

Multi-Objective Design Optimization of Ventilation Duct Systems: A Graph-Informed Hybrid Evolutionary Approach

Xiangming Liu, Bin Liu, Kunze Du, Da Gao, Nan Li*

School of Computer Science and Artificial Intelligence, Beijing Technology and Business University, Beijing, China

Abstract—Optimizing silencer placement in Heating, Ventilation, and Air Conditioning (HVAC) systems is a complex multi-objective problem due to conflicting objectives (noise, energy, cost) and intricate topological constraints. Conventional Multi-Objective Evolutionary Algorithms (MOEAs) often exhibit inefficient convergence on such problems due to their reliance on random search strategies. Addressing this challenging HVAC design problem requires a more informed approach. This paper proposes the G-HNSGA-III (Graph-Informed Hybrid NSGA-III), a novel framework that enhances the NSGA-III algorithm by embedding domain-specific knowledge from the system's Directed Acyclic Graph (DAG) topology. This is achieved through two core components that leverage heuristic search: a Graph-Informed Initialization (GINI) strategy to provide a high-quality starting population and a Graph-Informed Local Search (GILS) module for post-processing refinement. The performance of G-HNSGA-III was comprehensively benchmarked against the baseline NSGA-III and six other established MOEAs on a complex data center test instance. The results demonstrate a marked superiority, with G-HNSGA-III achieving a 38.4% higher mean Hypervolume (HV) than the baseline NSGA-III and a 99.3% Set Coverage (SC) dominance over MOEA/D. The framework consistently converged to the best-known Pareto front, achieving a final mean Inverted Generational Distance (IGD) of 0.0030. These findings validate that the proposed graph-informed strategies effectively accelerate convergence and enable the discovery of a higher-quality Pareto front, providing superior and practically applicable solutions for complex engineering design problems.

Keywords—*Multi-objective optimization; NSGA-III; graph-informed optimization; HVAC design; heuristic search; domain knowledge*

I. INTRODUCTION

Modern facilities with strict climate control requirements—spanning from commercial complexes to critical infrastructure like hospitals—rely heavily on HVAC systems for operational stability [1-2]. While essential, these systems create a fundamental operational conflict. On one hand, the mechanical noise from fans degrades the acoustic environment; on the other, the fan power required to drive airflow constitutes a major operational cost, directly opposing sustainability targets [3-4]. The trade-off is physical and unforgiving: installing silencers to attenuate noise inherently increases airflow resistance (pressure drop), which forces fans to consume more energy to maintain the same flow rate. Therefore, optimizing the placement of these

components is not merely a design preference but a central challenge in balancing energy efficiency with environmental quality.

This optimization task is compounded by the topological complexity of HVAC ductworks. Unlike simple linear systems, duct networks function as Directed Acyclic Graphs (DAGs), where the airflow dynamics are highly interdependent [5-6]. A design decision at a single upstream node propagates through the entire network, affecting pressure and noise levels at all downstream branches. Consequently, conventional design methodologies reliant on heuristics or trial-and-error fail to capture these global dependencies. They lack the mathematical rigor to navigate a search space characterized by discrete variables (e.g., specific silencer models) and strict physical constraints.

To address such high-dimensional engineering problems, Multi-Objective Evolutionary Algorithms (MOEAs) have become the established methodological standard [7-8]. However, their application to network-based systems faces a critical bottleneck: the quality of the initial population. Standard algorithms like NSGA-III typically rely on random initialization. In the strictly constrained solution space of an HVAC network, this stochastic approach often results in an inefficient search process, leading to slow convergence or stagnation in local optima—a limitation well-known in the field but rarely addressed with domain-specific logic.

We propose that embedding domain-specific topological knowledge into the evolutionary search can overcome these limitations. To this end, we introduce the Graph-Informed Hybrid NSGA-III (G-HNSGA-III) framework. This approach enhances NSGA-III by integrating DAG topology through two mechanisms: a Graph-Informed Initialization (GINI) strategy and a Graph-Informed Local Search (GILS) module.

The principal contributions of this paper are as follows:

- The proposal of a novel hybrid framework, G-HNSGA-III, that effectively embeds the topological knowledge from a DAG model into the evolutionary search process for HVAC optimization.
- The design and implementation of two specific algorithmic components: a GINI strategy to accelerate initial convergence and a GILS module to enhance the final solution set quality.

*Corresponding author.

- A comprehensive validation of the framework on a complex, real-world data center case study, demonstrating its significant performance superiority over the baseline NSGA-III and six other established MOEAs.

The remainder of this paper is organized as follows. Related work is given in Section II. Section III details the DAG-based system modeling and the pre-computation of topological metrics. Section IV presents the proposed G-HNSGA-III framework in detail. Section V describes the experimental setup, parameters, and evaluation metrics. Section VI presents and analyzes the comparative results. Finally, Section VII concludes the paper and discusses future research directions.

II. RELATED WORK

The dual challenge of acoustic control and energy efficiency is well-established in HVAC literature. Noise generation is an inevitable byproduct of mechanical ventilation, necessitating robust attenuation strategies to meet environmental standards [9-10]. Approaches to characterize and mitigate this noise vary widely, from detection algorithms based on recursive partitioning in office spaces [11] to structural analysis of duct vibration problems [12]. Parallel to acoustic concerns, energy optimization remains a priority. Research has extensively focused on improving fan operating efficiency [3], [13] and integrating low-noise mechanical ventilators [14-15]. These studies collectively illustrate that optimizing an HVAC system is never a single-variable problem; it is a complex balancing act between silencing the system and powering it.

To navigate these conflicting objectives, Multi-Objective Evolutionary Algorithms (MOEAs) serve as the dominant computational tool. Comprehensive reviews confirm their efficacy in handling the non-linear constraints typical of building engineering [7-8]. Specific applications have ranged from optimizing pressure regulators [16] to automating design workflows based on BIM data [17]. However, the "black-box" nature of standard MOEAs presents limitations. Without domain knowledge, these algorithms struggle with initialization. As noted in recent studies, random initialization in algorithms like NSGA-III heavily influences the final hypervolume convergence [18], and while adaptive penalty schemes can mitigate some stagnation issues [19], they do not fundamentally resolve the lack of structural awareness.

Recent advancements suggest that hybridizing evolutionary search with domain-specific models offers a way forward. For instance, combining MOEAs with LSTM networks has proven effective in dynamic routing problems [20-21]. More importantly, the value of topological information is gaining recognition. Graph-based representations, such as DAGs and Convolutional Networks, are increasingly used to model spatial-temporal dependencies in complex systems [22-23]. Unlike conventional approaches that treat HVAC optimization as a black-box problem, this paper synthesizes these advancements by explicitly embedding the DAG topological structure into the evolutionary framework. This distinct positioning allows the proposed G-HNSGA-III to overcome the initialization blindness and premature convergence flaws inherent in standard MOEAs.

III. DAG MODEL

A. System Modeling with Directed Acyclic Graphs

A Directed Acyclic Graph (DAG) is employed to model the ventilation system, as its structure inherently captures the unidirectional nature of airflow and acoustic propagation. The system is formally represented as $G = (V, E)$, where V is the set of nodes and E is the set of directed edges.

Each node $t_i \in V$ corresponds to a unique physical component, including air sources (e.g., fans), transport elements (e.g., ducts, elbows), functional units (e.g., silencers), and terminal points (e.g., rooms). The edge set $E = e_{ij}$ represents the direct physical connections, where an edge e_{ij} signifies that both airflow and acoustic energy propagate from node t_i to t_j . The graph's topology and the physical properties of its connections are encoded in a weighted adjacency matrix, A [23].

