BLADE: Bi-Level Bayesian Optimization for Metal-Density--Constrained Multi-Layer Package Power/Ground Plane Synthesis

Siyuan Liang[1], Shanyi Li[1], Leilei Jin[1], Yuan Pu[1], Yushen Zhang [2],
Zhen Zhuang[1⋆], Kai-Yuan Chao[3], Ulf Schlichtmann[2], Tsung-Yi Ho[1]
[1] The Chinese University of Hong Kong, [2] Technical University of Munich, [3] Siemens
⋆: Corresponding Author

Abstract

In advanced packages, large power/ground (P/G) planes are essential for power integrity but prone to warpage. While a uniform metal-density constraint combats warpage, mandatory degassing holes complicate adherence to it and degrade IR-drop. We propose BLADE, the first P/G plane synthesis methodology that enforces this constraint while optimizing power integrity. BLADE expands P/G planes from skeletons under metal density control to ensure connectivity, naturally leaving holes for degassing. Guided by a fast IR-drop evaluator, our bi-level Bayesian optimization framework efficiently navigates the design space. Experiments on six SiP cases demonstrate that BLADE simultaneously satisfies the metal-density constraint and achieves superior IR-drop performance.

1 Introduction

Moore's Law has long driven semiconductor advancement [2], yet rising scaling costs are shifting the industry toward advanced packaging solutions [3, 4]. Among these, Fan-Out Wafer-Level Packaging (FOWLP) stands out as a particularly promising technological approach. As illustrated in Figure 1, a typical FOWLP package integrates multiple metal layers dedicated to signal redistribution and power delivery. The RDL interconnects establish essential die-to-die connections [5], while remaining areas form the critical power delivery network (PDN).

Conventional wire-based PDNs cannot meet the demanding electrical requirements of modern packages, where large dimensions cause substantial IR drops [6]. Solid power/ground planes address this by providing continuous low-resistance current paths for minimal voltage drop. These structures additionally offer electromagnetic shielding and enhanced thermal dissipation capabilities [7–9]. However, CTE mismatch-induced warpage necessitates uniform metal density to effectively control thermal deformation [10–12], establishing a fundamental design constraint.

In practice, P/G planes are synthesized after RDL routing completion, resulting in highly irregular layout regions [13]. As shown in Figure 2, manual design struggles in these constrained areas, leading to poor current paths and elevated IR drop that worsens warpage. Moreover, the mandatory degassing holes for packaging processes [14, 15] further fragment planes and degrade IR performance [1], creating a critical co-design challenge spanning power integrity, manufacturing yield, and long-term reliability considerations.

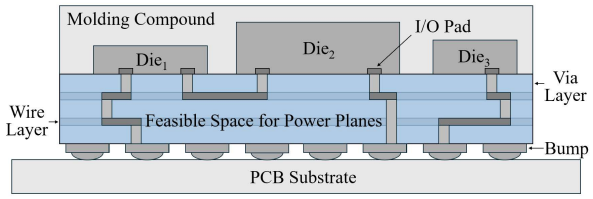

Figure 1: Schematic of 3D IC and multi-layer package power/ground planes.

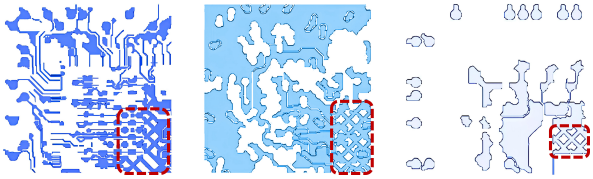

Figure 2: Manually designed power/ground planes with degassing holes, marked by red dashed lines [1].

While automation for signal routing and RDL design has advanced significantly, package P/G plane synthesis remains relatively unexplored. Existing research in PCB power plane design offers valuable insights, particularly SPROUT's grid discretization and graph-search approach [16]. Recent package-oriented methods include: Voronoi diagram construction for efficient plane generation [17], hierarchical genetic optimization with machine learning for copper isolation [18], and multi-agent game-theoretic strategies for balanced IR drop in multi-layer designs [13].

However, these methods face limitations in package applications. Most optimize P/G planes as idealized solid structures, neglecting essential manufacturing constraints like uniform metal density and degassing holes. Post-processing implementations of these requirements typically compromise electrical performance, underscoring the need for an integrated co-design methodology.

