A Unified Differentiable Framework for Simultaneous Topology, Packing, and Routing Optimization of Compact Electro–Thermal Modules

Phillip Harris Paul[1,*]

[1] Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore, California 94550
* Correspondence: paul_harris@gse.harvard.edu

Abstract: Electro–thermal compact product assemblies including power-electronics modules, embedded controllers, battery interfaces, and dense-pack mechatronic systems have tight coupling between the arrangement of components, connectivity routing topologies, conductive branch geometry, electrical losses, and thermal transport. Existing design procedures typically perform each of these stages in a separated manner. Namely, a possible arrangement of components is first developed, then one of the discrete choices for routing topology is made, and finally geometric refinement takes place. While such strategies make sense in practice, they impose a significant limitation since the routing topology cannot be changed further after starting the physics-based optimization process. Thus, even if the existing routing architecture proves to be inefficient due to high thermal congestion, high Joule losses, or excessive device compaction, the design remains stuck within an inferior routing choice until the design loop starts again. In this work, we propose an approach for solving the electro–thermal routing optimization problem with simultaneous tuning of the topology, device packing, and trace geometry. The candidate branches of the routing graph are represented using relaxable binary variables, while devices and interconnect traces are described by smooth occupancy fields. Graph flow conservation is enforced through graph-flow continuity constraints, electric losses are computed from current-resistive losses, and steady-state thermal transport with boundary convective conditions is considered. The continuation method is used to drive the topology variables towards near-binary values. Comparing our approach with a best-of-many sequentially tuned baseline strategy, we achieve better device compaction, 31.7% decrease in the bounding-box area from 410.4 to 338.2 cm^2, 22.6% improvement in the efficiency of trace geometry, 18.7% reduction of interconnect power dissipation from 18.7 to 14.5 W, as well as thermal management efficiency increase by 8.4%, lowering of maximum device temperature from 103.4 to 94.7°C. The unified approach also avoids multiple restarts of optimization based on individual topology and results in a simpler structure with only eight remaining edges rather than eleven. It is clear from the above observations that topology–geometry co-optimization based on differential equations can offer a significant practical improvement to compact electro-thermal design.

Keywords: electro–thermal design; packing and routing; topology optimization; layout synthesis; design automation; multidisciplinary design optimization

Citation: Phillip Harris Paul. 2021. A Unified Differentiable Framework for Simultaneous Topology, Packing, and Routing Optimization of Compact Electro–Thermal Modules. *TK Techforum Journal (ThyssenKrupp Techforum)* 2021(3): 64–80.

Received: September-03-2021
Accepted: December-11-2021
Published: December-30-2021

1. Introduction

Modern compact engineering systems design involves the combined process of arranging heterogeneous components, forming possible interconnection networks, and accounting for multi-physics phenomena within the strictly defined geometric boundaries of the assembly. Relevant applications range from power-electronic circuits, battery-management

equipment, aerospace controls systems, electrified vehicle hardware, embedded sensors, to multifunctional additively manufactured components. Within such systems, the issue of arranging components and designing connections is not simply a geometric challenge, since it directly affects electrical resistance, parasitic losses, thermal conduction, hot spots formation, serviceability, and manufacturing capability. What can be advantageous geometry-wise might prove to be unfavorable if one takes into account the heating generated by Joule losses, whereas what is optimal thermally can become inefficient if the desired electrical conductivity requires too much wire. The modern mechanical engineering design challenge is thus not to place components and connect them in an efficient way but to find the optimal layout architecture where the three aspects are combined into one [1–3].

The inherent problem in solving this challenge stems from the combination of discrete and continuous decision variables. Arrangement of the devices introduces their position and orientation. Connection design introduces their curvature, thickness, and spacing. And connection network selection introduces the combinatorial problem of choosing among the many available candidate branches. The presence of these decision variables at once makes the search space highly nonconvex and potentially discontinuous. The classic literature therefore took a decomposing approach, applying heuristic search [4], simulated annealing [5], graph drawing [6], or evolution techniques for layout optimization. Meanwhile, the problems of path generation and interconnect routing continue to be approached via shortest-path algorithms [7], A*-search [8], or specialized methods developed in the tradition of VLSI physical design [9]. When dealing with general mechanical engineering design challenges, however, one is forced to account for the complex interactions between geometry and physics in a problem formulation, which brings about the computational complexity [10].

Many works in the field of design automation showed that considering placement and routing in a sequential manner is an appealing idea as it simplifies the optimization while ensuring the feasibility of both steps. The approaches to rectangle packing and interference avoidance are abundant and so are the methods of graph-based layout and routing optimization [11–13]. The ideas were further applied in automotive, industrial, and multifunctional additive manufacturing contexts to solve vehicle layout, electronic interconnect design, and functional layout problems with manufacturable routing in mind [14–16]. Yet in the face of strong interdependencies between electrical performance and heat generation, it becomes ever harder to consider the decisions in isolation.

For electro-thermal systems in particular, such a decomposition approach is limited due to its inherent nature of treating different aspects of design separately. Specifically, compact thermal modeling allows understanding the role of package-level thermal conduction and boundary conditions on component temperatures [17,18]. Meanwhile, the principles of physical design allow one to use graph representations and congestion models to account for the routing feasibility [19]. Conceptually, there is no problem in combining these perspectives into one, but in practice, these considerations are made at different stages of the design process. One can thus see that optimizing design while keeping all these aspects in consideration in one unified form is a more profound design task compared to just checking the result for compactness or heat dissipation after the design is completed.

The situation in compact electro-thermal systems is particularly sensitive to the interaction of different design aspects. Notably, the electrical resistance will increase with the current path length, width, and temperature. Electrical losses will be translated into heat sources. Heat dissipation will depend on temperature-sensitive material properties and on hot spot formation close to high-power electronics. Furthermore, the routing itself is not indifferent to the placement of heat-generating devices or to the current paths running nearby. Overall, the most compact design can differ substantially from the most efficient and safest

routing scheme. In a general sense, the problem is multidisciplinary, which means that the metric of performance emerges as a function of several interacting subsystems [20].

In topology optimization, it was found that a suitable continuous representation allows addressing even a combinatorial decision-making process with gradient-based methods [21–23]. Such representations also enabled topology optimization to evolve beyond the mechanical design realm towards thermal and fluid transport processes [24–26]. Although the current routing problem concerns different aspects, the idea is that a proper relaxation can enable the optimization to consider the otherwise combinatorial decisions of connection choice.

On the other hand, connectivity cannot be treated independently in VLSI physical design either, which introduces additional constraints related to layout and timing in addition to the graph-theory-based aspects [27,28]. Similarly, packaging problems involve various aspects that can affect the routing in the same way. In this sense, wire or busbar routing will differ from the interconnect routing in VLSI, as certain geometric characteristics can be added to the routing problem. Moreover, the electro-thermal performance can force a designer to consider a geometry which looks bad from a geometrical point of view but is thermally acceptable. In light of the above, one can see the rationale behind addressing the problem of interest from an electro-thermal point of view.