B. Formal Expression and Matrix Construction for Nodes and Edges

Each node $v_i \in V$ in the DAG corresponds to a discrete physical component of the ventilation system (e.g., a fan, duct segment, or silencer). The quantifiable properties of each component are encapsulated in a multi-dimensional attribute vector, F_i . This vector is a concatenation of several attribute subset, which are based on established engineering principles [24]:

Location ($L_i \in \mathbb{R}^3$): The 3D spatial coordinates (x_i, y_i, z_i) .

Geometric ($G_i \in \mathbb{R}^k$): Component dimensions, such as duct length, width, and height $(L_{duct,i}, W_{duct,i}, H_{duct,i})$ or elbow parameters $(W_{elbow,k}, H_{elbow,k}, r_{elbow,k})$.

Fluid Dynamic ($V_i \in \mathbb{R}^{(k_V)}$): Parameters governing airflow, including surface roughness (ε_i), local loss coefficients (K_i), and fan performance curve coefficients.

Acoustic ($C_i \in \mathbb{R}^{(k_C)}$): Parameters governing sound, including natural attenuation $A_{nat,i}(f)$, flow-generated noise $L_{W,SGN,i}(f)$, and silencer insertion loss $IL_i(f)$.

Economic ($E_i \in \mathbb{R}^{(k_E)}$): Cost-related data, such as the procurement cost ($Cost_{initial,i}$).

The directed edges, $e_{ij} \in E$, formally represent the unidirectional connections between nodes t_i and t_j , defining the path for airflow and acoustic energy. The acyclic property of the graph is critical, as it mirrors the physical unidirectionality of the ventilation system. This structure ensures a partial ordering of nodes, which eliminates circular dependencies and guarantees the sequential convergence of network-based simulations.

Weighted Adjacency Matrix (A): This matrix encodes the topology of the system, defining the existence and properties (e.g., length) of the connections between any two nodes.

Component Attribute Matrix (D_{comp}): This matrix serves as a central repository for the intrinsic features of all available component types. Each row corresponds to a specific component model, while the columns parameterize its key attributes:

- Component ID
- Reference location vector
- Geometric and fluid dynamic parameter vector
- Acoustic and economic parameter vector

During optimization, a specific component from D_{comp} is assigned to a node in the graph. The process of translating a physical system layout into this formal DAG representation is shown in Fig. 2.

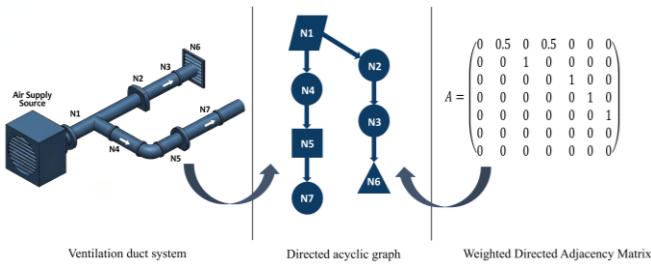


Fig. 1. Illustration of directed acyclic graph generation.

For example, consider a simple system composed of a duct segment, an elbow, and another duct segment. The attributes of these three components would be stored in the component attribute matrix, D_{comp} , as shown in Eq. (1).

$$D_{comp} = \begin{bmatrix} num & l & g & a \\ 1 & [0, 0, 0] & [2000, 1000, 4.9] & 11000 \\ 2 & [0, 10, 0] & [2000, 1] & 11000 \\ 3 & [0, 0, 10] & [2000, 1000, 1.5] & 11000 \end{bmatrix} \quad (1)$$

Each row in D_{comp} represents a unique component in the system. The first row, f_1 , corresponds to the first duct segment and contains its complete attribute vector. This vector begins with the component's unique identifier, 1. It is followed by the location vector $L_1 = [0, 0, 0]$, specifying its 3D spatial coordinates. The subsequent vector, $G_1/V_1 = [2000, 1000, 4.9]$, encapsulates its geometric and fluid dynamic properties, representing, for instance, a width of 2000 mm, a height of 1000 mm, and a length of 4.9 m. The final vector (omitted here for brevity) would contain the relevant acoustic and economic parameters, completing the row's specification for a single component.

A weighted directed adjacency matrix, $A = (a_{ij})$, is used to encode both the system's topology and the proportional distribution of acoustic energy. This $N \times N$ matrix is defined as shown in Eq. (2):

$$a_{ij} = \begin{cases} \omega_{ij}, & \text{if } e_{ij} = \langle t_i, t_j \rangle \in E \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

The weight ω_{ij} represents the fraction of acoustic energy that propagates from node t_i to a direct successor node t_j . At any junction where the duct branches, this fraction is determined by the ratio of the cross-sectional areas, as calculated in Eq. (3):

$$\omega_{ij} = S_j / \sum_{k \in \text{succ}(i)} S_k \quad (3)$$

where S_j is the cross-sectional area of the specific branch leading to node t_j , and the denominator is the sum of the areas of all branches immediately downstream from node t_i .

For example, consider the system illustrated in Fig. 1. Its topology and energy distribution are encoded in the weighted adjacency matrix A , shown in Eq. (4). An element $a_{12} = 0.5$ signifies that the branch from node 1 to node 2 receives 50% of the acoustic energy from node 1; similarly, $a_{14} = 0.5$ indicates an equal energy split to node 4. In contrast, an element $a_{23} = 1$ indicates that 100% of the energy from node 2 propagates to node 3, which is characteristic of a non-branching connection. All zero-valued elements denote the absence of a direct path between the corresponding nodes.

$$A = \begin{pmatrix} 0 & 0.5 & 0 & 0.5 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad (4)$$

C. Pre-computation of Structural Information

To enhance computational efficiency, key structural metrics are pre-computed from the DAG. As these metrics are inherent to the network topology and independent of any specific silencer configuration, they are calculated once and stored prior to initiating the optimization process. A primary metric is node depth, $Depth(t_i)$, which defines a node's hierarchical position relative to the system's source. The root node of the system, t_{source} (e.g., the fan), is assigned a depth of zero: $Depth(t_{source}) = 0$. For any other node t_j , the depth is calculated recursively, as expressed in Eq. (5), as the minimum depth of its immediate predecessors plus one:

$$Depth(t_j) = \min_{t_i \in \text{pred}(t_j)} \{Depth(t_i) + 1\} \quad (5)$$

where, $\text{pred}(t_j)$ is the set of immediate parent nodes of t_j .

Another pre-computed metric is the downstream total airflow, $Q_{downstream}(t_i)$. This value represents the cumulative volumetric flow demand of all terminal outlets reachable from node t_i . It is calculated via a reverse topological traversal, starting from the outlets. The recursive definition is provided in Eq. (6):

$$= \begin{cases} Q_{downstream}(t_i) \\ Q_{(demand,i)}, \text{if } t_i \in V_{outlets} \\ \sum_{(t_j \in \text{succ}(t_i))} Q_{downstream}(t_j), \text{otherwise} \end{cases} \quad (6)$$

where $Q_{(demand,i)}$ is the specified airflow requirement for a terminal outlet node t_i , and $\text{succ}(t_i)$ is the set of immediate successor nodes of t_i .