In this paper, we present **BLADE**, the first automated P/G plane synthesis tool for packages that simultaneously optimizes IR drop performance and satisfies uniform metal-density constraints. Our framework begins by building immutable skeletal structures to ensure connectivity between power sources and drains. P/G planes are then grown from these skeletons through a density-aware expansion process. The layout is divided into subregions, and expansion in each region halts as soon as a predefined metal density threshold is met. This strategy not only controls warpage but also naturally creates voids that function as degassing holes. To tackle the performance sensitivity to skeleton layout and expansion parameters, we employ a bi-level Bayesian Optimization (BO) framework that co-optimizes these variables hierarchically. Additionally, an efficient IR drop evaluator is implemented to enable fast design iterations.

This work is licensed under a Creative Commons Attribution 4.0 International License.
DAC '26, Long Beach, CA, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2254-7/2026/07
https://doi.org/10.1145/3770743.3803901

The main contributions of this work are as follows:

- We propose BLADE, the first P/G plane synthesis tool that co-optimizes IR drop and reliability while automatically forming degassing holes for manufacturability.
- We develop a bi-level BO framework that jointly improves skeleton design and plane expansion, achieving better performance than conventional BO.
- We introduce a fast IR drop evaluator that accelerates optimization without sacrificing accuracy.
- Extensive experiments show that BLADE delivers superior IR drop performance, scalability, and runtime efficiency compared to existing methods.

2 Preliminaries

2.1 Package Power/Ground Plane

The package PDN is a critical infrastructure responsible for delivering stable power from voltage regulators to integrated circuits, with its performance directly impacting overall system power integrity, signal integrity, and reliability. Unlike on-die PDNs that typically utilize mesh structures [19], package-level implementations predominantly rely on continuous metal P/G planes, which provide superior current distribution and inherent noise suppression.

P/G plane quality is predominantly assessed by DC IR drop performance [20]. This metric is critical as it directly determines the voltage margin at load circuits, where excessive drop can cause timing violations and functional failures [21]. Minimizing IR drop is therefore a primary synthesis objective.

A conventional strategy is to maximize metal coverage within the permissible regions left after RDL routing (blue areas in Figure 1). This approach offers multifaceted benefits: electrically, it reduces current path resistance for lower IR drop and provides electromagnetic shielding [22]; mechanically, it enhances structural rigidity and thermal dissipation [1].

However, this optimization involves fundamental trade-offs. Accommodating multiple independent power nets creates competing demands for space. Furthermore, stringent manufacturing rules, notably the uniform metal-density constraint for warpage prevention and the mandatory degassing holes, directly conflict with unimpeded metal coverage. This complex co-optimization challenge between electrical performance, manufacturing constraints, and reliability necessitates advanced automated synthesis solutions.

2.2 Bayesian Optimization

Bayesian optimization (BO) is a sample-efficient method for minimizing black-box objective functions $f(x)$ that lack analytical expressions or derivatives, requiring direct, and potentially noisy, evaluation. The goal is to find:

$$\hat{x} = \arg\min_{x \in \boldsymbol{X}} f(x), \tag{1}$$

where $\boldsymbol{X}$ is a bounded feasible domain. As illustrated in Figure 3, BO operates through two key components:

- A **surrogate model**, typically a Gaussian Process (GP) [23], that probabilistically approximates $f(x)$. Alternatives such as Random Forest, BART, and BMARS are also used in certain settings [24–26].
- An **acquisition function** that guides sampling toward regions likely to improve upon the current best value f^*. A common

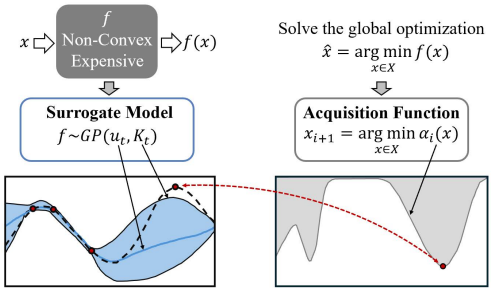

Figure 3: Schematic of Bayesian optimization (BO).

choice is Expected Improvement (EI) [27]:

$$\mathrm{EI}(x) = \mathbb{E}[\max(0, f^* - f(x))], \tag{2}$$

computed under the GP posterior. Variants such as logEI, PI, and LCB are also widely adopted [28, 29].

