Comparative Study of Scalarization and Pareto-Based Search for iSMR Core Loading Pattern Optimization

Byeong-hyeok Ha [a*], Tongkyu Park [a], Sung-kyun Zee [b]
[a]*FNC Technology, Heungdeok IT Valley Bldg. 20F, 13, Heungdeok 1-ro, Giheung-gu, Yongin-si, 16954, Korea*
[b]*FNC Technology, Daeduk Tech Business Center, 1004-2, 593, Daedukdae-ro, Yusung-gu, Daejeon-si, 34112, Korea*
[*]*Corresponding author: haqualux@fnctech.com*

Keywords: multi-objective optimization, scalarization, Pareto front, simulated annealing, loading pattern optimization

1. Introduction

Small modular reactors (SMRs) are investigated more frequently as flexible, safety-enhanced alternatives to conventional large-scale nuclear power plants. One of the promising SMRs, the innovative SMR (iSMR) has a vast and highly combinatorial design space due to extra constraints on fuel assembly (FA) arrangement introduced by its core safety and performance requirements. Core loading pattern design is inherently a multi-objective optimization (MOO) problem: the arrangement of fuel assemblies within the core must simultaneously satisfy competing requirements for power distribution safety margins and neutronic performance. In particular, the peak power peaking factor (Fr) must be minimized to maintain thermal-hydraulic safety margins, while the effective multiplication factor (keff) must be maximized to ensure sufficient reactivity for the operating cycle. Because these two objectives do not, in general, share a single optimal loading pattern, an effective optimization framework must be capable of exploring the trade-off relationship between them rather than collapsing the problem into a single criterion by assumption.

There are several algorithmic approaches to MOO problems, and they differ fundamentally in when and how they incorporate a decision-maker's preference between objectives. Scalarization-based approaches combine multiple objectives into a single weighted criterion prior to the search, effectively reducing the multi-objective problem to a sequence of single-objective searches. Pareto-front-based approaches, in contrast, operate directly on the notion of dominance between candidate solutions and attempt to construct an approximation of the full trade-off surface without committing to a fixed preference in advance. This study applies multi-objective simulated annealing (MOSA) to the iSMR core loading pattern optimization problem employing several distinct search strategies — scalarization under multiple normalization schemes and weight ratios, and Pareto-front archiving under multiple acceptance rules — and compares their relative performance against an exhaustively determined reference Pareto front.

2. Methods

2.1. Problem Definition

The optimization problem addressed in this study is the loading pattern design of an iSMR core, which requires fuel assemblies of five different types to be assigned to eleven core positions. Each of the eleven positions independently accepts one of the five FA types, so the full combinatorial design space consists of 5^{11} (48,828,125) unique loading patterns. This space is completely discrete and unstructured with respect to gradient information, making it an ideal candidate for metaheuristic search methods such as simulated annealing.

Two objective functions are considered. The first is the peak power peaking factor, Fr, which characterizes the maximum local power density relative to the core-average power density; lower Fr values correspond to larger thermal-hydraulic safety margins and are thus preferable (minimization objective). The second is the effective multiplication factor, keff, which characterizes the core's reactivity balance; higher values of keff are preferable, subject to practical operational constraints (maximization objective). Because these two objectives are not typically optimized simultaneously by the same loading pattern, the problem is treated as a genuine two-objective optimization task rather than reduced to a single criterion.

2.2. AI Surrogate Model

Evaluating Fr and keff for a candidate loading pattern using conventional core physics analysis codes requires a significant amount of computational time per evaluation, which makes exhaustive or near-exhaustive evaluation of the full 48,828,125-pattern design space computationally impractical using conventional methods. To overcome this limitation, we employed a previously trained AI surrogate model that predicts Fr and keff directly from the loading pattern configuration, lowering the cost per evaluation by several orders of magnitude compared to direct physics-based analysis.

Using this surrogate model, we evaluated all 48,828,125 possible loading patterns and applied non-

dominated sorting to the resulting (Fr, keff) pairs to construct a reference Pareto front representing the genuine optimal trade-off surface for this problem, as predicted by the surrogate model. The reference front contains 50 non-dominated solutions. This reference front serves as the ground-truth benchmark against which all subsequent search-based results are evaluated in Section 3.