The potential sound source strength of a node, $L_{W,potential}(t_k)$, estimates the maximum potential noise

contribution of each node t_k . Its value is dependent on the component type and is defined in Eq. (7) as:

$$L_{w,potential}(t_k) = \begin{cases} L_{w,rated}(t_k), & \text{if } t_k \text{ is a fan} \\ L_{w,SGN,max}(t_k), & \text{if } t_k \text{ is a duct component} \end{cases} \quad (7)$$

where $L_{w,rated}(t_k)$ is the fan's sound power level at its rated operating point. The term $L_{w,SGN,max}(t_k)$ is the maximum estimated self-generated noise for an aerodynamic component (e.g., a duct, elbow, or damper), calculated using the empirical formula presented in Eq. (8):

$$L_{w,SGN,max}(t_k) = L_{wc}(t_k) + 50 \log_{10}(V_{d,k}) + 10 \log_{10}(A_k) \quad (8)$$

Where $L_{wc}(t_k)$ is the component-specific sound power level, $V_{d,k}$ is the design velocity through the component, and A_k is its cross-sectional area.

Finally, the previously computed metrics (node depth, downstream airflow, and potential sound source strength) are aggregated into a composite Location Importance score, $I(t_k)$. This heuristic is designed to estimate the strategic value of placing a silencer at any given node t_k , providing critical domain knowledge to guide the optimization algorithm. This score is calculated via Eq. (9) as a normalized, weighted sum of the previously defined metrics:

$$I(t_k) = \alpha \cdot \frac{1}{Depth(t_k) + 1} + \beta \cdot \frac{Q_{downstream}(t_k)}{\max(Q_{downstream}(t_i) \mid t_i \in V)} + \gamma \cdot \frac{L_{w,potential}(t_k)}{\max(L_{w,potential}(t_j) \mid t_j \in V)} \quad (9)$$

where α , β , and γ are weighting coefficients constrained by $\alpha + \beta + \gamma = 1$. Each term in the sum is normalized to a range between 0 and 1 to ensure a balanced contribution.

Collectively, these pre-computed metrics— $Depth(t_i)$, $Q_{downstream}(t_i)$, and $I(t_k)$ —constitute the quantitative domain knowledge that informs the heuristic guidance for the Graph-Informed Initialization (GINI) and Local Search (GILS) strategies within the GINS-III framework.

IV. GRAPH-INFORMED HYBRID OPTIMIZATION FRAMEWORK

The core of this work is the Graph-Informed Hybrid Optimization Framework (GINS-III), designed to solve the multi-objective silencer placement problem. The framework is built upon the Non-dominated Sorting Genetic Algorithm III (NSGA-III) and enhances its performance by incorporating the pre-computed topological metrics described in Section II-C.

This integration is achieved through two primary components: a Graph-Informed Initialization (GINI) strategy and a Graph-Informed Local Search (GILS) mechanism. By systematically leveraging this domain knowledge, the GINS-III

framework guides the evolutionary search process to efficiently balance the conflicting objectives of acoustic attenuation, energy consumption, and economic cost.

A. Multi-Objective Optimization Preliminaries

In multi-objective optimization, a solution X_1 Pareto dominates X_2 (denoted $X_1 < X_2$) if it is superior in at least one objective without being inferior in any other. Formally, this relationship is defined in Eq. (10):

$$X_1 < X_2 \Leftrightarrow (\forall i, f_i(X_1) \leq f_i(X_2)) \wedge (\exists j, f_j(X_1) < f_j(X_2)) \quad (10)$$

where $f_i(X)$ is the value of the i -th objective function. The set of all such non-dominated solutions within the search space X constitutes the Pareto Front (PF), as mathematically represented in Eq. (11).

$$PF = \{X \in X \mid \nexists X' \in X, X' < X\} \quad (11)$$

Multi-Objective Evolutionary Algorithms (MOEAs) aim to find a high-quality approximation of this front. The foundational NSGA-II algorithm uses a crowding distance metric to promote diversity, but its efficacy degrades in many-objective (≥ 3 objectives) problems. To overcome this limitation, NSGA-III replaces the crowding distance calculation with a reference-point-based selection mechanism, which effectively maintains population diversity across high-dimensional trade-off surfaces and allows for superior performance on complex problems. The complete methodological workflow of the proposed approach is visually illustrated in Fig. 2.

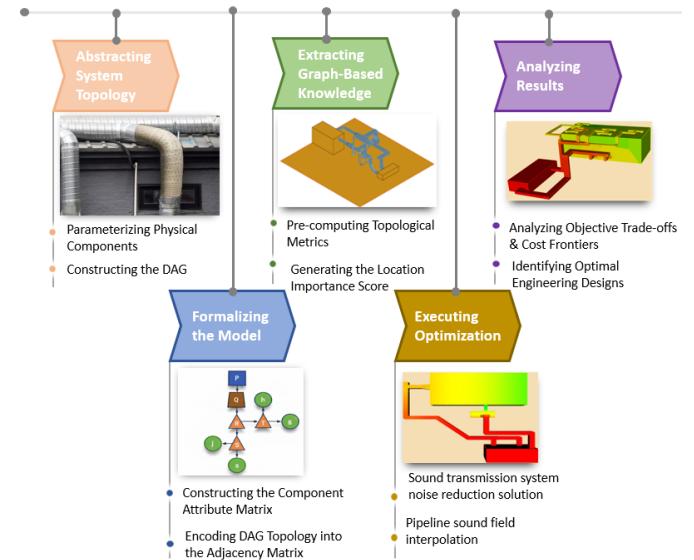


Fig. 2. Methodological workflow of the graph-informed optimization approach.

B. Framework Overview

The G-HNSGA-III framework is a two-phase hybrid architecture. It enhances NSGA-III's global exploration by embedding graph-based domain knowledge through two novel components: a Graph-Informed Initialization (GINI) strategy and a Graph-Informed Local Search (GILS) mechanism.

Phase 1: Global Exploration. This phase employs NSGA-III, but critically replaces its random initialization with our GINI

strategy. GINI leverages pre-computed topological metrics to generate a high-quality guided initialization population, thereby accelerating convergence from the outset.

Phase 2: Local Refinement. The preliminary Pareto front from the global search is then refined by the GILS module. As a post-processing step, GILS performs a targeted neighborhood search to improve solution convergence, operating independently of the main evolutionary loop.

C. Graph-Informed Population Initialization (GINI)

To mitigate the inefficiencies of random initialization in complex engineering problems, we developed the Graph-Informed Initialization (GINI) strategy. GINI replaces a purely random start by embedding pre-computed topological knowledge into the initial population using two complementary mechanisms:

1) *Heuristic seeding*: A fraction of the population (p_{seed}) is heuristically seeded. For individuals in this subset, decision variables corresponding to the most topologically important nodes (as determined by $I(t_k)$) are assigned a high-performance silencer model with a high probability.

2) *Graph-Informed Opposition-Based Learning (GIOBL)*: Unlike standard Opposition-Based Learning (OBL) which operates on numerical inversion, GIOBL performs a physical state inversion at these high-impact nodes. For a given solution x , its opposition x_{opp} is generated by flipping the silencer state (i.e., installed \leftrightarrow not installed) at these key locations, thereby improving the algorithm's exploratory efficiency.

D. Graph-Informed Local Search (GILS)

The Graph-Informed Local Search (GILS) is a post-processing module designed to refine the final Pareto front. It leverages the system's DAG topology to perform a targeted search, which improves solution convergence and diversity without interfering with the primary evolutionary operators of NSGA-III.

The core strategy of GILS is to focus its computational effort on a set of key solutions (S_{key}), which are strategically significant points on the Pareto front. These are identified as:

- Knee Points: Solutions that are closest to the ideal point in the normalized objective space, representing the most balanced trade-offs.
- Extreme Points: Solutions that represent the best performance for each individual objective.