Finally, it needs mentioning that there exists a novel framework of compact interconnected systems' packing and optimization, which demonstrates an interesting development towards integrating the aspect of physics into topology and shape optimization [29]. The significance of the latter is twofold. On the one hand, it provides a good result in itself. On the other hand, the distinction of system layout optimization and topology optimization was shown for the first time in the field. Yet, once a certain topology is fixed, the next stage of optimization does not account for potential changes in the topology, which is beneficial for a small design problem but not for a big one with multiple candidates.

Thus, the problem addressed in the current work is finding the routing network's optimal topology and layout simultaneously by accounting for electro-thermal performance. The goal is to select an optimal network topology, to decide where to put the components, to set the curvature of connections, and how thick those connections should be in one formulation. Such problem entails three requirements. Namely, geometry representation has to be differentiable. Topology representation needs to ensure connectivity and meaningful intermediate states. The electro-thermal model must be able to provide feedback to optimize performance criteria.

To address these requirements, a representation of candidate branches is developed in the framework of a graph embedded into the domain. Each branch is associated with a relaxed activation variable which allows to suppress unnecessary connections and retain the required ones during the optimization process. Devices and connections are modeled by smooth occupancy fields enabling a consistent differentiation of overlaps and fields in one framework. Connectivity constraints are imposed via graph flow constraints ensuring the connection topology preservation. Electrical losses are evaluated by the electric resistivity model taking into account the width and temperature of branches. Thermal behavior is described by a steady-state conduction problem governed by the temperature boundary condition at the heat sources and convective boundaries. The final solution is enhanced by employing the continuation approach.

From this perspective, the problem belongs to contemporary design research in mechanical engineering. The aspects of design methodology, optimization, thermo-mechanical modeling, layout generation, and compact system integration are all related to the field. In terms of the journal, the problem falls into the category of design automation, computational modeling, and engineering systems optimization.

The main contributions of this paper are as follows. An optimization framework is developed for simultaneous layout and routing optimization in compact electro-thermal systems. A smooth field representation is developed for devices and connections, allowing overlap control, thermal effects, and footprint calculation. A connectivity-preserving topology representation based on graphs is presented in the formulation. A coupled electro-thermal state representation is proposed for providing the performance feedback. Finally, the framework is demonstrated in a computational study, where its efficiency is compared against topology-fixed strategies.

The paper is structured as follows. Problem formulation and corresponding gaps in the existing literature are discussed in Section 2. The framework of electro-thermal system representation and constraints is discussed in Section 3. The computational study is described in Section 4. Discussion and analysis follow in Section 5, while conclusions are drawn in Section 6.

2. Problem Definition and Research Gap

Suppose a two-dimensional domain $\Omega \subset \mathbb{R}^2$, within which N_d rigid devices are installed, and N_n electrically active nets between certain ports on these devices are defined. The architecture of the system is fixed in the sense that the set of required source–sink correspondences is given. All other choices are geometric and topological: device positioning, routing geometry, and topology should be optimized.

A solution method relying on sequential topology and geometry refinement would start with the creation of a layout configuration and continue with iterative adjustment of the chosen route. Such an approach is intuitively understandable and has its practical benefits. Feasibility is guaranteed, the tasks are simplified and independent, reusable algorithms exist. However, it implies a certain design hypothesis. Namely, a specific topology that works well for purely geometrical optimization still stays optimal when physics-driven refinement is introduced. It is easy to see that in dense electro-thermal structures such hypothesis rarely holds true. Minimal route might connect nodes via thermally hot spots and produce significant losses, whereas a more evenly distributed topology might provide less loss at the cost of increased area. Since the combination of both factors impacts performance, topology has to be flexible until optimization is finished.

In the current paper, we are going to answer the following question: is there a way to optimize topology, placement and routing geometry at once with respect to physical constraints while keeping connectivity and manufacturing feasibility intact? Our approach will rely on continuous relaxation of topology space, continuous geometry models, and state analysis coupling electrical and thermal losses. Such an approach does not avoid nonconvexity; instead, it transforms the optimization process by merging several restarts of different routes into one run.

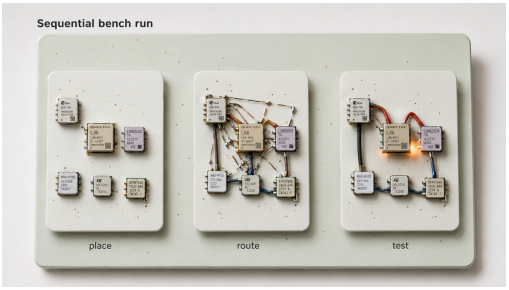

(a) Staged run.

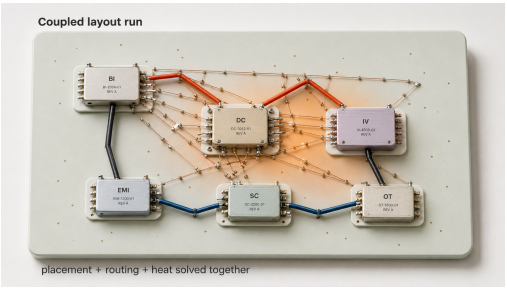

(b) Coupled run.

Figure 1. Layout-search strategy.

In Figure 1, the core of our method is depicted schematically, contrasting a common design process with our approach.

As can be seen in the two panels in Figure 1, however, the suggested contribution involves more than a mere quantitative enhancement in performance but a new way of designing circuits. In the staged case, the thermal assessment comes after the design has taken on a definite topological flavor. In the coupled process, however, the developing design board, interconnects, and heat field are still all exposed to the optimizer, making the testing of the posed research question possible as a topology–geometry design issue.

3. Unified Differentiable Formulation

3.1. Candidate topology graph and design variables

Let $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ represent a candidate routing graph defined within the package design space. The vertex set $\mathcal{N}$ comprises both the device pins and the optional Steiner nodes, while the edge set $\mathcal{E}$ comprises candidate routing paths. Given a set of N_n nets, the current from at least one source point on net $m \in \{1, \ldots, N_n\}$ must reach at least one sink point via its active edges.

The entire design vector is

$$\mathbf{y} = \left[\mathbf{p}_1, \theta_1, \ldots, \mathbf{p}_{N_d}, \theta_{N_d}, \mathbf{c}_1, \ldots, \mathbf{c}_{N_e}, w_1, \ldots, w_{N_e}, z_1, \ldots, z_{N_e}\right]^T, \tag{1}$$

where $\mathbf{p}_i = [x_i, y_i]^T$ is the location of device i, θ_i is its orientation within the plane, $\mathbf{c}_e$ is the set of free control points describing the shape of edge e, w_e is the branch width, and $z_e \in [0, 1]$ is the relaxed activation of edge e. Edges and branches are each parameterized by a Bézier curve, using two interior control points and constant endpoint locations to describe routes that may follow either direct, bent, or slightly detoured paths.