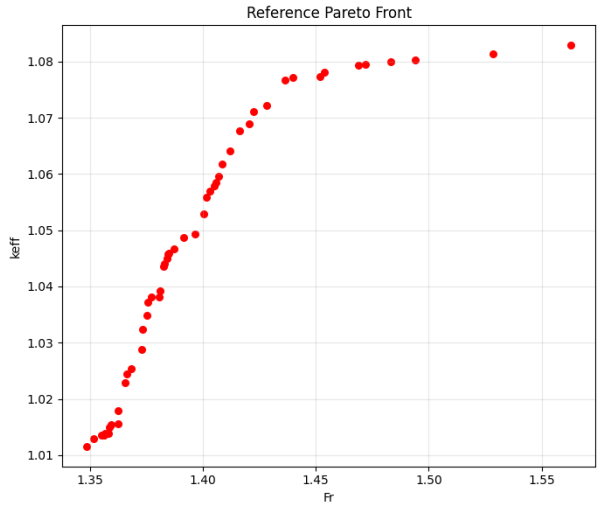


Fig. 1. Reference Pareto front obtained from exhaustive surrogate-model evaluation

2.3. Multi-Objective Simulated Annealing (MOSA)

Simulated annealing (SA) was used as the underlying single-solution metaheuristic for all search strategies in this study. At each iteration, a neighboring loading pattern was generated by randomly selecting one of the eleven core positions and replacing the FA type at that position with a different, randomly selected FA type from the remaining four types. This neighbor-generation rule was applied identically across all search configurations to ensure that differences in performance could be attributed to the search strategy rather than variations in the underlying move operator.

Two categories of multi-objective search strategy were implemented on top of this common SA framework.

(1) Scalarization-based search.

The two objectives were combined into a single scalar objective using a weighted sum after normalization. Two normalization schemes were used.

Min-Max normalization rescales each objective to the [0, 1] range:

$$Fr' = \frac{Fr - Fr_{min}}{Fr_{max} - Fr_{min}} \quad keff' = \frac{keff - keff_{min}}{keff_{max} - keff_{min}}$$

Z-score standardization rescales each objective using its mean (μ) and standard deviation (σ):

$$Fr' = \frac{Fr - \mu_{Fr}}{\sigma_{Fr}} \quad keff' = \frac{keff - \mu_{keff}}{\sigma_{keff}}$$

Fr_{min}, Fr_{max}, μ_{Fr}, σ_{Fr}, $keff_{min}$, $keff_{max}$, μ_{keff} and σ_{keff} were fixed in advance, computed from a reference dataset rather than updated online during each SA run, so that the normalization basis remained identical across all iterations and across all runs of a given configuration.

The normalized objectives were then combined into a single scalarized objective to be minimized:

$$E = w_{Fr} \cdot Fr' + w_{keff} \cdot (-keff')$$

For each normalization scheme, three weight ratios (w_{Fr}:w_{keff}) were tested — 1:1, 10:1, and 1:10 — yielding six scalarization configurations in total.

(2) Pareto-front archiving search.

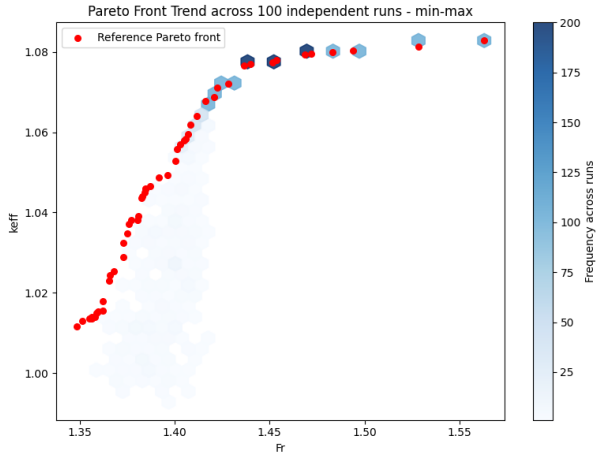
Rather than combining the objectives, this strategy evaluates dominance directly and keeps track of non-dominated solutions found during the search. Each objective was given its own independent temperature parameter, and an acceptance probability for each objective was computed separately at each iteration using its own temperature schedule. The two per-objective acceptance probabilities were then aggregated into a single acceptance decision using one of three rules: the minimum of the two probabilities, the maximum of the two probabilities, or their mean. Using the minimum results in a comparatively low aggregate acceptance probability, which favors convergence and exploitation; using the maximum yields a comparatively high aggregate acceptance probability, which favors broader exploration at the cost of additional search effort; the mean rule falls between these two extremes.