The GILS algorithm proceeds iteratively as follows:

1) *Preparation*: At the beginning of each iteration t , the current Pareto front, $PF_{current}$, is de-duplicated to create a set of unique solutions.

2) *Key solution identification*: The set of key solutions, S_{key} , is identified from the unique front based on the definitions above.

3) *Adaptive neighborhood search*: For each solution X on the front, a set of neighboring solutions, $N(X)$, is generated. This is done by applying "upgrade" (add/improve a silencer) or

"downgrade" (remove/lessen a silencer) operations at nodes along critical acoustic propagation paths in the DAG. The search depth is adaptive: solutions belonging to S_{key} are subjected to a more extensive neighborhood search.

4) *Evaluation and filtering*: All generated neighborhood solutions X' are evaluated by the physics simulation model. Any solution that violates system constraints is discarded.

5) *HV-based acceptance*: The remaining valid, non-dominated neighbors are considered for inclusion. A candidate solution X' is accepted and used to update the archive ($PF_{current}$) only if its inclusion increases the Hypervolume (HV) indicator of the overall set. That is, $PF_{current}$ becomes $PF_{current} \cup \{X'\}$ only if $HV(PF_{current} \cup \{X'\}) > HV(PF_{current})$.

6) *Termination criteria*: The GILS process terminates when one of two conditions is met: (a) a maximum number of refinement iterations, $T_{refinement}$, is reached, or (b) the improvement in the Hypervolume stagnates between iterations, such that $\Delta HV < \delta HV$.

E. The G-HNSGA-III Algorithm Workflow

The complete workflow of the G-HNSGA-III framework integrates the GINI strategy, the core NSGA-III evolutionary loop, and the GILS post-processing module. The step-by-step process is as follows:

1) *Initialization*: The topological metrics of the DAG are pre-computed (as described in Section II-C). Key parameters for the NSGA-III algorithm (e.g., population size N , max generations G_{max} , crossover/mutation rates) are set.

2) *Graph-Informed Population Generation (GINI)*: The initial population, P_0 (size N), is created using the GINI strategy (Section III-C). Each individual in P_0 is then evaluated using the physics-based simulation model to determine its objective function values.

3) *Evolutionary loop*: The evolutionary process iterates for generations $g = 0$ to G_{max} to drive the population toward convergence along the Pareto front. This process consists of three substeps:

a) *Variation*: An offspring population, Q_g (size N), is generated from the current population, P_g , using standard genetic operators (selection, crossover, and mutation).

b) *Evaluation*: All individuals in the offspring population Q_g are evaluated.

c) *Selection*: The parent and offspring populations are merged into a combined population $R_g = P_g \cup Q_g$ (size $2N$). A fast non-dominated sort is applied to R_g . The next generation, $P_{(g+1)}$, is constructed by selecting individuals from the best non-dominated fronts (F_1, F_2, \dots). The reference-point mechanism of NSGA-III is used to ensure diversity when selecting from the last front that cannot be fully included.

4) *Local refinement (GILS)*: Upon termination of the evolutionary loop, the final non-dominated front, PF_{final} , is passed to the GILS module. GILS iteratively refines this front to enhance the convergence and diversity of the solutions.

5) *Final output:* The refined Pareto-optimal solution set from GILS is returned as the final output of the algorithm. The

complete workflow described above is visually summarized in the flowchart in Fig. 3.

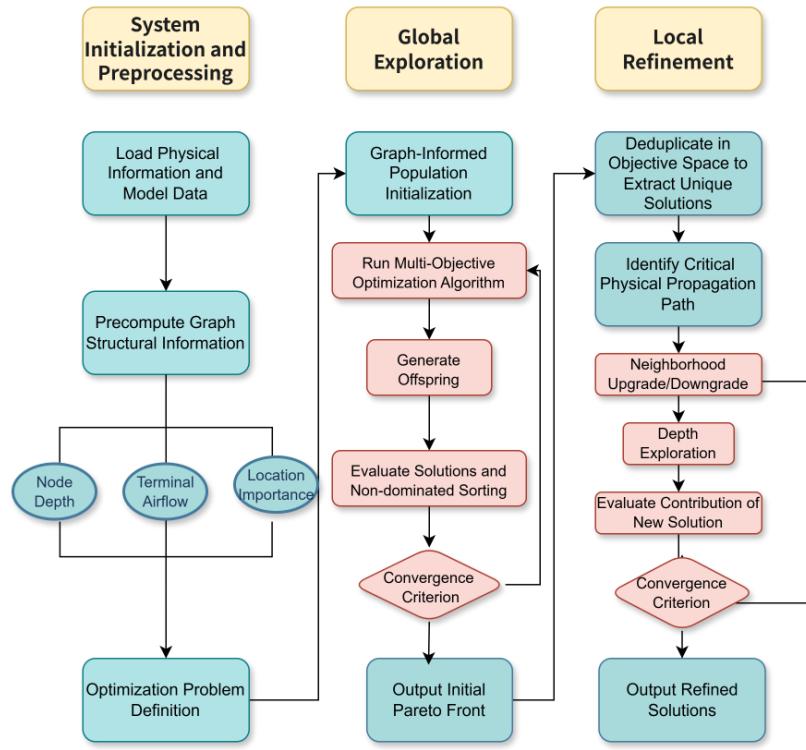


Fig. 3. Flowchart of the multi-objective optimization program.

V. EXPERIMENTAL SETUP AND VALIDATION

This chapter details the computational experiments designed to validate the performance of the proposed G-HNSGA-III framework. All experiments were conducted on a complex, large-scale data center air supply duct system, which serves as the primary test instance.

A. Computational Environment and Simulation Model

The simulation and optimization environment was developed in Python 3.9. The physics-based model for fluid dynamics, acoustics, and economics was implemented using the PyTorch library. All MOEAs, including our proposed framework, were implemented using the Platypa-opt library. The hardware platform consisted of an Intel Core i7-10700 CPU and 16 GB of RAM, running a Windows 11 operating system.

B. Test Instance: Data Center HVAC System

A complex and asymmetrical data center HVAC supply air system was used as the test instance to ensure the problem's practical relevance and difficulty. A partial schematic of the system's DAG topology is shown in Fig. 4. The key parameters and engineering constraints of this system are summarized in Table I.

C. Compared Algorithms and Parameter Settings

To validate the performance of the proposed G-HNSGA-III framework, it was benchmarked against a baseline NSGA-III and six other established MOEAs: NSGA-II, SPEA2, GDE3, MOEAD, and OMOPSO.

TABLE I. KEY PARAMETERS OF THE DATA CENTER TEST INSTANCE

Parameter Category	Parameter	Value	Unit
Topological Properties	Total nodes (components)	44	-
	Potential silencer locations	22	-
	Number of fans	2	-
	Total edges (paths)	40	-
Geometric Properties	Main duct cross-section	1500 x 800	mm
	Branch duct cross-section	500 x 500	mm
Operational Requirements	Number of terminal outlets	12	-
	Average outlet airflow demand	0.6	m³/s
Decision Variables	Available silencer models	18	-
Economic Factors	Electricity cost	0.85	RMB/kWh
	Annual operating hours	2500	h
Engineering Constraints	Max. main duct velocity	5.0	m/s
	Max. branch duct velocity	4.0	m/s
	Max. outlet air velocity	2.0	m/s
	Max. room noise criterion	40	dB(A)
	Max. fan static pressure	800	Pa

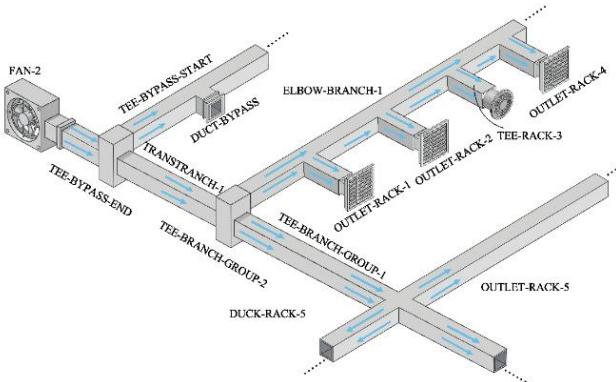


Fig. 4. Node processing of ventilation duct model.