The BO process iterates as follows: starting from a small initial dataset $\mathcal{D}_t$, the surrogate model $p(y \mid x, \mathcal{D}_t)$ is fit. New candidates are selected by optimizing the acquisition function S, which uses the surrogate as a computationally efficient proxy for f. Each selected point x_{i+1} is evaluated to obtain $f(x_{i+1})$, and the dataset is updated as $\mathcal{D}_{i+1} = \mathcal{D}_i \cup \{(x_{i+1}, f(x_{i+1}))\}$. The model is then retrained, and the loop repeats until an evaluation budget is exhausted. The best point $\hat{x}$ is returned from the final dataset $\mathcal{D}_T$, converging toward the global optimum with high sample efficiency.

3 Overall Flow

The proposed framework automates P/G plane synthesis through a structured two-stage process enhanced by BO. As shown in Figure 4, the flow begins with skeleton construction to establish fundamental connectivity between all power sources and sinks. This is followed by plane expansion where the wire-frame skeletons are iteratively widened into full planes while enforcing uniform metal density constraints. The voids formed during this density-controlled process naturally serve as manufacturing-required degassing holes.

The quality of the final P/G planes is predominantly governed by two critical factors: the skeleton topology and the expansion speed of each plane. Recognizing the hierarchical nature of this dependency, we develop a bi-level BO framework to systematically coordinate these parameters. In the outer BO, guidance point locations loc are determined and treated as compulsory connection vertices during skeleton construction, structurally guiding the creation of superior routing patterns. The inner BO then determines the optimal expansion speed v for each plane, minimizing IR drop variations and enhance overall performance.

The optimization proceeds through tight inter-layer coordination. For each guidance point configuration loc sampled by the outer BO, the inner BO performs comprehensive exploration to identify the corresponding optimal expansion speeds $\hat{v}$, returning the achieved cost $\hat{f}$ to update the outer model. This iterative refinement continues until convergence, ultimately producing optimized parameters $\hat{loc}$ and $\hat{v}$ that collectively enable high-quality P/G plane synthesis.

4 Metal-Density-Constrained Synthesis

4.1 Skeleton Construction

A skeleton is first generated for each P/G plane by routing wires to connect all source and sink vertices, including any compulsory guidance points. This fixed skeleton guarantees persistent connectivity for all subsequent design phases.

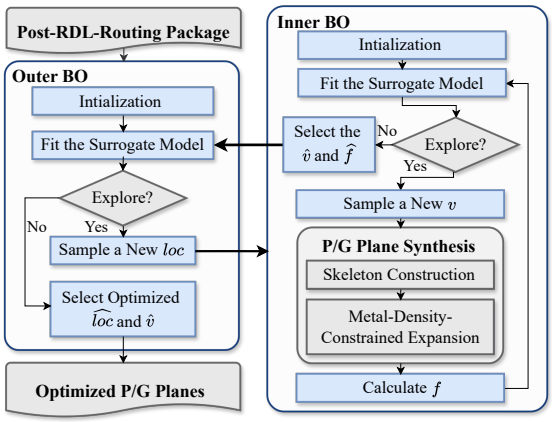

Figure 4: The proposed overall flow of BLADE.

While obstacle-avoiding rectilinear Steiner Minimum Trees (RSMT) offer minimal wirelength, they are computationally expensive and ill-suited for multi-layer designs [30]. Instead, we employ an efficient, Prim-like algorithm [31]. The process begins by connecting an arbitrary source-sink pair using the Lee algorithm [32]. The existing connection then acts as a seed for another Lee-like propagation through available routing space until another vertex is reached. A backtrace step then constructs the shortest path from this new vertex to the seed. This propagate-and-backtrace loop iterates until all vertices are incorporated into the connected skeleton.

Figure 5(b) illustrates the resulting skeletons for the example design after RDL routing (Figure 5(a)). Black lines denote RDL routes and obstacles; blue lines/polygons represent different P/G nets. The fixed skeletons ensure persistent source-to-sink connectivity throughout the following synthesis stages.

4.2 Metal-Density-Constrained Expansion

The wire-frame skeletons guarantee electrical connectivity but exhibit prohibitively high resistances and IR drops due to their minimal cross-sectional area. To transform them into robust, low-resistance P/G planes, we employ an iterative expansion process. is process is governed by a core radial expansion algorithm, manufacturing-aware constraints, and multi-layer integrity rules.

The core algorithm treats the skeleton of a P/G plane as an expansion front. This front expands outward isotropically and uniformly into adjacent, unclaimed space at a given expansion speed, v, which is a key variable optimized by the BO framework. Expansion terminates locally upon encountering package boundaries or pre-allocated routing space from other planes, following a first-come, first-served contention policy.