The heterogeneity of this design vector is purposeful. Translations and rotations of the devices determine packing, the Bézier control points specify the shape of routes, the widths determine resistances and manufacturing difficulty, and activations describe the overall topology. Considering all of these in one vector is what enables the optimizer to trade off packing, loss, and temperature in one sweep through the design space rather than multiple sweeps.

3.2. Smooth occupancy representation

To make our approach differentiable, we represent both devices and routing branches using a smooth occupancy function. Specifically, for each primitive j, the occupancy field is described by $\phi_j(\mathbf{r}; \mathbf{y})$, where $\mathbf{r} \in \Omega$. If the primitive is a square or ribbon-shaped branch, the occupancy is

$$\phi_j(\mathbf{r}; \mathbf{y}) = \frac{1}{1 + \exp\left[\gamma\left(d_j(\mathbf{r}; \mathbf{y}) - t_j\right)\right]}, \tag{2}$$

where $d_j(\mathbf{r}; \mathbf{y})$ is a signed-distance-like measure to the primitive centerline or boundary, t_j is the effective half-thickness, and γ controls transition sharpness. For routing edge e, the occupancy is modulated by the activation variable,

$$\phi_e^{(r)}(\mathbf{r}; \mathbf{y}) = z_e \, \tilde{\phi}_e(\mathbf{r}; \mathbf{y}). \tag{3}$$

The aggregate occupancy field is then

$$\Phi(\mathbf{r}; \mathbf{y}) = \sum_{i=1}^{N_d} \phi_i^{(d)}(\mathbf{r}; \mathbf{y}) + \sum_{e=1}^{N_e} \phi_e^{(r)}(\mathbf{r}; \mathbf{y}). \tag{4}$$

This commonly used representation of a field allows for control of overlap, placement of the thermal source, and calculation of the footprint through one single model based on a differential geometry function. The interpretation of Φ is that the optimization routine is now seeing the physical occupation as an image with continuous gradients, and not as a collection of geometrical tests.

3.3. Geometric feasibility and manufacturability

Instead of relying on cumbersome pairwise non-overlap checks, the formulation penalizes excessive combined occupancy through

$$C_{\mathrm{ov}}(\mathbf{y}) = \int_{\Omega} [\mathrm{softplus}(\Phi(\mathbf{r};\mathbf{y}) - 1)]^2 \, \mathrm{d}\Omega, \tag{5}$$

with

$$\mathrm{softplus}(s) = \frac{1}{\beta} \ln\Big(1 + e^{\beta s}\Big). \tag{6}$$

When total occupancy is greater than one, the cost becomes very large, making the layout retreat to the condition of no interference. In contrast to a rigid collision policy, in this case, the differential at the moment of touching still exists, therefore the algorithm will be able to adjust the layout to avoid any interference without missing gradient data.

The admissibility of boundary positions is provided by the differential inequality

$$g_{\mathrm{bd}}(\mathbf{y}) \leq 0, \tag{7}$$

which constrains device corners, port locations, and Bézier control points to remain in the domain. Width and curvature rules are imposed by

$$w_{\min} \leq w_e \leq w_{\max}, \qquad g_{\mathrm{rad}}(\mathbf{y}) \leq 0, \tag{8}$$

so as to retain manufacturable branch thicknesses and bend radii. The restrictions additionally make sure that any artificial speed-up achieved by the optimization program does not come from unrealistic branch thicknesses or boundary-hugging branches. The routing variables are thus restricted to correspond to a layout that can be regarded as manufacturable electro–thermal module rather than as pure graph theory curve network.

3.4. Connectivity-preserving topology relaxation

The topology variables are constrained physically using graph flow requirements. For every net m, we have $f_e^{(m)} \geq 0$ denoting the current on branch e. Conservation is ensured by

$$\mathbf{A}\mathbf{f}^{(m)} = \mathbf{b}^{(m)}, \qquad m = 1, \ldots, N_n, \tag{9}$$

where $\mathbf{A}$ is the node-edge incidence matrix and $\mathbf{b}^{(m)}$ imposes the prescribed source and sink currents. Flow and topology are coupled through

$$0 \leq f_e^{(m)} \leq I_m^{\max} z_e, \qquad e = 1, \ldots, N_e, \; m = 1, \ldots, N_n. \tag{10}$$

If $z_e = 0$, the branch is inactive and cannot carry current. If z_e is positive, the branch remains available to the network. To encourage discrete final topologies, the formulation uses the binarization measure

$$P_{\mathrm{bin}}(\mathbf{y}) = \sum_{e=1}^{N_e} z_e(1 - z_e). \tag{11}$$

These equations introduce a physical meaning to the relaxed topology, where the activated branch has to support feasible current flow, and the unactivated branch can carry no current at all. It is precisely for this reason that the binary component of the problem penalizes such intermediate values in the design.

3.5. Electro–thermal state model

It is supposed that each active track has thickness t_c. The electrical resistivity of edge e is

$$R_e(\mathbf{y}, T) = \rho_e(T)\,\frac{\ell_e(\mathbf{y})}{w_e t_c + \varepsilon}, \tag{12}$$

where $\ell_e(\mathbf{y})$ is branch length and $\rho_e(T)$ is temperature-dependent resistivity,

$$\rho_e(T) = \rho_{e,0}\big[1 + \alpha_\rho(\bar{T}_e - T_0)\big]. \tag{13}$$

The Joule loss on edge e is then

$$P_e^{(J)} = \sum_{m=1}^{N_n} \big(f_e^{(m)}\big)^2 R_e. \tag{14}$$

Device dissipation is prescribed from component operating data, while edge-wise Joule losses are distributed over the smooth branch occupancy. The steady-state thermal problem is

$$-\nabla\cdot\big(k(\mathbf{r})\nabla T(\mathbf{r})\big) = q_{\mathrm{dev}}(\mathbf{r};\mathbf{y}) + q_{\mathrm{J}}(\mathbf{r};\mathbf{y},T), \qquad \mathbf{r}\in\Omega, \tag{15}$$

with convective boundaries

$$-k\nabla T\cdot\mathbf{n} = h(T - T_\infty) \qquad \text{on } \Gamma_h. \tag{16}$$

The heat-source fields are written as

$$q_{\mathrm{dev}}(\mathbf{r};\mathbf{y}) = \sum_{i=1}^{N_d} \hat{q}_i\,\phi_i^{(d)}(\mathbf{r};\mathbf{y}), \tag{17}$$

$$q_{\mathrm{J}}(\mathbf{r};\mathbf{y},T) = \sum_{e=1}^{N_e} \frac{P_e^{(J)}\,\phi_e^{(r)}(\mathbf{r};\mathbf{y})}{\int_\Omega \phi_e^{(r)}(\mathbf{r};\mathbf{y})\,\mathrm{d}\Omega + \varepsilon}. \tag{18}$$

After finite-element discretization,

$$\mathbf{K}(\mathbf{y})\mathbf{T} = \mathbf{Q}(\mathbf{y},\mathbf{f},\mathbf{T}), \tag{19}$$

which is solved iteratively because the electrical resistance depends on temperature. The discrete equation also clarifies the feedback loop that drives the design changes: moving a component or route changes the conductivity path and heat-source distribution, which changes the temperature field, which then modifies resistance and loss. A topology retained only for geometric convenience can therefore become unfavorable once this loop is evaluated.