For a fair and rigorous comparison, all algorithms were executed under a consistent computational budget: a population size of 100 (105 for MOEAD, due to its reference vector requirements) and a termination criterion of 15,000 total function evaluations. To ensure statistical robustness, each experimental run was independently repeated five times.

The specific parameter configurations for each algorithm are detailed in Table II. Standard evolutionary parameters, such as crossover and mutation probabilities, were set to common values from the literature. The unique parameters for the GINI and GILS components of G-HNSGA-III were determined through a separate calibration process, as detailed in Section V-E.

TABLE II. PARAMETER CONFIGURATION FOR COMPARED ALGORITHMS

Algorithm	Parameter	Value
General Settings		
All Algorithms	Population / Swarm Size (N)	100 (105 for MOEAD)
	Crossover Operator	Simulated Binary (SBX)
	Mutation Operator	Polynomial
Algorithm-Specific Parameters		
NSGA-II/III, G-HNSGA-III, SPEA2	SBX crossover probability (p_c)	0.9
	Polynomial mutation probability(p_m)	0.1
	distribution index (η)	20.0
GDE3	Scaling Factor (F)	0.5
	Differential evolution crossover rate (CR)	0.9
MOEAD	Neighborhood Size (T)	20
OMOPSO	Swarm Size(N)	100
	Archive Size	100
G-HNSGA-III Specific Parameters		
GINI	seeding fraction (p_{seed})	0.12
	Importance ratio(r_{imp})	0.1
	Assignment probability (p_{assign})	0.8
GILS	Max refinement iterations (T_{ref})	10
	Max search depth (key nodes)	3
	Neighborhood search depth (normal nodes)	1

D. Performance Evaluation Metrics

To quantitatively evaluate the performance of the compared algorithms, three widely-used metrics were employed: Set Coverage (SC), Inverted Generational Distance (IGD), and Hypervolume (HV).

1) *Set Coverage (SC)*: The SC metric (also known as the C-metric) quantifies the Pareto dominance relationship between two solution sets, P_A and P_B . The value $SC(P_A, P_B)$ represents the fraction of solutions in P_B that are weakly dominated by at least one solution in P_A , as computed by Eq. (12).

$$SC(P_A, P_B) = \frac{|\{b \in P_B \mid \exists a \in P_A: a \succeq b\}|}{|P_B|} \quad (12)$$

where $a \succeq b$ denotes that solution a weakly dominates b . A value of $SC(P_A, P_B) = 1$ indicates that all solutions in P_B are dominated by P_A , signifying the superiority of algorithm A over B.

2) *Inverted Generational Distance (IGD)*: The IGD metric quantifies the quality of a solution set (P_{alg}) by computing the average Euclidean distance from a reference Pareto front (P^*) to P_{alg} , as formulated in Eq. (13).

$$IGD(P^*, P_{alg}) = \frac{\sum_{p^* \in P^*} \min_{p \in P_{alg}} d(p^*, p)}{|P^*|} \quad (13)$$

A lower IGD value indicates a superior solution set that is both closer to the true front and better distributed.

Since the true Pareto front is unknown for this complex engineering problem, a high-fidelity reference front P^* was constructed by aggregating the final non-dominated solution sets from all 5 independent runs of all 8 compared algorithms. A non-dominated sort was performed on this aggregate archive, yielding 786 unique solutions. This set serves as the best-known approximation of the true Pareto front for all IGD calculations.

3) *Hypervolume (HV)*: The HV indicator is a unary metric, fully compliant with Pareto dominance, that measures both the convergence and diversity of a solution set P_{alg} . It calculates the volume of the objective space dominated by P_{alg} relative to a reference point P_{ref} . A larger HV value signifies a better Pareto front approximation. In this study, the reference point P_{ref} was set to the nadir point (the vector of worst objective values) observed across all experimental runs, supplemented by a small offset to ensure that all generated solutions strictly dominate P_{ref} .

E. Hyperparameter Calibration for G-HNSGA-III

A sensitivity analysis was performed to optimize the core hyperparameters of the G-HNSGA-III framework. All experiments were conducted on the data center test instance, with each setting repeated 7 times under a consistent computational budget of 15,000 function evaluations.

1) *GILS max search depth*: The Local Refinement Exploration Depth for the GILS module was tested across integer values from 0 (disabled) to 6. As shown in Fig. 5, the

results confirm the module's efficacy. Both HV and IGD metrics identify a consistent optimum at a search depth of 3, achieving a peak HV of 0.358 and a minimum IGD of 0.00326. This setting was chosen as it achieves the peak demonstrated performance; increasing the search depth further leads to performance degradation while unnecessarily increasing the computational budget. Excessively deep local searches can cause the population to converge prematurely to a suboptimal region, stifling the global exploration required to find the true Pareto front.

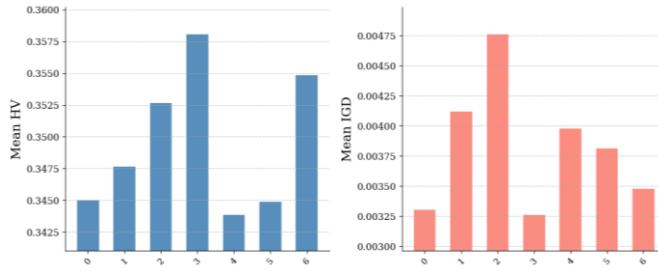


Fig. 5. Performance of the H-NSGAIII algorithm at different local refining exploration depths.

2) *GINI heuristic seeding fraction*: The Heuristic Seeding Fraction (p_{seed}) for the GINI module was tested from 0.0 (disabled) to 0.23. The results, presented in Fig. 6, demonstrate a clear peak in performance at a fraction of 0.12, which achieved the highest HV (0.360) and the lowest IGD (0.00384). Performance degrades beyond this point, likely because an excessive injection of guided solutions prematurely reduces the initial population's diversity and hinders broader exploration. Therefore, a value of 0.12 was selected as it provides the optimal balance between a guided initialization and sufficient initial diversity.

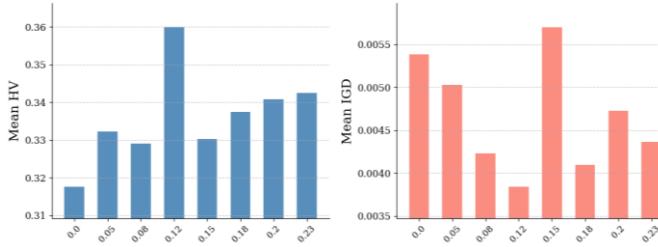


Fig. 6. Mean Hypervolume (HV) and Mean Inverted Generational Distance (IGD) of G-HNSGA-III under different settings for the initial population adjustment ratio parameter.

VI. RESULTS AND DISCUSSION

This chapter presents the quantitative performance comparison of the G-HNSGA-III framework against the baseline NSGA-III and six other established MOEAs.