To enforce uniform metal density and prevent manufacturing defects, we implement a dynamic region-locking system. The design layer is partitioned into a grid of subregions. We monitor the metal fill ratio for each subregion in real time. When a subregion meets a user-defined threshold, for example, the one marked by red dashed lines in Figure 5(c), it becomes locked, halting further expansion within this subregion. This guarantees a globally uniform metal density. Consequently, the voids formed in locked subregions inherently function as the manufacturing-required degassing holes.

For multi-layer scenarios, we enforce two critical vertical integrity rules. First, to support the RDL routes, we project all RDL routes downward as mandatory fill zones. In the most extreme case, i.e., all

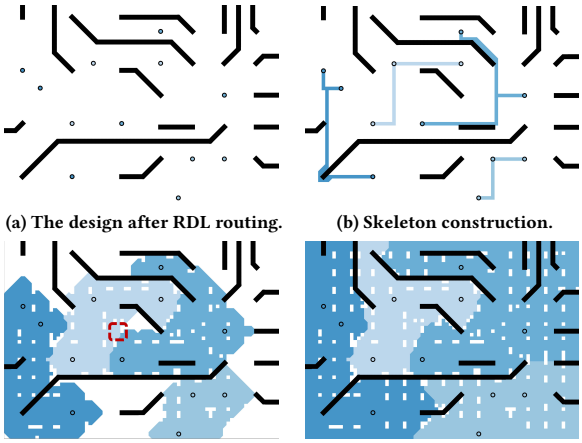

(a) The design after RDL routing. (b) Skeleton construction.

(c) P/G plane expansion. (d) Final P/G planes.

Figure 5: Schematic of the proposed metal-density-constrained P/G plane synthesis.

P/G planes fail to fill a certain zone, we insert dummy metal to ensure mechanical support. Second, to ensure vertical void alignment and prevent unsupported metal, for subregions sharing the same x and y coordinates, the precise void shape of the first locked subregion will be projected to all corresponding subregions across other layers. These projected areas are reserved as unreachable, guaranteeing perfectly aligned degassing holes and structural robustness.

5 Bi-Level BO for IR Drop Mitigation

5.1 Problem Model

The quality of P/G plane synthesis is critically dependent on the skeleton and the expansion speed. We posit that synthesizing planes with superior IR drop performance requires a two-pronged approach: insertion of guidance points $loc : \{\{x_{1_1}, \cdots, x_{n_m}\}, \{y_{1_1}, \cdots, y_{n_m}\}\}$ to shape high-quality skeletons, and the assignment of proper expansion speeds $v : \{v_1, \cdots, v_n\}$ for each plane. Here, n is the number of P/G planes, m is the user-specified number of guidance points for each P/G plane, $\{x_{i_j}, y_{i_j}\}$ is the location of the j-th guidance point for plane P_i, and v_i is the expansion speed of plane P_i.

With a sampled combination of $\hat{loc}$ and $\hat{v}$, we conduct P/G plane synthesis, and evaluate the IR drops IR Drop(P_i) of each plane. Then we treat the IR drop performance as the objective function f:

$$f(loc, v) = \begin{cases} \max\limits_{i,j} \triangle\text{IR Drop}(P_i, P_j) & \text{Successful Synthesis;} \\ \infty & \text{Failed Synthesis,} \end{cases} \tag{3}$$

where $\triangle$IR Drop(P_i, P_j) represents the IR drop variance between plane P_i and P_j. Minimizing this variance enhances reliability and promotes Nash equilibrium [13], reducing global peak IR drop through balanced resource allocation. Our goal is to efficiently find the optimized loc and v within limited number of attempts:

$$\hat{loc}, \hat{v} = \underset{\{loc,v\} \in D}{\text{argmin}}\ f(loc, v), |D| \leq T. \tag{4}$$

5.2 BO Configuration

Bi-Level Optimization. The co-optimization of guidance point locations (loc) and expansion speeds (v) presents a unique challenge due to their inherent hierarchical dependency and significantly different feasible ranges. Conventional monolithic BO formulation proves particularly unsuitable here, as it would treat these interdependent

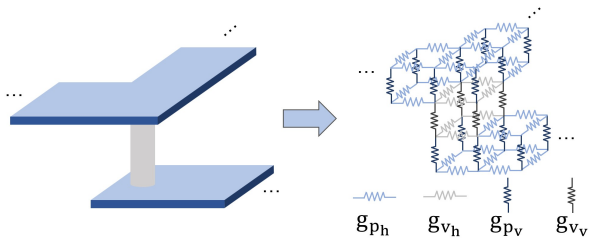

Figure 6: Transformation of P/G plane to equivalent circuit.

parameters as independent variables [33], resulting in both an explosion of the search space dimensionality and severely compromised sampling efficiency.