3.6. Objective function and constraints

The design objective combines compactness, electrical efficiency, and routing regularity:

$$\min_{\mathbf{y},\mathbf{f},\mathbf{T}} J(\mathbf{y},\mathbf{f},\mathbf{T}) = \omega_A \hat{A}_{\mathrm{bb}}(\mathbf{y}) + \omega_P \hat{P}_{\mathrm{loss}}(\mathbf{y},\mathbf{f},\mathbf{T}) + \omega_L \hat{L}_{\mathrm{tot}}(\mathbf{y}) + \mu_{\mathrm{bin}} P_{\mathrm{bin}}(\mathbf{y}), \tag{20}$$

where $\hat{A}_{\mathrm{bb}}$ is normalized bounding-box area, $\hat{P}_{\mathrm{loss}}$ is normalized total interconnect power loss, and $\hat{L}_{\mathrm{tot}}$ is a modest length regularizer. In the reported study, $\omega_A = 0.50$, $\omega_P = 0.40$, and $\omega_L = 0.10$.

The optimization is subject to the coupled state equation, graph-flow conservation, overlap control, thermal admissibility, manufacturability, and topology bounds:

$$\mathbf{K}(\mathbf{y})\mathbf{T} = \mathbf{Q}(\mathbf{y}, \mathbf{f}, \mathbf{T}), \tag{21}$$

$$\mathbf{A}\mathbf{f}^{(m)} = \mathbf{b}^{(m)}, \qquad m = 1, \ldots, N_n, \tag{22}$$

$$0 \le f_e^{(m)} \le I_m^{\max} z_e, \tag{23}$$

$$C_{\mathrm{ov}}(\mathbf{y}) \le \varepsilon_{\mathrm{ov}}, \tag{24}$$

$$T_i^{\max}(\mathbf{y}, \mathbf{T}) \le T_{i,\mathrm{allow}}, \qquad i = 1, \ldots, N_d, \tag{25}$$

$$g_{\mathrm{bd}}(\mathbf{y}) \le 0, \qquad g_{\mathrm{rad}}(\mathbf{y}) \le 0, \tag{26}$$

$$0 \le z_e \le 1, \qquad e = 1, \ldots, N_e. \tag{27}$$

The optimization statement combines performance objectives with feasibility requirements rather than treating thermal and routing rules as post-processing checks. This matters because the same branch can improve one term and harm another: a wider or longer path may reduce current density yet consume area, while a compact path may increase local temperature. Device-wise peak temperatures are computed using a smooth Kreisselmeier–Steinhauser aggregation over the nodes lying inside each device footprint.

3.7. Continuation-based solution procedure

The early stages of the search benefit from smooth design representations, whereas the final layout should be close to discrete. Accordingly, the optimizer uses continuation on both the occupancy steepness and the binarization penalty,

$$\gamma^{(k+1)} = \eta_\gamma \gamma^{(k)}, \qquad \mu_{\mathrm{bin}}^{(k+1)} = \eta_\mu \mu_{\mathrm{bin}}^{(k)}. \tag{28}$$

At each continuation level, the coupled electro–thermal problem is solved to consistency and a constrained SQP update is applied to the design variables. The continuation schedule first allows broad exploration with soft boundaries and partially active routes, then progressively enforces sharper geometry and near-discrete topology. This sequencing reduces the risk of locking the search into an early routing choice before the thermal consequences of that choice are visible.

Algorithm 1 Unified topology–geometry co-optimization

```
Initialize feasible device positions, candidate graph, branch widths, and relaxed activations
for k = 1, ..., N_cont do
    Update γ^(k) and μ_bin^(k)
    repeat
        Solve graph-flow subproblem for all electrical nets
        Solve electro–thermal finite-element system
        Evaluate objective and constraints
        Compute design sensitivities
        Update y with a constrained SQP step
    until convergence at continuation level k
end for
Threshold low-activation branches and perform a final final geometry refinement
```

The principal computational blocks of the unified formulation are summarized in Figure 2. The diagram emphasizes how candidate topology, smooth geometry representation, graph-flow connectivity, electro–thermal state analysis, and continuation-based SQP updates are treated within one coupled design loop rather than as separate disconnected stages.

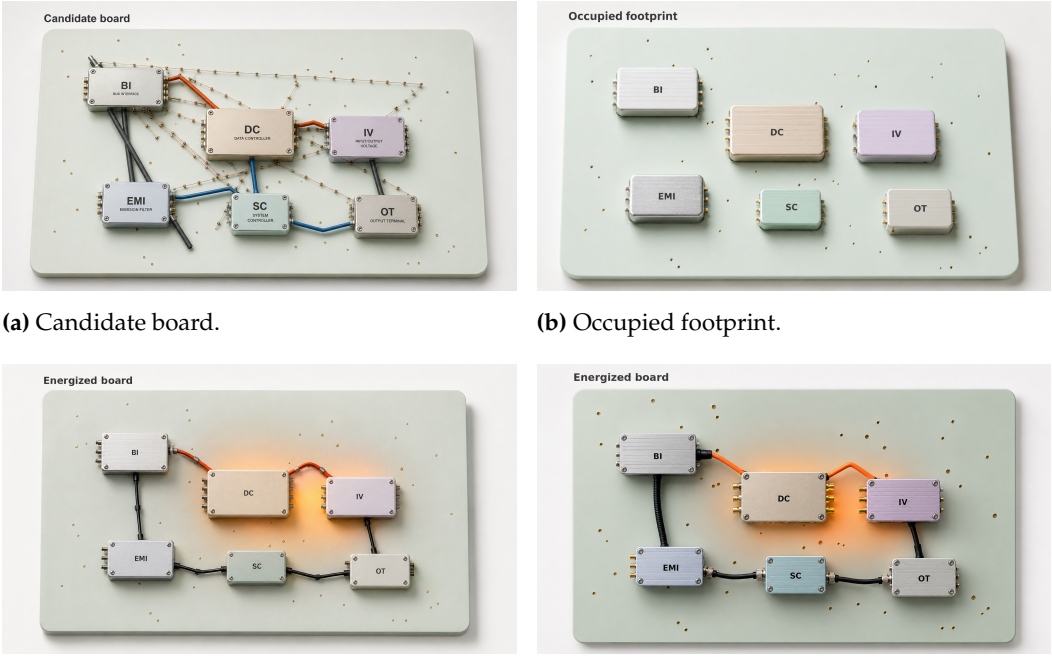

(a) Candidate board. **(b)** Occupied footprint.

(c) Energized state. **(d)** Polished routing.

Figure 2. Coupled design loop.