For all nine configurations (six scalarization variants and three Pareto-archiving variants), the maximum number of solution evaluations per SA run was capped at 50,000. Runs that satisfied the cooling-schedule termination criteria before reaching this cap terminated earlier; runs that did not converge under the cooling schedule were terminated upon reaching the 50,000-evaluation cap.

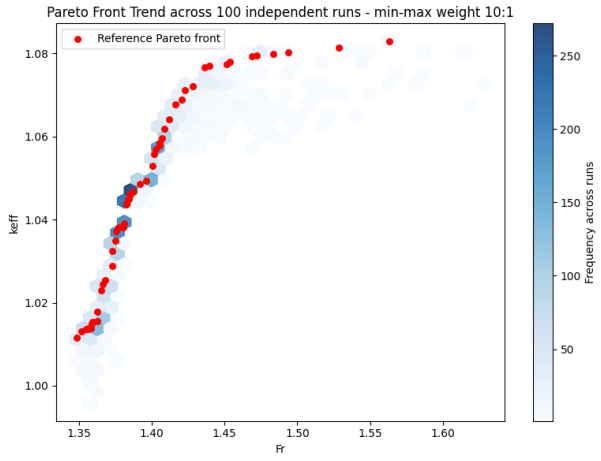
3. Analyses & Results

Each of the nine search configurations was executed independently 100 times. For each run, we kept track of the total number of solutions evaluated, the total number of temperature stages completed, and the number of discovered non-dominated solutions that perfectly matched a solution in the reference Pareto front. To

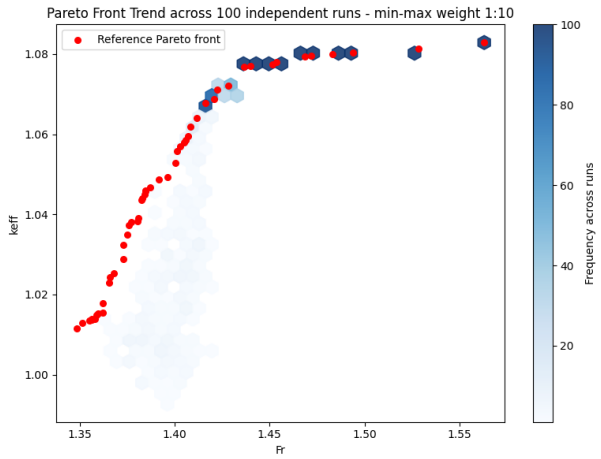
determine where each method succeeded or failed within the trade-off surface, the reference front was divided into two equally sized regions by rank: an Fr-dominant region (the 25 solutions with the smallest Fr) and a keff-dominant region (the 25 solutions with the largest keff). Match counts were computed separately within each region (Fr-dominant region and keff-dominant region). Table I provides summary statistics (mean ± standard deviation) for all 100runs.



