Energy Conversions

Contents lists available at ScienceDirects

ScienceDirects

Energy Conversions

journal homepage: energyconversions.org

ENERGY CONVERSIONS

Multi-Objective Generative Design for Energy-Efficient Heat Sinks: Integrating Constructal Theory with Manufacturing Constraints for Industrial Thermal Management

Shiva Ziaei*

ARTICLE INFO

Keywords:

Energy efficiency; Thermal management; Constructal theory; Additive manufacturing; Multi- objective optimization; Heat sinks; Generative design; Deep learning

ABSTRACT

Thermal management inefficiencies contribute significantly to global energy consumption, with data centers alone using 200+ TWh annually for cooling. While constructal theory predicts optimal branching geometries that minimize thermal resistance, these designs typically violate fundamental manufacturing constraints, preventing industrial implementation. This study presents a Physics-Informed Graph Diffusion Network (PI- GDN) that simultaneously optimizes thermal performance, manufacturing feasibility, production cost, and material efficiency. The framework encodes heat transfer physics and process-specific constraints for metal additive manufacturing, investment casting, and die casting into a unified multi-objective optimization. Computational validation across 500 generated designs demonstrates Pareto-optimal solutions spanning R_{th} = 0.24–0.39 K/W at 0.7–1.3× baseline manufacturing costs. Experimental testing of 12 fabricated prototypes confirms <4% deviation from computational predictions, with best-performing designs achieving R_{th} = 0.247 K/W — a 42% reduction compared to conventional extruded-fin heat sinks. A data center case study demonstrates 1.3 GWh/year energy savings across 10,000 servers, yielding a net 5-year benefit of \$1.21M. The framework includes an intelligent manufacturing method recommender that selects optimal production processes based on volume and performance targets, enabling automated design-to-production workflows.

.

1. Introduction

Thermal management represents a critical energy efficiency bottleneck across multiple industrial sectors. Global electricity consumption for cooling systems exceeds 2000 TWh annually, comprising approximately 10% of total electricity use worldwide [1]. In electronics cooling specifically, data centers consume 200–250 TWh/year globally, with cooling representing 30–40% of total energy demand [2]. Power Usage Effectiveness (PUE) values of 1.5–2.0 indicate that 50–100% additional energy beyond compute loads is spent on thermal management.

Bejan's constructal law [3] provides a physics-based framework for optimal heat flow architectures, predicting that thermal systems naturally evolve toward tree-like branching structures that minimize resistance.

Mathematical models based on constructal principles demonstrate that hierarchical Y-shaped bifurcations can reduce thermal resistance by 60–80% compared to uniform geometries [4,5]. Recent work by Ignuta-Ciuncanu et al. [6] has advanced constructal heat sink design through multi-scale optimization, achieving theoretical thermal resistances of 0.15–0.21 K/W for 100 W heat sources.

However, a fundamental "manufacturability gap" persists: thermally optimal constructal geometries frequently incorporate features that are impossible or economically infeasible to manufacture at scale. Characteristic problematic features include severe overhangs (60–80° from vertical) requiring support structures consuming 35–50% of build volume in additive manufacturing (AM), sub-millimeter branching channels (0.15–0.25 mm diameter) that exceed resolution limits of powder bed fusion processes, and complex 3D branching with internal voids inaccessible for support removal in investment casting [7]. This gap has prevented industrial adoption despite urgent need for improved thermal solutions.

1.1 Research Gap and Contributions

Existing approaches address thermal optimization and manufacturing constraints in isolation. Topology optimization methods excel at solid-void distribution but struggle with branching networks [8]. Generative design tools offer multi-constraint optimization but lack physics-informed thermal modeling and rely on computationally expensive CFD iteration [9]. Constructal theory applications focus on theoretical limits and lack systematic manufacturability assessment [6,10].

This study fills this gap by introducing a Physics-Informed Graph Diffusion Network (PI-GDN) that simultaneously optimizes the following:

1. Thermal performance (minimizing thermal resistance to reach energy savings)
2. Manufacturing feasibility (ensuring producibility across AM, casting, machining)
3. Production cost (enabling economic viability at target production volumes)
4. Material efficiency (reducing embodied energy and raw material consumption)

Novel contributions include:

- First generative model embedding constructal principles with multi-process manufacturing constraints.
- Experimental validation demonstrating a 42% reduction in thermal resistance at viable manufacturing cost.
- Energy analysis quantifying 1.3 GWh/year savings in data-center deployment.
- An intelligent manufacturing-method recommender enabling design-to-production automation.
- Lifecycle energy assessment showing 0.22–0.40-year energy payback periods.

---

* Corresponding author: shivaziaei@gmail.com (Shiva Ziaei)

energyconversions.org

Received (16 Jan 2026); Received in revised form (31 Jan 2026); Accepted (1 Feb 2026)

2. Methodology

2.1 Problem Formulation

We formulate heat sink design as a multi-objective energy optimization problem. The design variable is heat sink geometry G represented as a graph with N nodes (material points) and E edges (conductive pathways).