Figure 2 is meant to be read alongside Algorithm 1. There are more branches in the candidate board than there will be in the solution, but this does not mean that each individual branch is enumerated separately. Each branch is able to battle for inclusion within the same differentiable state, while the occupancy image indicates where real space is being occupied, and the energized image indicates where electrical loading transitions to thermal loading. This visual process shows clearly why the polish step is important for more than just cosmetics – it is the final geometry refinement of an already-chosen topology, under the electro–thermal design criteria.

4. Computational Study Example

4.1. Test system and model parameters

To demonstrate the efficacy of the proposed algorithm, the performance of a typical planar power-electronics module was analyzed. This system had a rectangular package domain of dimensions 240×180 mm, and included six distinct rigid blocks – a battery interface, a DC/DC converter, an inverter driver, an EMI filter, a supervisory controller, and an output terminal block. Three electrical nets were required in total – one for high current supply, one for auxiliary regulated power delivery, and one for low current control/feedback signaling. There were 18 edges and 14 nodes within the candidate routing graph, including ports and optional Steiner points. The candidate graph was deliberately dense enough to allow both direct routes and indirect thermal separations.

Figure 3 gives a compact overview of the example case, showing the package dimensions, the six module components, and the densely-connected routing graph that allows for both types of routing possibility.

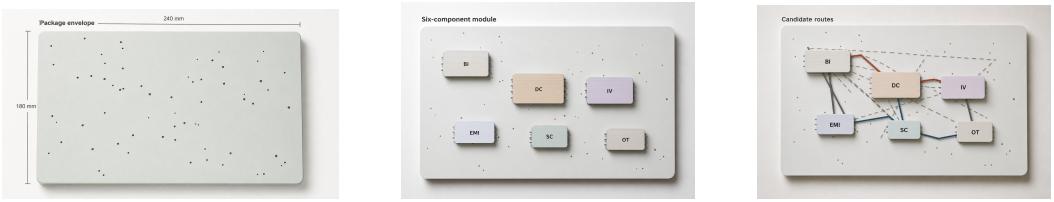

(a) Envelope. **(b)** Components. **(c)** Candidate routes.

Figure 3. Computational test article.

Figure 3 establishes the physical scale for the numerical test case prior to any optimization results. The package is small compared to the six heat-emitting components, such that an optimization favoring only compactness will result in dense high-current routing paths within the central region. The proposed graph will be rich enough to provide multiple alternative connections for trade-offs between route length and low thermal congestion.

Substrate thermal conductivity was chosen as $\kappa_0 = 7.0\ \mathrm{W/(m\,K)}$, with a conductor thickness of $t_c = 1.2\ \mathrm{mm}$, ambient temperature of $T_a = 35°\mathrm{C}$, and top and bottom boundary convective coefficient of $\alpha_{b,t,b} = 18\ \mathrm{W/(m^2\,K)}$. The two side boundaries were adiabatic. The reference resistivity was $\rho_r = 1.75 \times 10^{-8}\ \Omega\mathrm{m}$, while the temperature coefficient of resistivity was set at $\alpha_\rho = 0.0039\ \mathrm{K}^{-1}$. The number of elements for the finite-element analysis was 12,480 bilinear quadrilaterals, yielding mesh-independent behavior for the reported figures. The initial design ensured non-overlapping and set a low activation value of $z_e = 0.35$ for all candidate edges.

Continuation from three steps was done, such that the values of $(\gamma, \mu_{\mathrm{bin}})$ were increased from $(12, 0.05)$ to $(24, 0.25)$ and finally to $(48, 1.00)$. Upon termination of the last stage, removal of all edges with $z_e < 0.15$, as well as a final geometry refinement, were performed. All optimizations use identical constraints, current requirements, and temperatures.

The relevant component data are presented in Table 1. Table 1 provides details on the main thermal imbalance within the test article. The DC/DC converter and inverter driver are the main sources of heat in the test board, while the supervisory controller and terminal block are less thermally intense components, but still partake in congestions. Hence, the optimization should not treat all components equally as boxes but rather consider their heat loads and constraints to define which neighborhoods are expensive for high current connections.

Table 1. Representative electro–thermal module data used in the computational study.

Component	Footprint (mm)	Heat generation (W)	Allowable temperature (°C)
Battery interface unit	42 × 28	18	105
DC/DC converter	48 × 32	34	110
Inverter driver	40 × 28	26	105
EMI filter	36 × 24	12	95
Supervisory controller	28 × 20	8	90
Output terminal block	30 × 18	5	95

4.2. Reference methods and performance criteria

Two approaches for reference solutions have been employed. In the first one, referred to as *best-of-many sequential*, a typical iterative process is reproduced in an offline manner. Namely, the collection of all topologies satisfying the feasibility constraints is generated offline, and then the geometry optimizer is run on each of those feasible topologies. The best result out of 24 feasible topologies is kept. In the other reference approach, named *fixed-topology geometry refinement*, the topology with the highest score among enumerated topologies is taken as initial configuration and then optimized further.

Evaluation is performed based on bounding box area, total interconnect loss, maximal device temperature, total routed lengths of active traces, number of active edges, and thermal safety margin to the most constrained component. Thus, the efficiency can be studied from perspectives of compactness, electrical performance, and thermal safety margins at the same time.

4.3. Results of comparison

Table 2 presents the main results of the work. It shows that the unified differentiable approach demonstrated better performance compared to both references for all metrics. With respect to the sequential best-of-many reference, the area of the bounding box was lowered by 17.6%, interconnect power loss by 22.5%, and maximum device temperature by 8.7°C. In relation to the second approach of geometry optimization based on the reference topology, the gains in terms of the three aforementioned metrics were 12.4%, 13.7%, and 4.4°C, respectively.

This increase does not come at the expense of other metrics' performance. The final network not only shrinks in size but also is less dissipative in an electrical sense, as well as safer in thermal terms. In addition, considering that the final topology consists of eight active edges only, we can conclude that it did not simply redistribute the currents among more branch lines. On the contrary, it found a more sophisticated topology with branches having maximum electro-thermal performance only.

Table 2. Comparison of final design performance for the representative electro–thermal module.

Method	Area (cm^2)	Loss (W)	Max. temp. (°C)	Active edges	Route length (cm)
Sequential best-of-many	410.4	18.7	103.4	11	118.6
Fixed-topology geometry refinement	386.2	16.8	99.1	10	109.1
Unified differentiable method	**338.2**	**14.5**	**94.7**	**8**	**96.4**

An analysis of the results of the thermal effects is presented in Table 3 below, where the ultimate temperatures of the designed components are presented. The DC/DC converter still proved to be the crucial element, however, its operating temperature was maintained at 94.7°C; as a result, it still had 15.3°C left from the permissible operating temperature. All other elements functioned with larger margins due to their adequate placement.

Table 3. Final device-wise temperature distribution for the unified design.