A. Overall Performance Comparison and Solution Set Quality

All eight algorithms were independently executed 10 times under a consistent computational budget of 15,000 function evaluations. The statistical distributions of the final

Hypervolume (HV) and Inverted Generational Distance (IGD) metrics across these runs are summarized in Table III.

TABLE III. STATISTICAL RESULTS FOR HV AND IGD METRICS ACROSS 10 INDEPENDENT RUNS

Algorithm	Mean HV	Std HV	Median HV	Mean IGD	Std IGD	Median IGD
G-HNSGA-III	0.3891	0.0122	0.3910	0.0030	0.0008	0.0030
GDE3	0.2950	0.0124	0.2913	0.0154	0.0011	0.0152
MOEAD	0.1944	0.0229	0.1848	0.0239	0.0038	0.0234
MOGWO	0.1824	0.0477	0.1942	0.0119	0.0010	0.0119
NSGAII	0.2697	0.0197	0.2665	0.0193	0.0021	0.0197
NSGAI	0.2812	0.0173	0.2814	0.0202	0.0029	0.0195
OMOPSO	0.2655	0.0323	0.2626	0.0317	0.0067	0.0342
SPEA2	0.2608	0.0281	0.2519	0.0150	0.0033	0.0160

The statistical results confirm the superior performance of the G-HNSGA-III framework. For Hypervolume (HV), G-HNSGA-III achieved the highest mean value of 0.3891, a 38.4% improvement over the baseline NSGA-III (0.2812). This result, combined with a low standard deviation (0.0122), demonstrates superior overall solution quality and high reliability. Similarly, the framework obtained the lowest mean Inverted Generational Distance (IGD) of 0.0030, confirming the closest and most complete approximation of the best-known Pareto front. An exceptionally low standard deviation of 0.0008 further underscores its robust stability.

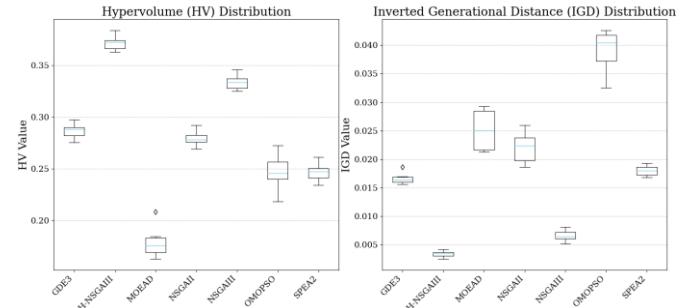


Fig. 7. Boxplot distributions of the final Hypervolume (HV) and Inverted Generational Distance (IGD) metrics for all compared algorithms.

The statistical distributions of the final metric values are visualized using boxplots in Fig. 7 to provide a clearer comparison of algorithm robustness and consistency.

The plots visually affirm the superiority of G-HNSGA-III. For the HV metric (Fig. 8, left), the G-HNSGA-III boxplot is located in a distinctly higher region than all competitors, indicating a consistently better overall performance. Furthermore, its compact interquartile range (IQR) and short whiskers signify high stability and low variance across the independent runs. A similar pattern is observed for the IGD metric (Fig. 8, right), where the G-HNSGA-III boxplot is positioned significantly lower and exhibits the tightest

distribution. Collectively, these plots confirm the findings from Table IV: G-HNSGA-III not only achieves a superior median performance but also does so with the highest level of consistency.

Fig. 8 illustrates the final non-dominated solution sets obtained by each algorithm, plotted in the three-dimensional objective space (Noise Level, Total Cost, and Total Pressure Drop). This visualization highlights the differences in the algorithms' ability to approximate the true Pareto front.

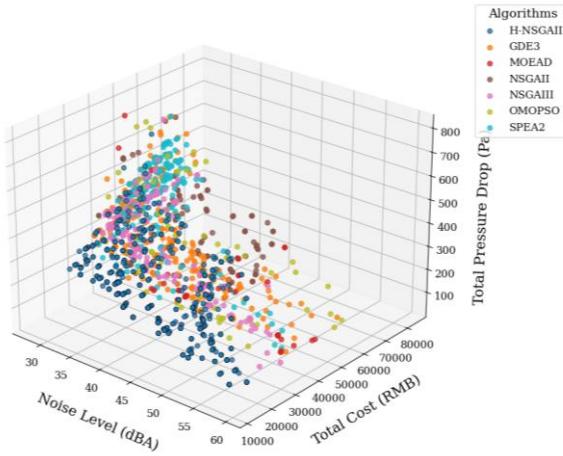


Fig. 8. Comparison of the distribution of solution sets of each algorithm in the three-dimensional target space.

The solution set produced by G-HNSGA-III (cyan) provides a visibly superior approximation of the Pareto front. It extends further toward the ideal point (lower values on all three axes) and achieves a broader, more uniform coverage of the trade-off surface compared to the other algorithms. The solutions from the competing MOEAs are, for the most part, either dominated by or co-located with only a subset of the G-HNSGA-III solutions. This visualization directly corroborates the quantitative HV and IGD results, confirming that the proposed framework excels in both convergence (approaching the optimal region) and diversity (covering a wide range of trade-off solutions).

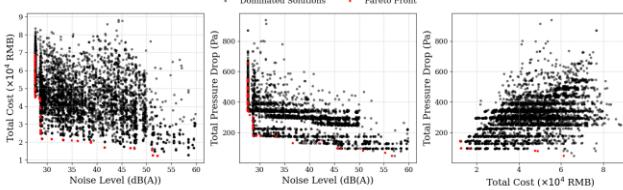


Fig. 9. 2D Pareto front projections of the solution sets of each algorithm under different target pairs.

Fig. 9 projects the aggregated solution set onto three 2D objective planes to visualize the trade-off relationships, with the 2D non-dominated front highlighted in red. The plots confirm the expected conflicts between objectives. The Noise-Cost plane (a) demonstrates that improved acoustic performance requires higher capital investment. The Noise-Pressure Drop plane (b) shows that noise reduction typically increases aerodynamic resistance, leading to higher energy use. Finally, the Cost-Pressure Drop plane (c) identifies the set of most efficient designs that achieve the lowest pressure drop for a given cost.

These visualizations affirm the conflicting, multi-dimensional nature of the problem space. The Pareto front approximated by G-HNSGA-III (shown in Fig. 8) provides the most effective navigation of these trade-offs.

B. Algorithmic Behavior and Dominance Analysis

To analyze the dynamic performance of the algorithms, the HV and IGD metrics were recorded throughout the search process. Fig. 10 plots the metric trajectories from a representative run for each algorithm, illustrating their convergence behavior.

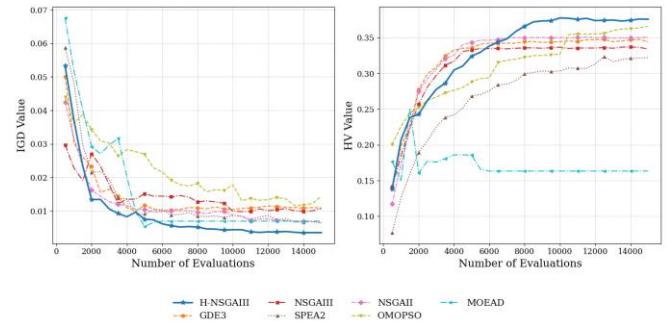


Fig. 10. Comparison of convergence curves of each algorithm on HV and IGD metrics.