To address this structural challenge, we formulate the problem as a bi-level optimization task. In this framework, the outer BO selects optimal guidance point locations (loc) to direct the P/G skeleton construction process, while the inner BO, operating conditionally on the established skeleton structures, optimizes the expansion speed (v) assignments. This hierarchical decomposition explicitly captures the fundamental causal relationship between structural configuration (loc) and expansion policy (v), thereby enabling more efficient and effective exploration of the complex design space.

Surrogate Model. The BO of both the levels employs the Gaussian process (GP) as the surrogate model. Assuming the function g is a realization of a GP with mean function μ and covariance kernel K, we denote it as $g \sim \mathcal{GP}(\mu, K)$. The choice of a kernel function K can drastically affect the quality of the surrogate model.

To better control the smoothness of the function being learned, we employ the Matern kernel as the kernel function [34]. It evaluates the similarity between two data points $\mathbf{x}_i$ and $\mathbf{x}_j$:

$$k_\gamma(d) = \sigma^2 \frac{2^{1-\gamma}}{\Gamma(\gamma)} \left(\frac{\sqrt{2\gamma}\, d}{l} \right)^\gamma K_\gamma \left(\frac{\sqrt{2\gamma}\, d}{l} \right). \tag{5}$$

Here, the kernel is parameterized by the Euclidean distance $d = \|\mathbf{x}_i - \mathbf{x}_j\|$, length-scale l, variance σ^2, and smoothness γ. The modified Bessel function of the second kind, K_γ, dictates the covariance decay, blending exponential and polynomial behavior. We select $\gamma = 5/2$, as it offers a computationally efficient closed-form solution and models functions with a realistic degree of smoothness.

Acquisition Function and Optimization Techniques. To efficiently sample from the GP posterior while balancing exploration and exploitation, the Expected Improvement (EI) acquisition function is employed at both optimization levels. Computational efficiency is further enhanced through three key measures: 1) in each round of the outer exploration, only one guidance point is determined for each P/G plane, if more guidance points are needed, more rounds of explorations will be conducted, significantly reducing solution space ; 2) a dynamic boundary calculation that constrains the solution space for guidance point insertion based on source and drain locations; 3) an early termination scheme that skips inner exploration when outer configurations are deemed infeasible.

5.3 Efficient IR Drop Evaluation

We develop a fast SPICE-like IR drop evaluation for each set of parameter during optimization. As shown in Figure 6, we grid the P/G plane and models it as an equivalent circuit. The conductance $g_{i,j}$ between nodes i and j is defined by one of five possible values, corresponding to different physical scenarios: 0 (no connection), g_{p_v} (vertical connection in planes), g_{p_h} (horizontal connection in planes), g_{v_v} (vertical connection in vias), and g_{v_h} (horizontal connection in vias). Based on Kirchhoff's laws [35], we construct the following linear system of the circuit model:

$$\boldsymbol{A}\,\boldsymbol{x} = \boldsymbol{b}, \tag{6}$$

where $\boldsymbol{A}$ represents the conductance among the n nodes, $\boldsymbol{A}_{ii} = \sum_{j!=i}^{n} |g_{i,j}|$. The variable $\boldsymbol{x}$ and the constant $\boldsymbol{b}$ are defined as:

$$\boldsymbol{x} = [V_1, \cdots, V_n]^\top, \quad \boldsymbol{b} = [-I_1, \cdots, -I_n]^\top, \tag{7}$$

where V_i represents the electric potential at node i, and I_i represents the current flowing out from node i. If there are any voltage sources of V_k at node k, then $\boldsymbol{A}_{k,k} = 1$, $\boldsymbol{A}_{k,i} = \boldsymbol{A}_{i,k} = 0$ where $i \neq k$, and $b_k = V_k$.

$\boldsymbol{A}$ is symmetric positive definite due to its irreducibility and weak diagonal dominance [36]. Its sparsity, with a maximum of seven non-zero entries per row from the grid connectivity, enables a highly efficient solution. Consequently, we efficiently solve the equations via sparse Cholesky decomposition [37].