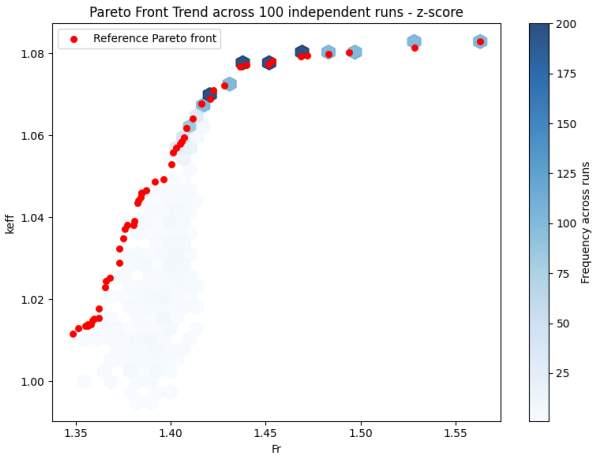
(a) Min-Max 1:1



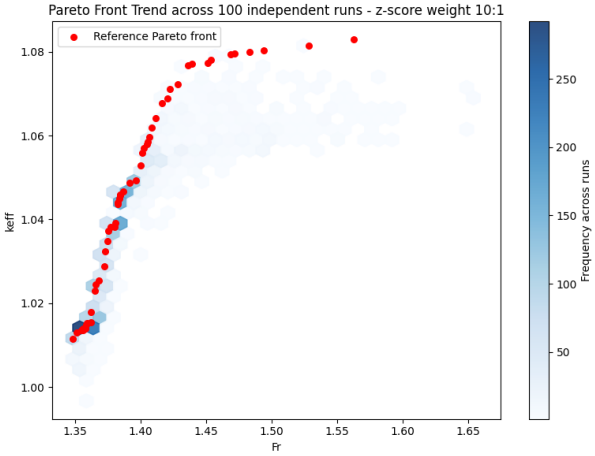
(b) Min-Max 10:1



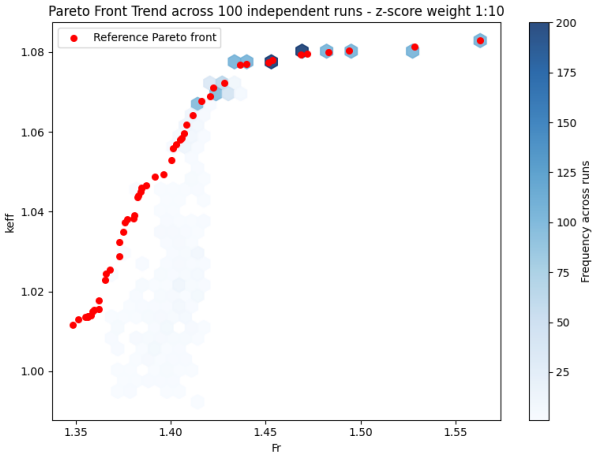
(c) Min-Max 1:10



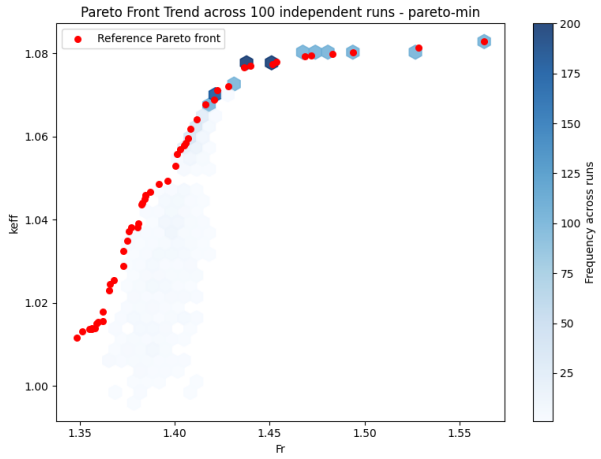
(d) Z-score 1:1



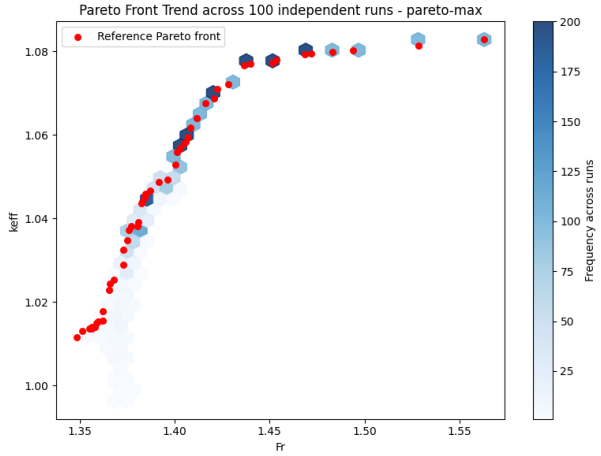
(e) Z-score 10:1



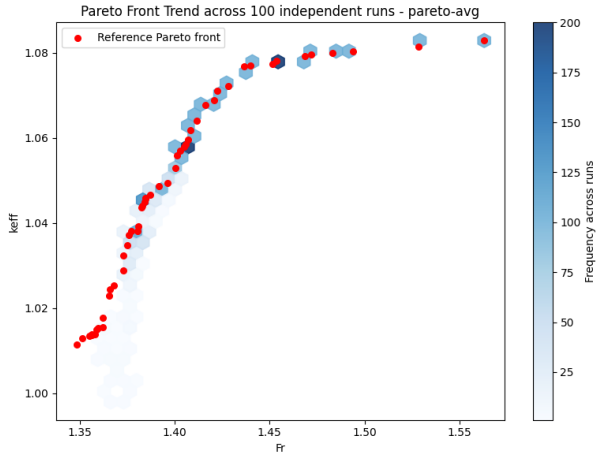
(f) Z-score 1:10



(g) Pareto-min



(h) Pareto-max



(i) Pareto-avg

Fig. 2. Hexbin density plots of the non-dominated solutions recovered across the 100 runs for each method, overlaid on the reference Pareto front

Table I. Search performance across nine MOSA configurations (mean ± standard deviation over 100 independent runs). The reference Pareto front (50 solutions) is split into an Fr-dominant region (25 solutions with the smallest Fr) and a keff-dominant region (25 solutions with the largest keff).

	Total Stages	Total Samples	Overall Match	Fr-dominant	keff-dominant
Min-Max 1:1	69.79 ± 8.96	19,051.07 ± 4,482.99	14.47 ± 0.70	0.00 ± 0.00	14.47 ± 0.70
Min-Max 10:1	73.11 ± 7.48	17,539.07 ± 3,720.22	19.53 ± 4.73	11.50 ± 3.45	8.03 ± 4.85
Min-Max 1:10	50.87 ± 1.91	10,469.00 ± 869.81	12.13 ± 1.10	0.00 ± 0.00	12.13 ± 1.10
Z-score 1:1	59.68 ± 3.71	14,147.46 ± 1,755.49	14.79 ± 0.96	0.02 ± 0.14	14.77 ± 0.94
Z-score 10:1	69.77 ± 6.52	15,684.55 ± 3,105.32	15.08 ± 6.24	12.73 ± 5.72	2.35 ± 2.12