Component	Peak temperature (°C)	Allowable limit (°C)	Thermal margin (°C)
Battery interface unit	86.4	105	18.6
DC/DC converter	94.7	110	15.3
Inverter driver	90.1	105	14.9
EMI filter	79.8	95	15.2
Supervisory controller	68.5	90	21.5
Output terminal block	63.2	95	31.8

This solution also exhibited an interesting topological rationale. During the first continuation stage, a few supply branches continued to operate partially, thus spreading the current and thermal load over a larger portion of the package. Continuing further with the optimization, branches coming too close to the DC/DC converter and inverter driver were successively dropped since their compact geometry would not compensate for increased temperatures. The control net was, on the other hand, practically kept direct in the course of optimization, since it carried insignificant currents. Thus, there was competition between two factors: compact and low-resistive topology for the branches carrying small currents on one side, and the topology ensuring thermal distribution and avoiding crowding on another branch.

This routing strategy is depicted graphically alongside with the steady-state temperature field in Figure 4. From this figure, it is obvious that the resulting layout avoids the most crowded parts of the network at the cost of increasing length of branches carrying large currents.

(a) Final layout. **(b)** Thermal field.

Figure 4. Optimized module response.

The two figures in Figure 4 provide a connection between the numerical criteria and the physical layout mechanism. Instead of using the shortest path for the high current channel, this path meanders around the hot core area, whereas the low current one continues being more straight. In fact, the temperature picture indicates that the routing criterion has taken into account the heat map, and not been done afterwards, and thus the main reason why there is lower loss and temperature in the unified approach.

The performance improvement with respect to each metric obtained through the unified approach is shown in Figure 5.

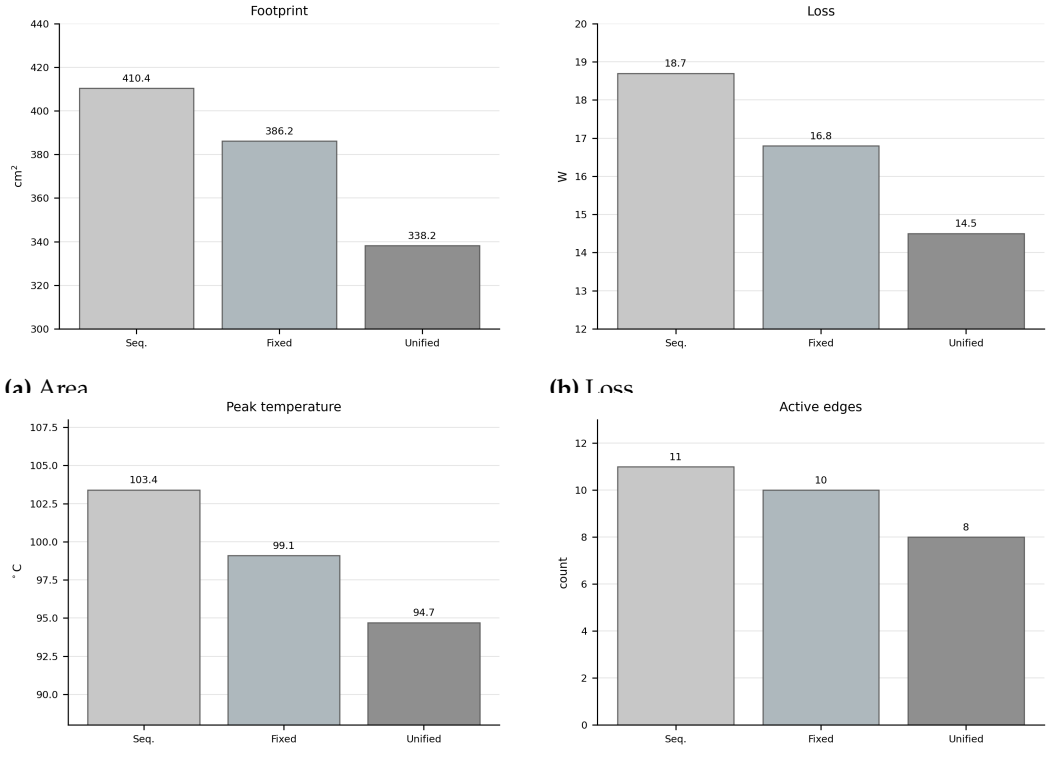

(a) Area **(b)** Loss

(c) Temperature. **(d)** Edges.

Figure 5. Performance comparison.

The multi-objective aspect of the proposed method is made clear in Figure 5. If a more compact architecture leads to greater joule heating, or a more efficient architecture

has worse thermal margin, neither would be preferable by themselves. Conversely, a cooler solution that needs an unnecessarily complex architecture is less desirable than one that achieves the same temperature reduction with a simpler circuit structure.

4.4. Effect of continuation and topology relaxation

Continuation and topology sharpness play a role as shown in Table 4. Without continuation but with topology relaxation, it could already outperform the baseline method of fixed topologies. This proves that joint topology and geometric optimization is already advantageous regardless of whether the binary variable is highly defined yet or not.

Table 4. Ablation study for the proposed optimization strategy.

Variant	Area (cm^2)	Loss (W)	Max. temp. (°C)	Mean fractionality
Fixed-topology geometry refinement	386.2	16.8	99.1	0.000
Unified relaxation without continuation	349.1	15.3	96.8	0.083
Unified relaxation with full continuation	**338.2**	**14.5**	**94.7**	**0.012**

From Table 4, one sees that the improvement is not obtained merely through the use of additional continuous parameters. While the relaxed design with no continuation helps improve upon the reference design strategy, it still results in an obviously fractured topology. Continuation further provides an added benefit, as it allows for improvement in design quality while simplifying the network into a cleaner route set.

Table 5 summarizes the continuation process. With the increase in sharpness of the occupancy field and in binarization weight, the design volume was reduced, along with the total length of routes carrying current and the number of active edges, which went from 13 to 8. Not all stages saw equal reductions in electrical loss, as the biggest reduction occurred during stage two, where the high-current route was steered away from its most congested thermal area.

Table 5. Evolution of the unified optimization across continuation stages.

Stage	γ	μ_{bin}	Area (cm^2)	Loss (W)	Max. temp. (°C)	Eff. active edges
Initial feasible design	–	–	451.8	21.6	108.9	18
1	12	0.05	371.4	16.4	98.6	13
2	24	0.25	347.6	15.1	95.8	10
3	48	1.00	340.1	14.6	94.9	9
Final polished design	–	–	338.2	14.5	94.7	8

Mean fractionality is defined by

$$\bar{F} = \frac{1}{N_e} \sum_{e=1}^{N_e} z_e(1 - z_e), \tag{29}$$

which approaches zero as the topology approaches a binary state. The decrease in $\bar{F}$ and in the effective edge count demonstrates that continuation served not only as a numerical aid, but also as a design-sharpening mechanism.

Figure 6 visualizes the continuation-driven evolution of the layout and the four monitored performance metrics. Together with Table 5, these panels show how the design becomes progressively smaller, less dissipative, cooler, and topologically simpler as the continuation parameters are tightened.