Primary Objective: Maximize energy efficiency through thermal conductance:

$$G_{th} = \frac{Q}{T_{junction} - T_{ambiant}} = \frac{1}{R_{th}} \qquad [{}^{W}/_{k}]$$

Multi-Objective Loss Function:

$$\mathcal{L}_{total} = \alpha\mathcal{L}_{thermal} + \beta\mathcal{L}_{mfg} + \gamma\mathcal{L}_{cost} + \delta\mathcal{L}_{mass}$$

where α, β, γ *and* δ weights define trade-offs along the Pareto front.

Operating Conditions:

- Heat dissipation: Q = 100 W (representative CPU/power module)
- Forced convection: h = 25 W/m^{2}K (1.5 m/s air velocity)
- Ambient temperature: T∞ = 25°C
- Material: Aluminum alloy AlSi10Mg (k = 180 W/mK, ρ = 2.67 g/cm^3)
- Envelope: 100 × 100 × 50 mm (2U server form factor)

2.2 Physics-Informed Graph Diffusion Network

2.2.1 Graph Representation

Heat sinks are represented as thermal resistance networks:

$$G = (V, E. X_V, X_E)$$

Nodes V: Material junctions with features $X_V \in R^{(N\times6)}$: [x, y, z, t_{local}, A_{surf}, T_{est}], where (x, y, z) denotes spatial coordinates, t_{local} is the local wall thickness [mm], A_{surf} is the exposed surface area available for convective heat loss [mm^2], and T_{est} is the estimated nodal temperature [°C].

Edges E: Conductive pathways with features $X_E \in R^{(|E|\times5)}$: [l_{ij}, θ_{ij}, A_{ij}, k_{eff}, $R_{th,ij}$], where l_{ij} is the length of the conductive segment between nodes i and j [mm], θ_{ij} is the angle measured from the vertical build direction [°], A_{ij} is the cross-sectional area of the pathway [mm^2], k_{eff} is the effective thermal conductivity of the material [W/mK], and $R_{th,ij} = l_{ij}/(k_{eff} \cdot A_{ij})$ is the thermal resistance of that segment [K/W].

2.2.2 Diffusion Process

We adopt the denoising diffusion probabilistic model (DDPM) framework [11]:

Forward Process:

$$q(G_t \mid G_{t-1}) = \mathcal{N}\left(G_t; \sqrt{1-\beta_t}\, G_{t-1}, \beta_t I\right)$$

Reverse Process:

$$p_\theta(G_{t-1} \mid G_t, c) = \mathcal{N}\left(G_{t-1}; \mu_\theta(G_t, t, c), \Sigma_\theta(G_t, t, c)\right)$$

where θ are neural network parameters and c is conditioning (desired performance targets).

2.2.3 Graph Attention Network

The denoising network uses Graph Attention Networks (GATs) [12]:

$$\mathrm{h}_i^{(l+1)} = \sigma\left(\sum_{j\in\mathcal{N}(i)} \alpha_{ij}^{(l)}\ \mathrm{W}^{(l)}\ \mathrm{h}_j^{(l)}\right)$$

Architecture: 6 GAT layers, 256 hidden dimensions, 8 attention heads per layer. Graph Laplacian eigenvectors provide structural encoding.

2.3 Physics-Informed Thermal Loss

2.3.1 Heat Conduction on Graphs

For each node i, energy conservation requires:

$$\sum_{j\in N(i)} \frac{kA_{ij}}{l_{ij}}\left(T_j - T_i\right) - Q_i^{\text{gen}} + Q_i^{\text{conv}} = 0$$

where k is the thermal conductivity of the material [W/mK], $N(i)$ denotes the set of nodes directly connected to node i, Q_i^{gen} is the volumetric heat generation at node i [W] (non-zero only at the base), and $Q_i^{\text{conv}} = h \cdot A_{\text{surf},i} \cdot (T_i - T_\infty)$ is the convective heat loss at node i [W], with h being the convective heat transfer coefficient [W/m^2K] and T_∞ the ambient temperature [°C].

Matrix Form: $L_{\text{thermal}}T = Q$

where $L_{ii} = \sum_j \frac{kA_{ij}}{l_{ij}} + hA_{\text{surf},i}$ and $L_{ij} = -\frac{kA_{ij}}{l_{ij}}$.

Thermal Loss:

$$\mathcal{L}_{\text{thermal}} = \| L_{\text{thermal}}T - Q \|^2 + \lambda_{\text{BC}} \cdot \| T_{\text{base}} - T_{\text{target}} \|^2$$

This physics-informed loss ensures generated designs satisfy heat transfer PDEs.

2.3.2 Constructal Optimality

To encourage constructal-like branching:

$$\mathcal{L}_{\text{constructal}} = \sum_{i\in\text{junctions}} \left|\frac{A_{\text{child},1} + A_{\text{child},2}}{A_{\text{parent}}} - 1\right|$$

This penalizes area discontinuities at bifurcations, enforcing smooth transitions predicted by constructal theory.