Initially, $\boldsymbol{A}$ is factorized into the product of a lower triangular matrix $\boldsymbol{L}$ and its transpose. To minimize the creation of new non-zero elements during factorization, a permutation matrix $\boldsymbol{P}$ is utilized to reorder the rows and columns:

$$\boldsymbol{P}^\top \boldsymbol{A} \boldsymbol{P} = \boldsymbol{L}\boldsymbol{L}^T. \tag{8}$$

Then the equations can be solved by:

$$\boldsymbol{x} = \boldsymbol{P}\boldsymbol{L}^{\top -1}\boldsymbol{L}^{-1}\boldsymbol{P}^\top \boldsymbol{b}, \tag{9}$$

by maximizing the sparsity of $\boldsymbol{L}$, the solving is much faster than direct solving, which will be shown in Section 6.4.

At the end, the DC IR drop of P/G plane P_i is derived by taking the maximum difference between V_j and V_{source}:

$$\text{IR Drop}(P_i) = \max_j (V_{source} - V_j). \tag{10}$$

Table 1: Benchmark statistics, including the width (W), the height (H), the number of layers (N_L), the number of P/G nets (N_N), the number of P/G bumps (N_B), and the ratio of free space for P/G planes (R).

Benchmarks	W (μm)	H (μm)	N_L	N_N	N_B	R
#1	2500	2500	1	2	9	0.90
#2	2500	2500	1	3	9	0.90
#3	3500	3500	1	4	16	0.89
#4	2500	2500	2	4	17	0.95
#5	3500	3500	2	5	14	0.92
#6	5000	5000	3	5	23	0.92

6 Experimental Results

We comprehensively assesses the proposed methods across P/G plane synthesis and IR drop analysis. We compare different synthesis approaches in terms of IR drop performance and runtime, validate the advantages of our bi-level BO framework over conventional BO, and demonstrate the efficiency of our IR drop evaluation method.

6.1 Experiment Setup

We evaluate BLADE against three state-of-the-art P/G plane synthesis methods: SPROUT [16], which employs grid discretization and graph-search techniques; the Voronoi-diagram-based approach [17], utilizing geometric partitioning via Voronoi tessellation; and the game-theoretic method [13], implementing multi-agent optimization for balanced IR drop in multi-layer designs.

Table 2: Comparison of P/G planes synthesized by different methods.

Benchmarks	SPROUT [16]			Voronoi-Diagram-based [17]			Game-Theoretic [13]			BLADE		
	max IR Drop	△ IR Drop	Time (s)	max IR Drop	△ IR Drop	Time (s)	max IR Drop	△ IR Drop	Time (s)	max IR Drop	△ IR Drop	Time (s)
#1	1.0	1.0	0.45	1.1	1.3	**0.21**	0.85	0.28	51.88	**0.84**	**0.0068**	403.96
#2	1.0	1.0	0.39	0.82	0.71	**0.19**	0.79	0.10	71.29	**0.78**	**0.11**	319.22
#3	1.0	1.0	0.77	1.04	0.62	**0.41**	0.84	0.19	421.48	**0.79**	**0.12**	914.39
#4	1.0	1.0	1.54	1.37	1.5	**1.41**	0.98	0.75	302.17	**0.79**	**0.40**	1774.91
#5	1.0	1.0	**2.62**	0.73	0.63	2.79	0.71	0.40	407.99	**0.70**	**0.31**	2748.70
#6	1.0	1.0	**16.70**	0.98	0.97	18.8	1.0	0.95	5207.06	**0.91**	**0.53**	6752.62
Average	1.0	1.0	**3.75**	1.01	0.96	3.97	0.86	0.45	1076.98	**0.80**	**0.25**	2152.30

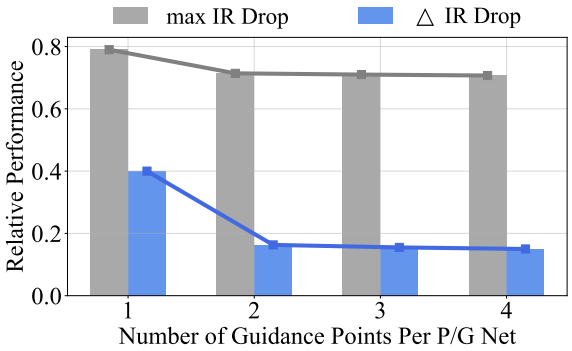

Figure 7: Impact of guidance point count on IR drop performance of benchmark #4.