Z-score 1:10	51.43 ± 1.77	10,708.96 ± 843.16	12.42 ± 1.01	0.01 ± 0.10	12.41 ± 1.02
Pareto-min	54.33 ± 5.74	12,867.02 ± 2,880.84	14.13 ± 1.00	0.00 ± 0.00	14.13 ± 1.00
Pareto-max	282.58 ± 4.93	50,000 ± 0	26.19 ± 2.05	2.92 ± 1.65	23.27 ± 1.02
Pareto-avg	162.29 ± 2.48	50,000 ± 0	24.93 ± 2.10	1.97 ± 1.62	22.96 ± 1.00

The results show a strong and, in some cases, near-total bias toward the keff-dominant region. Five of the nine configurations — Min-Max 1:1, Z-score 1:1, Min-Max 1:10, Z-score 1:10, and Pareto-min — recovered essentially none of the Fr-dominant solutions (0.00–0.02 out of 25 on average), while recovering 12–15 of the 25 keff-dominant solutions. This near-total failure to reach the Fr-dominant region is consistent with the underlying structure of the design space: across the full 48,828,125-pattern space, Fr and keff have a positive relationship in which improving one objective tends to also improve the other for the majority of candidate loading patterns, so that unguided local search naturally drifts toward the densely populated keff-favoring portion of the space. A genuine trade-off emerges only near the narrow boundary corresponding to the true Pareto front; this boundary represents a small, relatively difficult-to-reach subset of the entire search landscape.