The G-HNSGA-III framework (solid blue line) demonstrates superior convergence behavior on both metrics. Its IGD curve (left) exhibits the most rapid descent, reaching a final value of ~ 0.003 , which signifies an efficient search for a high-quality front. Concurrently, its HV curve (right) sustains growth longer than any competitor, diverging after $\sim 2,000$ evaluations to achieve the highest final HV of ~ 0.37 . In contrast, algorithms like MOEA/D stagnate early, indicating premature convergence.

For a direct comparison based on Pareto dominance, the Set Coverage (SC) metric was computed. Fig. 11 presents the full pairwise SC matrix as a heatmap, while Fig. 12 isolates the bidirectional dominance between G-HNSGA-III and each competitor for a focused analysis.

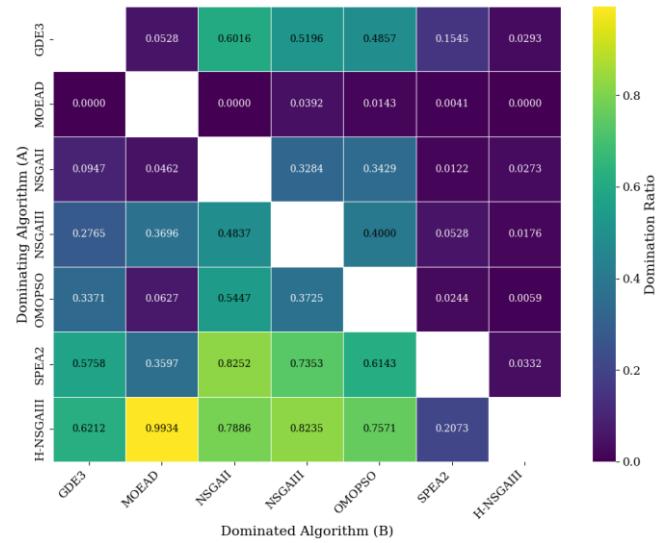


Fig. 11. Heatmap of the pairwise Set Coverage (SC) matrix across all algorithms.

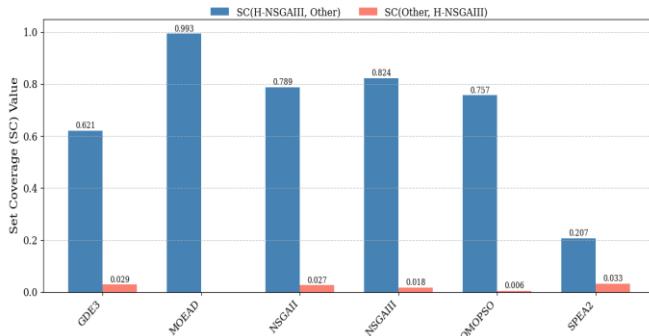


Fig. 12. Comparison of set coverage (SC) between G-HNSGA-III and other algorithms.

The results unequivocally demonstrate the dominance of G-HNSGA-III. As shown in Fig. 12, G-HNSGA-III dominates a very high percentage of the solutions generated by the other algorithms. For example, it dominates 99.3% of the solutions from MOEA/D (SC (G-HNSGA-III, MOEA/D) = 0.9934) and over 82.3% of the solutions from the baseline NSGA-III (SC (G-HNSGA-III, NSGA-III) = 0.8235).

Conversely, the orange bars show that the solution sets from the other algorithms dominate very few, if any, of the solutions produced by G-HNSGA-III. This direct quantitative evidence confirms that the graph-informed strategies enable G-HNSGA-III to discover a Pareto front that is demonstrably superior to those found by the other established MOEAs.

C. Engineering Significance Analysis of Representative Solutions

The high-quality Pareto front generated by G-HNSGA-III provides a rich set of optimal trade-off solutions for engineering decision-making. To illustrate this, Table IV details five representative solutions selected from the final front, including the extremes for each objective and a balanced solution identified using the TOPSIS method [25].

TABLE IV. PERFORMANCE COMPARISON OF TYPICAL SOLUTIONS OBTAINED BY G-HNSGA-III

Solution Type	Noise Level (dBA)	Total Cost (RMB)	Total Pressure Drop (Pa)	Number of Silencers
Lowest Noise Solution	28.67	45688.09	277.20	11
Lowest Cost Solution	49.86	19482.96	187.39	5
Lowest Pressure Drop Solution	45.08	37490.16	126.60	9
Min Cost (Noise \leq 50.0dBA)	46.78	19812.29	176.37	5
Best Compromise (TOPSIS)	38.84	22970.92	178.76	6

The 'Best Compromise' solution represents a practically balanced design. Relative to the 'Minimum Cost' solution, it achieves a substantial 11.1 dBA noise reduction (a 22% improvement) with a moderate 18% increase in cost and negligible impact on pressure drop. Analysis of the system configuration reveals that this acoustic improvement is achieved by deploying a single additional silencer at a critical upstream

location, which was correctly identified by the framework's graph-based importance metrics. This result empirically validates the capacity of G-HNSGA-III to leverage domain knowledge for discovering highly efficient trade-offs.

To facilitate practical project planning, the minimum cost required to achieve specific noise targets was extracted from the global Pareto front and plotted as the 'Minimum Cost Frontier' in Fig. 13. This curve quantifies the necessary capital investment to meet varying acoustic criteria. A distinct non-linear trend is observed, characterized by a 'knee' point in the 35–40 dBA range. Beyond this threshold, each incremental decibel of noise reduction incurs a substantially higher cost, demonstrating the principle of diminishing returns. This insight provides actionable data for cost-benefit analysis in real-world engineering contexts.

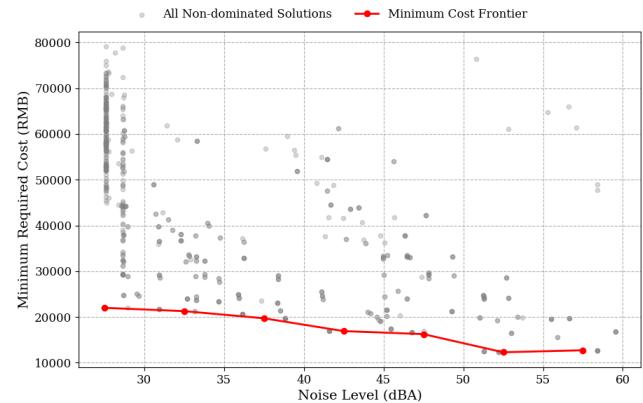


Fig. 13. Noise-cost quantitative trade-off analysis.

VII. CONCLUSION AND FUTURE WORK

This paper introduced G-HNSGA-III, a hybrid optimization framework that enhances NSGA-III by embedding topological domain knowledge from a Directed Acyclic Graph (DAG) model of an HVAC system. The framework's core innovations, a Graph-Informed Initialization (GINI) strategy and a Graph-Informed Local Search (GILS) module, were shown to be highly effective.

Experimental validation demonstrated that G-HNSGA-III significantly outperforms the baseline NSGA-III and six other established MOEAs across all standard metrics (HV, IGD, and SC). The results confirm that the proposed graph-informed heuristics effectively guide the search, leading to a superior Pareto front that offers a range of practical, high-value engineering trade-off solutions.

While the current physical model could be expanded and the framework's generalizability tested on other HVAC typologies, the results confirm the validity of our approach. Future work will focus on integrating surrogate models, such as Graph Neural Networks (GNNs), to improve computational efficiency and adapting this graph-informed optimization paradigm to other network-based engineering domains.

ACKNOWLEDGMENT

This research is based upon work supported by the National Natural Science Foundation of China (No. 62277001).

