. Christensen, S. Terlizzi, Advancements in multiphysics microdepletion analysis of an evinci^{tm}-like microreactor leveraging openmc-crab workflow, Tech. rep., Idaho National Laboratory (INL), Idaho Falls, ID (United States) (03 2024). doi:10.2172/2341461.
URL https://www.osti.gov/biblio/2341461

[67] D. Q. Mayne, Model predictive control: Recent developments and future promise, Automatica 50 (12) (2014) 2967–2986. doi:https://doi.org/10.1016/j.automatica.2014.10.128.
URL https://www.sciencedirect.com/science/article/pii/S0005109814005160

[68] Midcontinent Independent System Operator, Market reports: Real-time cleared offers, real-Time Cleared Offers market-report data; accessed 2026-08-13 (2026).
URL https://www.misoenergy.org/markets-and-operations/real-time--market-data/market-reports/

[69] Nuclear Energy Institute, Cost competitiveness of micro-reactors for remote markets, Report, Nuclear Energy Institute (Apr. 2019).
URL https://www.nei.org/resources/reports-briefs/cost-competitiveness-micro-reactors-remote-market

[70] J. A. Vitali, J. G. Lamothe, J. Toomey, Charles J., V. O. Peoples, K. A. McCabe, Study on the use of mobile nuclear power plants for ground operations, Tech. rep., U.S. Army, appendix C, Table C.5 (October 2018).
URL https://www.holosgen.com/wp-content/uploads/2019/03/G4-MNPP-Study.pdf

[71] J. A. Evans, R. T. Sweet, P. G. Medvedev, A. R. Wagner, C. Parisi, T. L. Lange, E. Perez, F. Rice, J.-F. Jue, E. Woolstenhulme, D. D. Keiser, Y. Arafat, Uranium-zirconium hydride nuclear fuel performance in the nak-cooled marvel microreactor, Journal of Nuclear Materials 598 (2024) 155145. doi:10.1016/j.jnucmat.2024.155145.

[72] Japan Atomic Energy Agency Nuclear Data Center, JENDL-4.0 cross section table for ^{235}u, Nuclear Data Center, Japan Atomic Energy Agency, accessed 2026-07-30; thermal $(n, fission)$ cross section at 0.0253 eV is 585.1 b (2026).
URL https://wwwndc.jaea.go.jp/cgi-bin/Tab80WWW.cgi?lib=J40&iso=U235

[73] Idaho National Laboratory, Nakfluidproperties, MOOSE Fluid Properties documentation, accessed 2026-08-09 (2024).
URL https://mooseframework.inl.gov/releases/moose/2024-11-11/source/fluidproperties/NaKFluidProperties.html