All baseline methods inherently lack support for both the uniform metal-density constraint and the mandatory degassing holes. To ensure a fair and consistent comparison, we apply uniform post-processing to all baseline results, enforcing metal-density requirements and introducing necessary degassing features absent in their native implementations. This standardized treatment guarantees that all evaluated methods satisfy identical manufacturability standards while maintaining methodological equity.

The experimental setup comprises six benchmarks covering diverse design scenarios, with detailed statistics provided in Table 1. Benchmarks #1 through #5 follow established synthesis methodologies from [38, 39], while Benchmark #6 adapts the pkg1 case from [40] to represent a complex multi-layer industrial system-in-package (SiP) configuration.

All experiments were executed on a computational server equipped with an Intel Xeon Gold 6226R processor and 256 GB of RAM.

6.2 Performance Comparison

For a clear comparative analysis, we use SPROUT's performance as the base and report the relative performance of all methods in Table 2. Here, we utilize two metrics for the IR drop performance: the peak IR drop (max IR Drop), and the IR drop variance across different P/G planes (△IR Drop). Lower values for both metrics indicate superior performance.

IR Drop Performance. As shown in Table 2, BLADE consistently achieves the lowest maximum IR drop and smallest IR drop variation across all benchmark cases. Compared to SPROUT, BLADE reduces the maximum IR drop by 20% and improves the IR drop variance by up to 75%. When evaluated against the most competitive baseline, the game-theoretic approach, BLADE still demonstrates a 7% improvement in maximum IR drop and enhances the IR drop variance by up to 45%.

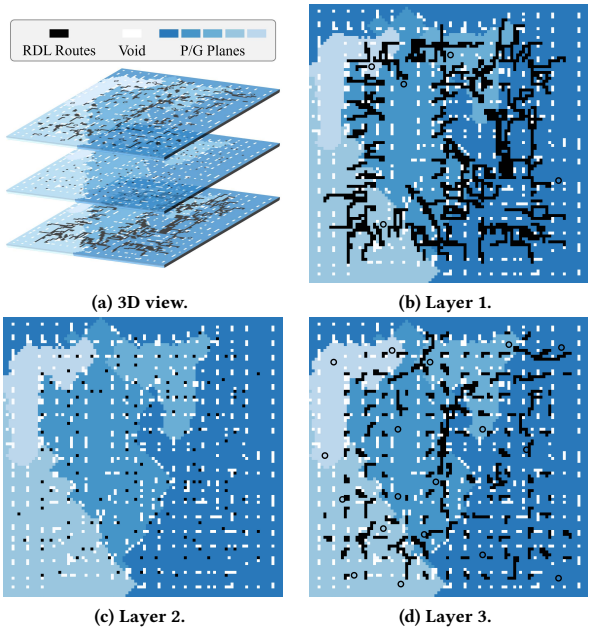

Figure 8: P/G planes synthesized for benchmark #6, fulfilling the uniform metal-density-constraint.

These results were obtained with just one guidance point per P/G net, which is the default configuration of BLADE. As shown in Figure 7, further investigation using benchmark #4 reveals that IR drop performance improves progressively with additional guidance points, reaching optimal performance at four points per net. This flexible framework enables users to balance performance and computational cost effectively. The computational overhead remains manageable due to BLADE's linear time complexity with respect to guidance point count, as each bi-level BO process handles only one guidance point per P/G net.

Scalability Analysis. BLADE maintains consistent performance across varying problem dimensions, from single-layer designs (benchmarks #1-#3) to complex multi-layer arrangements (benchmarks #4-#6). The method effectively handles increasing P/G net counts while adapting to diverse space utilization ratios. For the most challenging benchmark #6, BLADE generates high-quality P/G planes despite dense RDL routing, as shown in Figure 8. The synthesized designs achieve superior IR drop performance while maintaining uniform void distribution through built-in metal density control. Notably, voids across different layers are perfectly aligned vertically, significantly reducing fabrication complexity and ensuring both reliability and effective degassing capability.

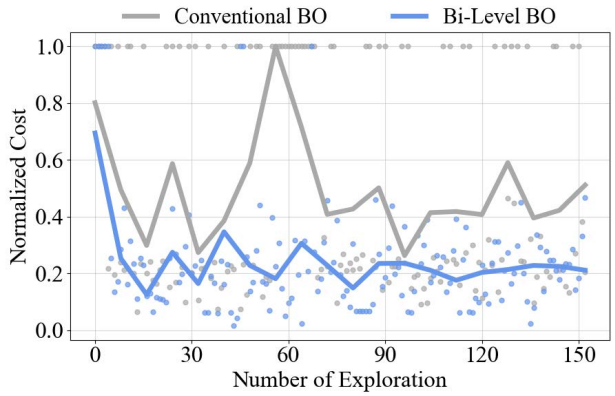

Figure 9: Comparison of the exploration efficiency between conventional BO and the bi-level BO.