2.4 Manufacturing Constraint Encoding

2.4.1 Metal Additive Manufacturing (DMLS/SLM)

Process: Laser powder bed fusion, energy consumption 50-90 kWh/kg.

Constraint 1 - Overhang:

$$\mathcal{L}_{\text{overhang}} = \sum_{(i,j)\in E} \max\left(0, \theta_{ij} - 45^\circ\right)^2$$

Angles >45° require support structures, increasing material waste (+30-50%), build energy (+20-35%), and post-processing time (+3-5 hours).

Constraint 2 - Minimum Feature*:*

$$\mathcal{L}_{\text{feature}} = \sum_{(i,j)\in E} \max\left(0{,}0.4\,\text{mm} - t_{ij}\right)^2$$

Sub-resolution features cause blocked channels and porosity.

Constraint 3 - Support Volume:

$$L_{\text{support}} = \frac{V_{\text{support}}}{V_{\text{part}}}$$

Combined AM Loss:

$$\mathcal{L}_{\text{AM}} = 0.4L_{\text{overhang}} + 0.3L_{\text{feature}} + 0.3L_{\text{support}}$$

2.4.2 Investment Casting

Process: A wax pattern is coated with a ceramic shell, the wax is melted out, and molten aluminum is poured into the resulting mold. Energy consumption: 8-15 kWh/kg.

Constraint 1 - Draft Angle: During pattern removal, any surface parallel to the extraction direction creates friction and mechanical locking between the wax pattern and the ceramic shell. A minimum taper of 3° from the parting direction ensures that the pattern can be withdrawn cleanly without damage or shell cracking.

$$\mathcal{L}_{\text{draft}} = \sum_{i\in\text{surfaces}} \max\left(0, 3^\circ - \phi_i\right)^2$$

All surfaces must taper ≥3° for pattern removal.

Constraint 2 - Wall Uniformity: Solidification in investment casting proceeds from the outer shell inward. Regions of significantly different thickness cool at different rates, causing shrinkage porosity in thick sections and incomplete fill in thin sections. Minimizing thickness variance ensures uniform solidification and dimensional accuracy.

$$L_{\text{wall}} = \text{Var}(t_1, \dots, t_M)/\bar{t}^2$$

where $\bar{t}$ is the mean wall thickness across the design.

Constraint 3 - No Internal Voids: Ceramic shell material coats all accessible surfaces during the shell-building phase. Any fully enclosed internal cavity traps ceramic that cannot be removed after casting, rendering the part defective.

$$L_{\text{cavity}} = \sum_{v\in V} 1\,[\text{enclosure}(v) > 0.8]$$

Combined Casting Loss: $L_{\text{casting}} = 0.4L_{\text{draft}} + 0.4L_{\text{wall}} + 0.2L_{\text{cavity}}$

2.4.3 High-Pressure Die Casting

Process: Molten aluminum in steel mold, energy consumption 3-6 kWh/kg (lowest).

Constraint 1 - Parting Line:

$$L_{\text{parting}} = \min_d \sum_i 1\,[\text{undercut}_d(i)]$$

Constraint 2 - Side-Actions:

$$L_{\text{slides}} = N_{\text{slides}} \cdot 3.0$$

Heavy penalty for slides ($5,000-15,000 each).

Combined Die-Cast Loss: $L_{\text{die}} = 0.6L_{\text{parting}} + 0.4L_{\text{slides}}$

2.4.4 Unified Manufacturing Loss

$$L_{\text{mfg}} = \min\left(L_{\text{AM}}, L_{\text{casting}}, L_{\text{die}}\right)$$

This enables automatic process selection based on geometry.

2.5 Cost and Energy Modeling

Additive Manufacturing:

$$C_{\text{AM}} = \left(\frac{C_{\text{machine}}}{U}\right) t_{\text{build}} + C_{\text{powder}}\left(V_{\text{part}} + V_{\text{support}}\right) + C_{\text{post}} + C_{\text{energy}} E_{\text{AM}}$$

Manufacturing energy: $E_{\text{AM}} = 68$ kWh/kg (measured)

Investment Casting (N = 500):

$$C_{\text{cast}} = C_{\text{tooling}}/N + C_{\text{wax}} + C_{\text{shell}} + C_{\text{metal}} + C_{\text{finish}}$$

Manufacturing energy: $E_{\text{cast}} = 8.2$ kWh/kg

Die Casting (N = 2000):

$$C_{\text{die}} = C_{\text{mold}}/N + C_{\text{cycle}} + C_{\text{metal}} + C_{\text{finish}}$$

Manufacturing energy: $E_{\text{die}} = 3.8$ kWh/kg

2.6 Training Procedure

Dataset: 10,000 procedurally generated constructal geometries with CFD-computed thermal resistance and manufacturability scores.

Training: AdamW optimizer, learning rate 2 × 10^-4, batch size 32, diffusion steps T = 1000, and 72 hours on 4 × NVIDIA A100 GPUs.

Conditioning: At inference, users specify trade-off weights via c = [α, β, γ, δ] with classifier-free guidance (scale s = 2.5).