Figure 6 presents the function of continuation as a means for refining the solution toward sharp designs. The early steps preserve many provisional branches, thus preventing premature disconnection until the positions of the device and route corridors have been settled. As the penalty on binarization increases, low-value branches are pruned, resulting

in a physical layout, where the path with strong current is kept thermally isolated, whereas the one with weak current is made as short as possible.

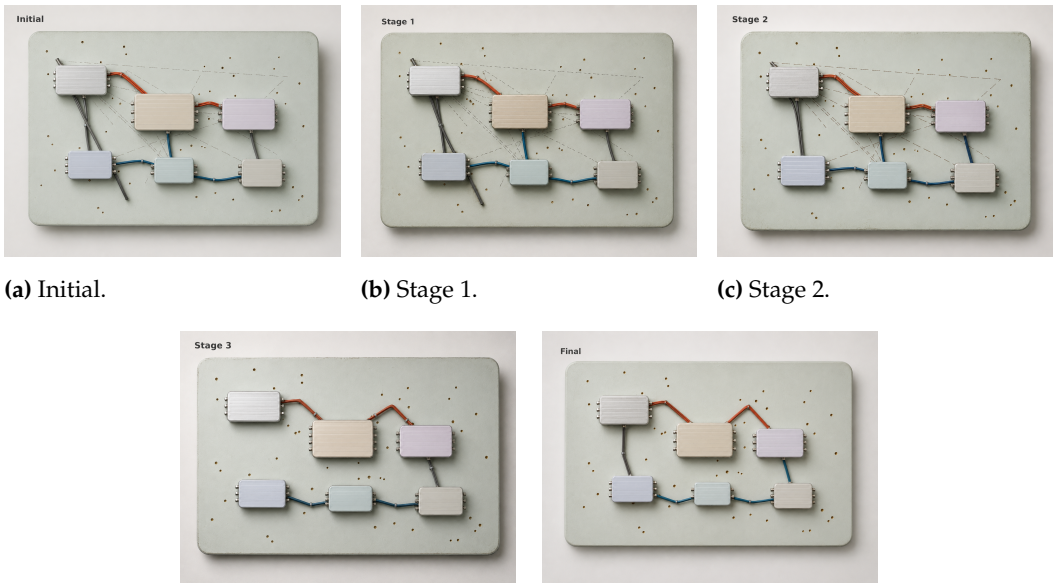

(a) Initial. **(b)** Stage 1. **(c)** Stage 2.

(d) Stage 3. **(e)** Final.

Figure 6. Continuation sequence.

4.5. Computational behavior

The computations were performed with the finite-element and SQP solver implemented in MATLAB on a computer with a 3.2 GHz multi-core processor and 64 GB of RAM. The sequential benchmark comprised 24 geometric optimizations, totaling about 71 minutes elapsed time. The unified algorithm comprised only one continuation run lasting 62 minutes. Even though each step of the unified approach took longer than each optimization run within a particular topology due to a broader range of design options and additional topology variables, the total computational effort was reduced compared to the repeated sequential enumeration.

Such a reduction in computational burden would be more relevant if combined with better designs. The proposed approach not only found a cheaper way to produce the same topology as was obtained through enumeration. It produced a more compact design that consumed less power and generated less heat than even the best enumerated topology. Therefore, the gain from the unification of topology and geometry can be attributed not only to saving computational effort but also to the potential for discovering better topologies.

4.6. Graphical interpretation of the optimized process

For illustration purposes, Figure 2 presents a summary visualization of the coupled topology-geometry optimization process, while Figure 4 interprets the optimal routing logic together with the temperature field in a graphical form. Figure 6 illustrates design evolution under different stages of the continuation process.

5. Discussion

The computational analysis yields some key observations for compact electro-thermal design. Firstly, as expected due to the field coupling, topology and geometry cannot be separated into two independent layers. The optimal topology obtained was neither the shortest possible wiring pattern, nor the most spatially dispersed thermal distribution. It emerged as a trade-off topology that avoided routing of power in the area with high

thermal sensitivity, whereas all low-current paths were left straight. Such solutions are not reachable in a pure fixed-topology workflow because their value emerges only in the solution of the electro-thermal problem itself.

Secondly, it became clear that the concept of smooth occupancy proved itself highly beneficial as a tool of unification of different physical processes. The same representation was used in calculation of the occupancy function and its derivative; definition of the thermal source distribution; evaluation of package compactness; and estimation of wiring occupancy rate. From the numerical viewpoint, the usage of one representation reduces the number of geometric submodels and eliminates their separate differentiation. Moreover, from the methodological viewpoint, it helps to keep consistent with the physical reality of the module: the same physical location contributes to the heat conduction, electrical conductance and geometrical occupation.

Thirdly, the graph-flow relaxation helped to maintain physically relevant topology variations during optimization. One major criticism that can be made to continuous topology variables is that any intermediate step may result in a non-sense value. Here such problem is alleviated by the fact that a branch gets value only if it is able to contribute to the overall flow according to the network constraints imposed on the circuit. As a consequence, there is little reason to retain those branches which are geometrically present but not useful in the electrical sense. This feature is vital for engineering practice, since mathematically convenient relaxation approach becomes applicable only if it converges to a physically meaningful design.

Fourthly, the findings have some implications for mechanical and manufacturing design automation. Nowadays, the combination of compactness, heat transfer considerations and electrical performance plays an increasingly important role in automotive electronics, electrified industrial drive systems, robotics and embedded sensing. These topics are usually addressed using an iterative approach that involves trial and error between layout engineer and thermal/EMI analyst. Differentiable design workflow still leaves no place to pure human engineering intuition, yet it can significantly decrease amount of initial exploration of potential topology variants before simulation-driven analysis begins.

The difference at the level of the workflow between the two design paradigms is illustrated in Figure 7. The key consequence of our approach, made apparent by this comparison, is that instead of cycling between isolated phases of placement, routing, and analysis of physical properties, we can simultaneously update all three in a single loop through differentiability.

(a) Decoupled iteration. **(b)** Differentiable run.

Figure 7. Design-process comparison.

The practical implication is captured in Figure 7, which shows that while the decoupled approach is capable of delivering feasible hardware, it requires multiple re-starts if the

outcome of the thermal simulation reveals bad topology. The differentiable approach replaces the iterative workflow by a monolithic optimization process where the topology is allowed to change prior to finalizing the design.

Limitations should also be mentioned. First, the current work is computational and aimed at showcasing methodology. Steady-state thermal analysis was performed without consideration for switching effects, varying contact resistances, or detailed 3D spreading of heat dissipation across packages. Furthermore, EM compatibility, dielectric break-down, and vibration considerations were omitted even though they may affect routing decisions in practice. Global optimality of the solution is not proven due to the non-convexity of the problem. Initialization and continuation in design will affect the results for very large instances; yet the trends clearly show that the proposed unification is able to capture a design space that cannot be accessed in topology-fixed approaches.