REFERENCES

[1] G. Huang, S. T. Ng, D. Li, and Y. Zhang, "State of the art review on the HVAC occupant-centric control in different commercial buildings," *J. Build. Eng.*, vol. 96, Art. no. 110445, Nov. 2024, doi: 10.1016/j.jobe.2024.110445.

[2] A. Shahajan, C. H. Culp, and B. Williamson, "Effects of indoor environmental parameters related to building heating, ventilation, and air conditioning systems on patients," *medical outcomes: A review of scientific research on hospital buildings**, *Indoor Air*, vol. 29, no. 2, pp. 161–176, Mar. 2019, doi: 10.1111/ina.12531.

[3] O. Alves, E. Monteiro, P. Brito, and P. Romano, "Measurement and classification of energy efficiency in HVAC systems", *Energy Build.*, vol. 130, pp. 408–419, Oct. 2016, doi: 10.1016/j.enbuild.2016.08.070.

[4] N. Asim et al, "Sustainability of heating, ventilation and air-conditioning (HVAC) systems in buildings—an overview", *Int. J. Environ. Res. Public Health*, vol. 19, no. 2, p. 1016, Jan. 2022, doi: 10.3390/ijerph19021016.

[5] D. Mavrokapanidis, G. N. Lolis, K. Katsigarakis, I. Korolija, and D. Rovas, "Semi-automated extraction of HVAC system topology from imperfect Building Information Models", in *Proc. 18th IBPSA Conf. (BS 2023)*, Shanghai, China, Sept. 2023. doi: 10.26868/2522708.2023.1392.

[6] M. Wang, G. N. Lolis, D. Mavrokapanidis, K. Katsigarakis, I. Korolija, and D. Rovas, "A knowledge graph-based framework to automate the generation of building energy models using geometric relation checking and HVAC topology establishment", *Energy Build.*, vol. 325, Art. no. 115035, Dec. 2024, doi: 10.1016/j.enbuild.2024.115035.

[7] R. Behmanesh, I. Rahimi, and A. H. Gandomi, "Evolutionary many-objective algorithms for combinatorial optimization problems: a comparative study", *Arch. Comput. Methods Eng.*, vol. 28, no. 2, pp. 673–688, Mar. 2021, doi: 10.1007/s11831-020-09415-3.

[8] S. Sharma and V. Kumar, "A comprehensive review on multi-objective optimization techniques: past, present and future", *Arch. Comput. Methods Eng.*, vol. 29, no. 7, pp. 5605–5633, Nov. 2022, doi: 10.1007/s11831-022-09778-9.

[9] D. Ouis, M. A. Hassannain, A. Alshibani, and A. M. Ghaithan, "Noise from heating, ventilation, and air conditioning (HVAC) systems: A review of its characteristics, effects and control", *J. Build. Eng.*, vol. 112, Art. no. 113770, Oct. 2025, doi: 10.1016/j.jobe.2025.113770.

[10] J. Harvie-Clark, N. Conlan, W. Wei, and M. Siddall, "How loud is too loud? Noise from domestic mechanical ventilation systems", *Int. J. Vent.*, vol. 18, no. 4, pp. 303–312, Oct. 2019, doi: 10.1080/14733315.2019.1615217.

[11] G. Iannace, G. Ciaburro, and A. Trematerra, "Heating, ventilation, and air conditioning (HVAC) noise detection in open-plan offices using recursive partitioning", *Buildings*, vol. 8, no. 12, p. 169, Dec. 2018, doi: 10.3390/buildings8120169.

[12] X. Ruan and Y. Yin, "Analysis of common problems of noise and vibration in building heating ventilation air conditioning design", *E3S Web Conf.*, vol. 283, p. 01051, 2021, doi: 10.1051/e3sconf/202128301051.

[13] M. Piwowarski and D. Jakowski, "Areas of fan research—a review of the literature in terms of improving operating efficiency and reducing noise emissions", *Energies*, vol. 16, no. 3, p. 1042, Jan. 2023, doi: 10.3390/en16031042.

[14] Y. B. Alexandrov, V. A. Sychenkov, R. R. Khaliulin, W. M. Yousef, and S. A. Semichev, "Ventilation systems noise reduction issue", in *2020 International Conference on Dynamics and Vibroacoustics of Machines (DVM)*, Samara, Russia, Sept. 2020, pp. 1–6. doi: 10.1109/DVM49764.2020.9243922.

[15] M. A. Oliveira, L. Bragança, S. M. Silva, D. Paixão, and J. António, "Noise reduction performance of a low energy façade-integrated mechanical ventilator", *Build. Acoust.*, vol. 28, no. 2, pp. 119–139, June 2021, doi: 10.1177/1351010X20966185.

[16] S. Li, J. Zhang, J. Hou, B. Zhang, L. Yang, and M. Zheng, "Multi-objective optimization of pressure regulators in buildings," *HVAC systems*, *J. Build. Eng.*, vol. 76, Art. no. 107260, Oct. 2023, doi: 10.1016/j.jobe.2023.107260.

[17] H. Wang et al, "BIM-based automated design for HVAC system of office buildings—An experimental study," *Build. Simul.*, vol. 15, no. 7, pp. 1177–1192, July 2022, doi: 10.1007/s12273-021-0883-7.

[18] J. Glamsch, T. Rosnitschek, and F. Rieg, "Initial population influence on hypervolume convergence of NSGA-III," *Int. J. Simul. Model.*, vol. 20, no. 1, pp. 123–133, Mar. 2021, doi: 10.2507/IJSIMM20-1-549.

[19] X. Xu, D. Cheng, D. Wang, Q. Li, and F. Yu, "An improved NSGA-III with a comprehensive adaptive penalty scheme for many-objective optimization," *Symmetry*, vol. 16, no. 10, p. 1289, Oct. 2024, doi: 10.3390/sym16101289.

[20] D. Ma, S. Zhou, Y. Han, W. Ma, and H. Huang, "Multi-objective ship weather routing method based on the improved NSGA-III algorithm," *J. Ind. Inf. Integr.*, vol. 38, p. 100570, Mar. 2024, doi: 10.1016/j.ji.2024.100570.

[21] Z. Li et al, "GA-LSTM and NSGA-III based collaborative optimization of ship energy efficiency for low-carbon shipping," *Ocean Eng.*, vol. 312, Art. no. 119190, Nov. 2024, doi: 10.1016/j.oceaneng.2024.119190.

[22] A. Behravan, R. Obermaisser, D. H. Basavegowda, and S. Meckel, "Automatic model-based fault detection and diagnosis using diagnostic directed acyclic graph for a demand-controlled ventilation and heating system in Simulink," in *2018 Annu. IEEE Int. Syst. Conf. (SysCon)*, Vancouver, BC, Canada, Apr. 2018, pp. 1–7. doi: 10.1109/SYSCON.2018.8369614.

[23] Z. Sheng, Y. Xu, S. Xue, and D. Li, "Graph-Based Spatial-Temporal Convolutional Network for Vehicle Trajectory Prediction in Autonomous Driving," *IEEE Trans. Intell. Transport. Syst.*, vol. 23, no. 10, pp. 17654–17665, Oct. 2022, doi: 10.1109/TITS.2022.3155749.

[24] 2021 ASHRAE Handbook—Fundamentals, Inch-Pound ed., Peachtree Corners, GA: ASHRAE, 2021.

[25] V. Pandey, Komal, and H. Dincer, "A review on TOPSIS method and its extensions for different applications with recent development," *Soft Comput.*, vol. 27, no. 23, pp. 18011–18039, Dec. 2023, doi: 10.1007/s00500-023-09011-0.