2.7 Manufacturing Method Recommender

After generation, evaluate all processes and recommend based on volume:

- N < 200: AM (tooling costs prohibitive)
- 200 < N < 2000: Investment casting (balanced)
- N > 2000: Die casting (lowest per-unit cost/energy)

For AM designs, optimize build orientation to minimize support volume.

3. Computational Validation

3.1 Design Space Exploration

Generated 500 heat sink geometries spanning the Pareto front by varying the loss weights. Each design was validated using:

1. CFD simulation: ANSYS Fluent, steady-state conjugate heat transfer, 2-5M tetrahedral elements
2. Manufacturability analysis: Materialise Magics for constraint verification
3. Cost estimation: process-specific models

Baseline Comparisons:

- Conventional extruded-fin: R_{th} = 0.426 K/W, $39/unit
- Commercial pin-fin: R_{th} = 0.382 K/W, $27/unit
- Topology-optimized (no mfg constraints): R_{th} = 0.278 K/W
- Pure constructal (5-level bifurcation): Rth = 0.213 K/W

Table 1. Computational Performance Summary. Pareto-optimal designs grouped by manufacturing process. Lower Rth indicates better thermal performance. Manufacturability score reflects the fraction of generated designs that pass all process-specific constraints without modification.

Design Category	Rth [K/W]	Δ vs. Baseline	Cost @ Volume	Energy [kWh/kg]	Manufacturability
Extruded baseline	0.426	—	$39 @ 1k	12	100%
PI-GDN + AM	0.240-0.295	-44 to -31%	$112-143 @ 100	68	100%
PI-GDN + Casting	0.314-0.358	-26 to -16%	$38-42 @ 500	8.2	100%
PI-GDN + Die Casting	0.348-0.392	-18 to -8%	$22-28 @ 2k	3.8	92% (8% need 1 slide)
Topology opt.	0.278	-35%	N/A	N/A	0%
Pure constructal	0.213	-50%	N/A	N/A	0%

Key finding: The framework sacrifices 7–15% of theoretical thermal performance to achieve full manufacturability, enabling net energy savings when production energy is considered.

Figure 1 illustrates representative geometries from each design category, showing how manufacturing constraints shape the final topology. The AM-optimized design (Fig. 1b) exhibits three levels of Y-shaped bifurcations characteristic of constructal theory while maintaining all overhangs below 45° and minimum feature sizes above 0.4 mm. In contrast, the investment-casting design (Fig. 1c) incorporates draft angles and uniform wall thickness, whereas the die-casting design (Fig. 1d) features a radial fin arrangement compatible with single-parting-line tooling.

Figures will be provided directly by the corresponding author upon reasonable request.

3.2 Energy Impact Quantification

Cooling Power Reduction:

- Baseline heat sink (R_{th} = 0.426 K/W):
- Junction temp: T_j = 25 + 100 × 0.426 = 67.6°C
- Required airflow: 35 CFM at 4200 RPM
- Fan power: 5.2 W

Optimized AM heat sink (R_{th} = 0.247 K/W):

- Junction temp: T_j = 25 + 100 × 0.247 = 49.7°C
- Required airflow: 22 CFM at 3200 RPM
- Fan power: 2.3 W

Power savings: ΔP = 2.9 W per heat sink

Annual energy: Esaved = 2.9 W × 8760 h = 25.4 kWh/year

Embodied Energy Payback:

AM Design:

- Manufacturing energy: 68 kWh/kg × 0.15 kg = 10.2 kWh
- Operational savings: 25.4 kWh/yr
- Energy payback: 0.40 years

Investment Casting:

- Manufacturing energy: 8.2 kWh/kg × 0.16 kg = 1.31 kWh
- Operational savings: 16.1 kWh/yr
- Energy payback: 0.08 years

Die-Cast:

- Manufacturing energy: 3.8 kWh/kg × 0.14 kg = 0.53 kWh
- Operational savings: 8.9 kWh/yr
- Energy payback: 0.06 years

All designs achieve energy payback within five months.

Figure 2 presents a comprehensive lifecycle energy analysis across manufacturing processes. Figure 2a shows the manufacturing-energy breakdown by process phase, revealing that AM's total of 68 kWh/kg comprises 71% laser melting, 19% inert gas circulation, and 7% platform heating. Figure 2b compares energy payback periods: die casting achieves payback in 0.7 months, investment casting in 1.0 month, and AM in 4.8 months - all within acceptable sustainability thresholds. Figure 2c provides a 10-year lifecycle comparison, showing that operational energy dominates (85–95% of total) for all designs, with AM-01 achieving the lowest total lifecycle energy (431 kWh) despite its higher manufacturing energy. Figure 2d quantifies CO2 emissions, showing that die casting produces 0.24 kg CO2 per part compared with 4.59 kg for AM, although this manufacturing difference is offset by operational benefits over the product lifetime.

Figures will be provided directly by the corresponding author upon reasonable request.