Assigning a 10:1 weight ratio to Fr substantially changed this picture for both normalization schemes. Min-Max 10:1 recovered 11.50/25 (46%) of the Fr-dominant region, and Z-score 10:1 recovered 12.73/25 (51%) — representing roughly half of the Fr-dominant solutions. This improvement, however, came at a sharp cost to keff-dominant coverage, which fell to 8.03/25 (32%) for Min-Max 10:1 and to just 2.35/25 (9%) for Z-score 10:1. Rather than broadening overall coverage, extreme weighting effectively traded coverage of one region for the other. This trade-off was accompanied by markedly higher run-to-run variance than the corresponding 1:1 configurations (e.g., standard deviation of 5.72 for Z-score 10:1's Fr-dominant match versus 0.14 for Z-score 1:1). This suggests that, rather than all runs converging to a similar, moderate outcome,

some individual runs under extreme weighting failed nearly entirely in one region while others succeeded.

In line with its relatively conservative, exploitation-favoring behavior, the minimum-acceptance-probability rule among the Pareto-archiving configurations replicated the same near-total keff-dominant bias seen in the equally weighted and keff-favoring scalarization methods (0.00/25 in the Fr-dominant region). Aside from the 10:1-weighted scalarization methods, the maximum- and mean-acceptance-probability rules were the only two configurations that recovered any Fr-dominant solutions without an explicit weighting toward Fr, reaching 2.92/25 (12%) and 1.97/25 (8%) respectively, while maintaining strong keff-dominant coverage (23.27/25 and 22.96/25, or 93% and 92%). This improved regional coverage, however, was obtained at a substantially higher evaluation cost: both configurations reached the 50,000-evaluation cap in every one of the 100 runs, roughly 3–5 times the evaluation count typically consumed by the scalarization methods and Pareto-min, which terminated earlier via the cooling schedule. Whether this modest gain in Fr-dominant coverage reflects an inherent advantage of the max/mean acceptance rule or is mostly due to the larger evaluation budget cannot be entirely separated from the present results and would require a budget-matched comparison to resolve.

4. Conclusions

This study applied multi-objective simulated annealing to the iSMR core loading pattern optimization problem, comparing six scalarization-based search configurations (Min-Max and Z-score normalization with weight ratios of 1:1, 10:1, and 1:10) against three Pareto-front archiving configurations (minimum-, maximum-, and mean-based acceptance probability aggregation). An AI surrogate model was employed to thoroughly evaluate the whole 48,828,125-pattern design space, allowing all nine search configurations to be benchmarked directly against a verified reference Pareto front of 50 solutions, evenly divided into Fr-dominant and keff-dominant regions.

The results show that five of the nine configurations — all equally weighted or keff-favoring scalarization variants, plus the minimum-acceptance Pareto-archiving variant — recovered almost none of the Fr-dominant region, a pattern attributable to a positive global correlation between Fr and keff across the majority of the design space, which directs unguided search toward the keff-favoring majority. Reweighting the scalarization objective toward Fr (10:1) recovered roughly half of the Fr-dominant region but sharply reduced keff-dominant coverage and increased run-to-run variance, indicating a direct trade-off rather than an overall improvement. Pareto-archiving with maximum or mean acceptance was the only approach that retained strong keff-dominant coverage while recovering a modest, non-zero fraction of the Fr-dominant region, but it did so with substantially more solution evaluations than any other configuration, leaving it unclear whether this benefit is due to the acceptance rule itself or simply a larger search budget. Taken together, these findings indicate that no single search configuration examined in this study uniformly dominates the others across both regions of the trade-off surface. Practical core loading pattern optimization may benefit from combining complementary search strategies, adaptively adjusting weighting or acceptance criteria during the search, or conducting a budget-matched comparison to isolate the effect of acceptance-rule choice from that of search length.

Acknowledgements

This work was supported by the Innovative Small Modular Reactor Development Agency grant funded by the Korea Government (MCEE) (No. RS-2023-00259289).

REFERENCES

[1] Kirkpatrick, S., Gelatt Jr, C. D., & Vecchi, M. P, Optimization by simulated annealing. science, 220(4598), 671-680, 1983.
[2] T. Park, H. G. Joo, C. H. Kim, and H. C. Lee, Multiobjective Loading Pattern Optimization by Simulated Annealing Employing Discontinuous Penalty Function and Screening Technique, Nuclear Science and Engineering, vol.162, pp.134, 2009.
[3] J. Kwon, T. Park, and S. K. Zee, Development of CNN-based surrogate models to predict power distributions in i-SMR, Annals of Nuclear Energy, 213, 111108, 2025.