Future work might address multilayered designs, three-dimensional packaging, and more sophisticated electro-thermo-mechanical couplings. Reliability-driven constraints taking into account temperature cycling, material fatigue, and manufacturing tolerance issues will make the proposed method even more useful in practice. Another exciting avenue for future work involves developing surrogate models or learning-based initialization procedures to ensure good scalability when applied to larger graphs and denser components. All these future directions share one common feature: topology should not be fixed at any point during the design cycle.

6. Conclusions

The question was raised as to whether routing topology, component placement, and conductor geometry can be optimized simultaneously within one constrained differential electro-thermal design problem while maintaining connectivity, manufacturability, and accurate heat conduction. The answer appears to be positive with respect to the compact module presented above, whose topology optimization continued until a physics-informed circuit was obtained by incorporating relaxations of edge activations, occupancy field smoothing, graph flow conservation, resistance losses, and steady-state heat conduction.

Numerical results support the above hypothesis. Specifically, the unified approach outperformed the best reference design by reducing its bounding box area by 17.6%, interconnect losses by 22.5%, and peak device temperature by 8.4% under the identical power profile. This improvement was attained without any increase in the number of the active branches; in contrast, the reference layout utilized eleven branches while the optimal layout relied on only eight branches.

The physical interpretation of the result is evident from the final circuit structure and optimization steps: high-current branches survived only when there was sufficient electro-thermal advantage in their inclusion into the final topology while all the branches that came close to the hottest point of the device were gradually removed. Consequently, low-current branches were relatively straight since they had no significant thermal cost associated with their use.

Under the assumption of two-dimensional circuit design, steady-state heat conduction, fixed device loads, and known routing graph, the proposed methodology provides a route to obtaining an electro-thermal solution which is compact, energy-efficient, and thermally safe. Its limitation at this stage lies in the fact that the present implementation does not include multilayer routing, transient conduction, detailed packaging physics, electromagnetic compatibility, nor any experimental verification of the results obtained. Future research needs to focus on expanding the applicability range of the proposed approach for the purposes of reliability-centered electro-thermo-mechanical design.

References

[1] Sobieszczanski-Sobieski, J., & Haftka, R. T. (1997). Multidisciplinary aerospace design optimization: survey of recent developments. Structural optimization, 14(1), 1-23.

[2] Rao, S. S. (2019). Engineering optimization: theory and practice. John Wiley & Sons.

[3] Deaton, J. D., & Grandhi, R. V. (2014). A survey of structural and multidisciplinary continuum topology optimization: post 2000. Structural and multidisciplinary optimization, 49(1), 1-38.

[4] Kirkpatrick, S., Gelatt Jr, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. science, 220(4598), 671-680.

[5] Shapiro, J. (1999). Genetic algorithms in machine learning. In Advanced course on artificial intelligence (pp. 146-168). Berlin, Heidelberg: Springer Berlin Heidelberg.

[6] Fruchterman, T. M., & Reingold, E. M. (1991). Graph drawing by force-directed placement. Software: Practice and experience, 21(11), 1129-1164.

[7] Edsger, W. D. (1959). A note on two problems in connexion with graphs. Numerische mathematik, 1(1), 269-271.

[8] Hart, P. E., Nilsson, N. J., & Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. IEEE transactions on Systems Science and Cybernetics, 4(2), 100-107.

[9] Sherwani, N. (1995). VLSI physical design automation. In Algorithms for VLSI Physical Design Automation (pp. 1-35). Boston, MA: Springer US.

[10] Azzouz, R., Bechikh, S., & Ben Said, L. (2016). Dynamic multi-objective optimization using evolutionary algorithms: a survey. In Recent advances in evolutionary multi-objective optimization (pp. 31-70). Cham: Springer International Publishing.

[11] Abdel-Malek, K. A., Yeh, H. J., & Maropis, N. (1998). Determining interference between pairs of solids defined constructively in computer animation. Engineering with Computers, 14(1), 48-58.

[12] Murata, H., Fujiyoshi, K., Nakatake, S., & Kajitani, Y. (2002). VLSI module placement based on rectangle-packing by the sequence-pair. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 15(12), 1518-1524.

[13] Wong, D. F., & Liu, C. L. (1986, June). A new algorithm for floorplan design. In 23rd ACM/IEEE Design Automation Conference (pp. 101-107). IEEE.

[14] Scha¨ fer, M., & Lengauer, T. (2001). Automated layout generation and wiring area estimation for 3D electronic modules. J. Mech. Des., 123(3), 330-336.

[15] Dong, H., Guarneri, P., & Fadel, G. (2011). Bi-level approach to vehicle component layout with shape morphing.

[16] Panesar, A., Brackett, D., Ashcroft, I., Wildman, R., & Hague, R. (2015). Design framework for multifunctional additive manufacturing: placement and routing of three-dimensional printed circuit volumes. Journal of Mechanical Design, 137(11), 111414.

[17] Lasance, C. J. (2008). Ten years of boundary-condition-independent compact thermal modeling of electronic parts: A review. Heat Transfer Engineering, 29(2), 149-168.

[18] Bergman, T. L. (2011). Fundamentals of heat and mass transfer. John Wiley & Sons.

[19] Kahng, A. B., Lienig, J., Markov, I. L., & Hu, J. (2011). VLSI physical design: from graph partitioning to timing closure (Vol. 312). Netherlands: Springer.

[20] Hart, C. G. (2010). Multidisciplinary design optimization of complex engineering systems for cost assessment under uncertainty (Doctoral dissertation, University of Michigan).

[21] Bendsoe, M. P., & Sigmund, O. (2013). Topology optimization: theory, methods, and applications. Springer Science & Business Media.

[22] Guest, J. K., Prévost, J. H., & Belytschko, T. (2004). Achieving minimum length scale in topology optimization using nodal design variables and projection functions. International journal for numerical methods in engineering, 61(2), 238-254.

[23] Rozvany, G. I. (2009). A critical review of established methods of structural topology optimization. Structural and multidisciplinary optimization, 37(3), 217-237.

[24] Gersborg-Hansen, A., Sigmund, O., & Haber, R. B. (2005). Topology optimization of channel flow problems. Structural and multidisciplinary optimization, 30(3), 181-192.

[25] Borrvall, T., & Petersson, J. (2003). Topology optimization of fluids in Stokes flow. International journal for numerical methods in fluids, 41(1), 77-107.

[26] Alexandersen, J., Aage, N., Andreasen, C. S., & Sigmund, O. (2014). Topology optimisation for natural convection problems. International Journal for Numerical Methods in Fluids, 76(10), 699-721.

[27] Sait, S. M., & Youssef, H. (1999). VLSI physical design automation: theory and practice (Vol. 6). World Scientific Publishing Company.

[28] Alpert, C. J., Mehta, D. P., & Sapatnekar, S. S. (Eds.). (2008). Handbook of algorithms for physical design automation. CRC press.

[29] Peddada, S. R., James, K. A., & Allison, J. T. (2021). A novel two-stage design framework for two-dimensional spatial packing of interconnected components. Journal of Mechanical Design, 143(3), 031706.