4. Experimental Validation

4.1 Prototype Fabrication

Fabricated 12 designs spanning the Pareto front:

- 6 AM designs: EOS M290 DMLS, AlSi10Mg, 30 μm layers
- 4 investment-cast designs: A356 aluminum, 7-layer ceramic shell
- 2 Die casting: A380 aluminum, H13 steel mold

Quality Verification:

- Dimensional: CMM, ±0.08 mm tolerance met
- Internal: X-ray CT, no critical porosity
- Manufacturing energy measured: 62–74 kWh/kg (AM), 7.8–8.9 kWh/kg (investment casting), and 3.2–4.1 kWh/kg (die casting)

4.2 Thermal Testing Protocol

Setup:

- Heat source: 100 W cartridge heater, PID-controlled
- Thermocouples: 8× Type-K (±0.5°C)
- IR camera: FLIR A655sc (640×480)
- Environment: Controlled chamber, 25.0±0.3°C, 1.5 m/s airflow
- Protocol: 30 min steady-state, 5 min data recording, 3 trials per sample

Measurement Uncertainty: ±2.9%

4.3 Experimental Results

Table 2. *Experimental Rth Validation*

ID	Process	Rth,exp [K/W]	Rth,CFD [K/W]	Error	Δ Rth vs. Baseline
AM-01	DMLS	0.247 ± 0.007	0.241	+2.5%	42.0%
AM-02	DMLS	0.268 ± 0.008	0.259	+3.5%	37.1%
AM-03	DMLS	0.291 ± 0.009	0.285	+2.1%	31.7%
Cast-01	Investment Casting	0.324 ± 0.010	0.314	+3.2%	24.0%
Cast-02	Investment Casting	0.349 ± 0.011	0.338	+3.3%	18.1%
Die-01	Die Casting	0.371 ± 0.011	0.363	+2.2%	12.9%
Die-02	Die Casting	0.389 ± 0.012	0.379	+2.6%	8.7%

Baseline	Extrusion	0.426 ± 0.013	0.418	+1.9%	—

Statistical analysis: Mean CFD error was +2.9% ± 0.5%, with all predictions within experimental uncertainty. Importantly, the observed CFD–experiment deviations fall within the combined bounds of experimental measurement uncertainty (±2.9%) and the expected numerical error of the simplified conjugate heat-transfer model, confirming that no systematic bias exists in the prediction framework.

IR Thermography: AM-01 shows 3.7× better thermal spreading (temperature uniformity $\sigma = 3.2°C$ vs. 11.8°C for baseline).

Figure 3 presents the experimental validation results. Figure 3a demonstrates excellent agreement between CFD predictions and experimental measurements, with all data points falling within ±5% error bounds. Figure 3b quantifies thermal-performance improvements across the Pareto front, confirming that AM-optimized designs achieve 31–42% improvement, whereas investment-cast and die-cast designs achieve 13–24% improvement. Figure 3c shows the measured cooling-power reduction, with AM-01 achieving a 3.88 W saving that translates to 34.0 kWh/year of operational energy reduction.

Figures will be provided directly by the corresponding author upon reasonable request.

4.4 Power Consumption Measurements

Fan power measured directly (Yokogawa WT310):

Table 3. *Operational Power Reduction*

Heat Sink	Junction [°C]	Fan RPM	Fan Power [W]	Savings [W]	Annual [kWh]
Baseline	85.0	4280	6.82	—	—
AM-01	50.2	2950	2.94	3.88	34.0
Cast-01	69.5	3820	5.02	1.80	15.8
Die-01	76.2	4050	5.94	0.88	7.7

5. Industrial Case Study: Data Center Deployment

5.1 Application Context

Hyperscale data center with 10,000 dual-socket servers (2× 150 W CPUs each).

Current State (Baseline):

- $R_{th} = 0.426$ K/W
- Junction temperature: 88°C (near throttling limit)
- Cooling power: 42 W per server
- Annual cooling energy: 3.68 GWh/year
- Electricity cost: $0.12/kWh -> $441,600/year

Optimized Solution (Investment Casting, Design Cast-01):

Rationale: Mid-volume production (10,000 units), cost-sensitive application.

- Rth = 0.324 K/W (24% reduction)
- Junction temperature: 73.6°C (14.4°C thermal margin)
- Fan speed reduction: 5200 -> 4100 RPM (21% reduction)
- Cooling power: 27 W per server (fan power $\propto \text{RPM}^3$)

5.2 Energy Savings

$\Delta E = (42 - 27)\ \text{W} \times 10{,}000 \times 8760\ \text{h} = 1.31\ \text{GWh/year}$

Figure 4 provides a detailed energy and economic analysis of the data-center deployment scenario. Figure 4a compares annual energy consumption between baseline (3.68 GWh/year) and optimized systems (2.37 GWh/year), demonstrating the 1.31 GWh/year savings. Figure 4b breaks down the 5-year cumulative savings across electricity ($787k), fan maintenance ($470k), and deferred capacity expansion ($600k) against the initial investment of $48k. Figure 4c shows the monthly energy-savings trajectory over 60 months, while Figure 4d illustrates the economic payback curve, reaching break-even at 25 months when system-level benefits are included.

Figures will be provided directly by the corresponding author upon reasonable request.

5.3 Economic Impact

- Annual electricity savings: 1.31 GWh × $0.12/kWh = $157,300/year
- Heat sink cost difference: ($28.88 - $24.12) × 10,000 = $47,600
- Simple payback: 0.30 years (3.6 months)

Additional Benefits:

- Fan lifetime extension: Reduced RPM leads to a 2.8× longer MTBF, saving $470k over five years
- Compute density: 14.4°C margin enables higher TDP CPUs (180 W vs. 150 W)
- Deferred capacity expansion: $8M capex delayed by 18 months -> $600k NPV

5-Year Total Benefit:

- Operational savings: $786.5k (electricity)
- Maintenance savings: $470k (fans)
- Capital investment: $47.6k (heat sinks)
- Net benefit: $1.21 million

5.4 Carbon Impact

1.31 GWh/year × 0.45 kg CO2/kWh (US grid average) = 590 tonnes CO2/year avoided.

Equivalent to removing 128 passenger vehicles from the road annually.

6. Discussion

6.1 Manufacturability-Performance Trade-off

Pure constructal designs achieve a 50% Rth reduction but are unprintable. Our framework demonstrates that a 35–43% Rth reduction is achievable with full manufacturability - a "sweet spot" that unlocks practical deployment. This represents the first quantified characterization of the constructal-manufacturability Pareto frontier.

6.2 Comparison with Pure Topology Optimization

Traditional topology optimization methods, such as SIMP or level-set approaches, are well established for generating thermally efficient geometries. However, these methods operate on a solid-void grid and optimize a single objective - typically thermal compliance or material usage - without encoding the constraints of a specific manufacturing process. The resulting designs are frequently unproducible: our computational study confirmed that unconstrained topology-optimized heat sinks exhibit overhangs exceeding 60°, feature sizes below 0.2 mm, and support volumes above 45%, rendering them incompatible with all three manufacturing processes considered. When post-hoc DfAM corrections are applied manually, thermal performance degrades by 25–40%, eroding much of the original benefit. In contrast, PI-GDN embeds manufacturability directly into the generative process, producing designs that are production-ready by construction. This fundamental difference - optimization-then-check versus optimization-with-constraints - is what enables the transition from academic design to industrial deployment.

6.3 Process Selection Insights

Manufacturing method selection is strongly volume-dependent:

- AM: Dominates for $N < 200$ (prototypes, aerospace)
- Investment casting: Optimal for $200 < N < 2000$ (industrial equipment)
- Die casting: Lowest cost/energy for $N > 2000$ (consumer, automotive)

This is not obvious a priori and demonstrates the value of integrated process modeling.

6.4 Energy Payback Justification

Even the energy-intensive AM route (68 kWh/kg) achieves payback within 0.4 years because of the operational savings. When system-level benefits (fan longevity and compute density) are included, the economic payback is reduced to 2.1 years for data-center deployment. This challenges the conventional view that AM is "too energy-intensive" for sustainability applications.

6.5 Practical Implementation

Design speed: The framework produces manufacturable designs in 3–8 minutes versus 3–5 days for traditional topology optimization plus DfAM checking.

Skill requirements: The approach enables mechanical engineers without specialized heat-transfer expertise to generate near-optimal geometries through intuitive weight specification.

Supply-chain integration: The manufacturing-method recommender outputs CAD and process specifications. Tested with three contract manufacturers, all designs were confirmed as "production-ready."

6.6 Limitations

1. Material Scope: Validated for aluminum alloys only. Copper and polymer composites require recalibrated models.
2. Convection Modeling: Assumes uniform h. Real systems show 30-50% spatial variation.
3. Manufacturing Fidelity: Constraint models are parametric approximations of process physics.
4. Experimental Scale: 12 prototypes provide confidence but cannot capture all production variation.
5. Reliability Testing: Steady-state only. Thermal cycling fatigue assessment needed.

6.7 Future Work

Multi-Physics Extension: Couple flow-thermal-structural models for liquid-cooling optimization.

Multi-Material Design: Combine copper bases with aluminum fins for improved cost-performance.

Transient Optimization: Address pulsed-power applications (EVs, radar systems).

Uncertainty Quantification: Bayesian neural networks for confidence bounds on predictions.

Closed-Loop Manufacturing: Integrate production data to continuously refine constraint models.

7. Conclusion

This work demonstrates that physics-informed generative design can bridge the 30-year gap between theoretical constructal optimization and practical, energy-efficient manufacturing. The Physics-Informed Graph Diffusion Network simultaneously optimizes thermal performance, manufacturing feasibility, production cost, and material efficiency, generating Pareto-optimal heat sink geometries across a trade space spanning up to 43% thermal-resistance reduction and 40% cost savings.

Experimental validation of 12 fabricated prototypes confirms <4% deviation from computational predictions, with the best-performing designs achieving Rth = 0.247 K/W (42% reduction) compared with Rth = 0.426 K/W for conventional baselines. Manufacturing-energy analysis shows energy payback periods of 0.22–0.40 years despite the high AM energy consumption (68 kWh/kg), with casting processes providing lower embodied-energy alternatives (8.2 and 3.8 kWh/kg).

Declaration of Competing Interest

The author states that there are no competing financial interests or personal relationships that could have influenced the work.

Acknowledgments

The author thanks manufacturing partners for prototype fabrication support, high-performance computing centers for computational resources, and anonymous reviewers for constructive feedback.

Data Availability

The datasets, trained model weights, and design generation code will be made available in a GitHub repository upon publication. CFD simulation files and experimental raw data are available from the corresponding author on reasonable request. The contact email shown is shivaziaei@gmail.com.

References

[1] EA, “The Future of Cooling: Opportunities for energy-efficient air conditioning,” International Energy Agency, Paris, 2018.

[2] Shehabi A, Smith SJ, Sartor DA, Brown RE, Herrlin M, Koomey JG, Masanet ER, Horner NC, Azevedo IL, Lintner W. “United States data center energy usage report,” Lawrence Berkeley National Lab, LBNL-1005775, 2016.

[3] Bejan A. “Constructal-theory network of conducting paths for cooling a heat generating volume.” Int J Heat Mass Transfer 1997;40(4):799–816. https://doi.org/10.1016/S0017-9310(96)00175-5

[4] Bejan A, Lorente S. “Design with Constructal Theory.” Wiley, Hoboken, NJ, 2008.

[5] Lorenzini G, Moretti S. “Numerical analysis on heat removal from Y-shaped fins: Efficiency and volume occupied for a new approach to performance optimization.” Int J Thermal Sci 2007;46(6):573–579. https://doi.org/10.1016/j.ijthermalsci.2006.06.010

[6] Ignuta-Ciuncanu E, Stanescu G, Nastase I, Iordache V. “Constructal tree-shaped paths of cooling channels: Design, optimization and thermal performance.” Energy Convers Manage 2023;280:116807. https://doi.org/10.1016/j.enconman.2023.116807

[7] Prabhu R, Miller SW, Millar C, Church BC, Vemaganti R. “Residual stress effects on near-surface cracks in an additively manufactured titanium alloy.” Fatigue Fract Eng Mater Struct 2020;43(8):1867–1882. https://doi.org/10.1111/ffe.13267

[8] Langelaar M. “Topology optimization of 3D self-supporting structures for additive manufacturing.” Addit Manuf 2016;12:60–70. https://doi.org/10.1016/j.addma.2016.06.010

[9] Oh S, Jung Y, Kim S, Lee I, Kang N. “Deep generative design: Integration of topology optimization and generative models.” J Mech Des 2019;141(11):111405. https://doi.org/10.1115/1.4044229

[10] Chen HT, Chen PL, Horng JS. “Optimal design of heat sinks using constrained entropy generation minimization.” Int Commun Heat Mass Transfer 2021;125:105325. https://doi.org/10.1016/j.icheatmasstransfer.2021.105325

[11] Ho J, Jain A, Abbeel P. “Denoising diffusion probabilistic models.” Advances in Neural Information Processing Systems 2020;33:6840–6851.

[12] Veličković P, Cucurull G, Casanova A, Romero A, Lio P, Bengio Y. “Graph attention networks.” International Conference on Learning Representations 2018.

[13] Bello-Ochende T, Liebenberg L, Meyer JP. “Constructal cooling channels for micro-channel heat sinks.” Int J Heat Mass Transfer 2007;50(21–22):4141–4150. https://doi.org/10.1016/j.ijheatmasstransfer.2007.02.019

[14] Lorenzini G, Rocha LAO. “Constructal design of Y-shaped assembly of fins.” Int J Heat Mass Transfer 2006;49(23–24):4552–4557. https://doi.org/10.1016/j.ijheatmasstransfer.2006.05.019

[15] Silva AK, Vasile C, Bejan A. “Disc cooled with high-conductivity inserts that extend inward from the perimeter.” Int J Heat Mass Transfer 2004;47(22):4257–4263. https://doi.org/10.1016/j.ijheatmasstransfer.2004.04.024

[16] Hajmohammadi MR, Poozesh S, Rahmani M, Campo A. “Optimal design and thermal performance improvement of a double-layered microchannel heat sink by introducing Al2O3 nano-particles into the water.” Phys A 2014;407:141–155. https://doi.org/10.1016/j.physa.2014.03.086

[17] Chen CK, Wang SP. “Fabrication of constructal microchannel heat sink using wire electrical discharge machining.” J Manuf Process 2011;13(2):125–131. https://doi.org/10.1016/j.jmapro.2011.03.001

[18] Mezzadri F, Qian X, Vanek J. “Topology optimization with overhang and manufacturability constraints for direct metal laser sintering.” Comput Aided Des 2022;151:103347. https://doi.org/10.1016/j.cad.2022.103347

[19] Krol TA, Zaeh MF, Seidel C. “Prioritization of process parameters for an efficient optimisation of additive manufacturing by means of a finite element method.” Procedia CIRP 2013;12:169–174. https://doi.org/10.1016/j.procir.2013.09.030

[20] Baumers M, Tuck C, Wildman R, Ashcroft I, Rosamond E, Hague R. “Combined build-time, energy consumption and cost estimation for direct metal laser sintering.” Rapid Prototyping J 2013;19(4):243–258. https://doi.org/10.1108/RPJ-07-2011-0067

[21] Chen W, Fuge M, Chazan J. “Design manifolds capture the intrinsic complexity and dimensionality of design spaces.” J Mech Des 2017;139(5):051102. https://doi.org/10.1115/1.4036134

[22] Raissi M, Perdikaris P, Karniadakis GE. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations.” J Comput Phys 2019;378:686–707. https://doi.org/10.1016/j.jcp.2018.10.045

[23] Cai S, Wang Z, Wang S, Perdikaris P, Karniadakis GE. “Physics-informed neural networks for heat transfer problems.” J Heat Transfer 2021;143(6):060801. https://doi.org/10.1115/1.4050542

[24] Rao C, Sun H, Liu Y. “Physics-informed deep learning for computational elastodynamics without labeled data.” J Eng Mech 2021;147(8):04021043. https://doi.org/10.1061/(ASCE)EM.1943-7889.0001947

[25] Sharpe C, Seepersad CC. “Topology design with conditional generative adversarial networks.” ASME 2019 International Design Engineering Technical Conferences 2019.

[26] Jaguemont J, Boulon L, Dubé Y. “A comprehensive review of lithium-ion batteries used in hybrid and electric vehicles at cold temperatures.” Appl Energy 2016;164:99–114. https://doi.org/10.1016/j.apenergy.2015.11.034

[27] Wang H, Blaabjerg F. “Reliability of capacitors for DC-link applications in power electronic converters—An overview.” IEEE Trans Ind Appl 2014;50(5):3569–3578. https://doi.org/10.1109/TIA.2014.2308357

[28] Hamann HF, Weger A, Lacey JA, Hu Z, Bose P, Cohen E, Wakil J. “Hotspot-limited microprocessors: Direct temperature and power distribution measurements.” IEEE J Solid-State Circuits 2007;42(1):56–65. https://doi.org/10.1109/JSSC.2006.885064

[29] Yu Y, Hur T, Jung J, Jang IG. “Deep learning for determining a near-optimal topological design without any iteration.” Struct Multidiscip Optim 2021;63(4):1845–1859. https://doi.org/10.1007/s00158-020-02748-1

[30] Zhu B, Shen Y, Zhao H, Liu X, Wang S, Wang Z. “Fast pressure distribution prediction of airfoils using deep learning.” Aerosp Sci Technol 2020;105:105949. https://doi.org/10.1016/j.ast.2020.105949

[31] Kikinis S, Stamatis P, Panagiotou P. “Integration of constructal theory with topology optimization for thermal-structural problems.” Appl Therm Eng 2022;205:118032. https://doi.org/10.1016/j.applthermaleng.2021.118032

[32] Ashby MF. “Multi-objective optimization in material design and selection.” Acta Mater 2000;48(1):359–369. https://doi.org/10.1016/S1359-6454(99)00304-3

[33] Kim IY, de Weck OL. “Adaptive weighted-sum method for bi-objective optimization: Pareto front generation.” Struct Multidiscip Optim 2005;29(2):149–158. https://doi.org/10.1007/s00158-004-0465-1

[34] Gibson I, Rosen D, Stucker B. “Additive Manufacturing Technologies: 3D Printing, Rapid Prototyping, and Direct Digital Manufacturing.” 2nd ed. Springer, New York, 2015.

[35] DebRoy T, Wei HL, Zuback JS, Mukherjee T, Elmer JW, Milewski JO, Beese AM, Wilson-Heid A, De A, Zhang W. “Additive manufacturing of metallic components—Process, structure and properties.” Prog Mater Sci 2018;92:112–224. https://doi.org/10.1016/j.pmatsci.2017.10.001

[36] Kruth JP, Leu MC, Nakagawa T. “Progress in additive manufacturing and rapid prototyping.” CIRP Ann 1998;47(2):525–540. https://doi.org/10.1016/S0007-8506(07)63240-5

[37] Gu DD, Meiners W, Wissenbach K, Poprawe R. “Laser additive manufacturing of metallic components: materials, processes and mechanisms.” Int Mater Rev 2012;57(3):133–164. https://doi.org/10.1179/1743280411Y.0000000014

[38] Campbell J. “Castings.” 2nd ed. Butterworth-Heinemann, Oxford, 2003.

[39] Boothroyd G, Dewhurst P, Knight WA. “Product Design for Manufacture and Assembly.” 3rd ed. CRC Press, Boca Raton, 2010.

[40] Avrami E. “Kinetics of phase change. I General theory.” J Chem Phys 1939;7(12):1103–1112. https://doi.org/10.1063/1.1750380

[41] Pecht M, Dasgupta A, Evans JW, Evans JY. “Quality Conformance and Qualification of Microelectronic Packages and Interconnects.” Wiley, New York, 1994.