Computational Efficiency. BLADE achieves practical computational efficiency with execution times under 10,000 seconds across all benchmarks, performing comparably to the game-theoretic approach while being substantially faster than manual design. Although SPROUT and Voronoi-based methods demonstrate faster execution, BLADE's significantly superior electrical performance justifies the additional computational investment. The framework maintains scalable performance across design complexity, particularly in multi-layer configurations, delivering substantial quality improvements with manageable runtime increases.

6.3 Effectiveness of Bi-Level BO

We evaluate our bi-level optimization framework against conventional BO [33] using benchmark #1 as a representative test case. Results in Figure 9 clearly demonstrate that our hierarchical approach achieves significantly faster convergence and superior final solution quality compared to conventional single-level BO.

The performance improvement originates from the explicit handling of parameter hierarchy within our bi-level framework. While conventional optimization treats all design parameters uniformly within a single search space, our methodology strategically decomposes the problem into two logically separated yet coordinated levels. The upper level specializes in optimizing guidance point locations, while the lower level focuses on determining optimal expansion parameters for each P/G plane.

Experimental validation confirms that our bi-level framework substantially enhances both optimization efficiency and solution quality for P/G plane synthesis. By properly addressing the inherent hierarchical relationships among design parameters, our method enables more thorough design space exploration and ultimately produces superior P/G network designs compared to conventional single-level optimization approaches.

6.4 Efficiency of IR Drop Evaluation

To demonstrate the effectiveness of our method, we benchmark it against two established approaches: the node-based method [41] and the Jacobi-based method [13]. Table 3 presents a comprehensive comparison, demonstrating the superior performance of our proposed approach across various problem scales. Since our method is an accelerated solving approach for the original linear system, it achieves perfect accuracy (ACC = 1.0) consistently across all benchmarks from 200 to 20,000 nodes. This matches the accuracy of the Jacobi-based method while outperforming the node-based.

Table 3: Comparison of different IR drop evaluations

Number of Nodes	Node-based [41]		Jacobi-based[13]		Proposed	
	Time(s)	ACC	Time(s)	ACC	Time(s)	ACC
200	1.00	0.99	0.22	**1.0**	**0.0021**	**1.0**
800	9.92	**1.0**	1.59	**1.0**	**0.0074**	**1.0**
1800	29.79	0.99	3.99	**1.0**	**0.029**	**1.0**
3200	43.17	0.99	2.19	**1.0**	**0.081**	**1.0**
5000	22.46	0.99	1.36	**1.0**	**0.19**	**1.0**
7200	32.48	0.99	2.76	**1.0**	**0.39**	**1.0**
9800	47.30	0.98	6.02	**1.0**	**0.72**	**1.0**
12800	64.06	0.99	10.29	**1.0**	**1.2**	**1.0**
16200	83.65	0.99	15.51	**1.0**	**1.9**	**1.0**
20000	104.94	0.98	22.14	**1.0**	**2.8**	**1.0**

The computational efficiency of our approach shows even more substantial advantages. For the smallest case (200 nodes), the proposed method completes in just 0.0021 seconds—approximately 476× faster than the node-based method and 105× faster than the Jacobi-based approach. This performance advantage persists throughout the scaling range: for 20,000 nodes, our method requires only 2.8 seconds compared to 104.94 seconds (node-based) and 22.14 seconds (Jacobi-based), representing speedups of 37× and 8× respectively.

The scalability of our approach is particularly noteworthy. While competing methods exhibit irregular execution time patterns, with the node-based method showing significant fluctuations and the Jacobi-based method demonstrating non-monotonic scaling behavior, our proposed technique maintains predictable performance growth with increasing problem size. This consistent scaling behavior, combined with perfect accuracy and substantial speed advantages, establishes our method as an effective solution for large-scale DC IR drop analysis in modern circuit design.

7 Conclusion

This paper presents BLADE, a pioneering automated framework for synthesizing power/ground (P/G) planes in advanced packages. BLADE directly addresses the critical co-design challenge of optimizing IR-drop performance while strictly adhering to the uniform metal-density constraint essential for reliability requirement. Our method, based on skeletal structures and region-based constrained expansion, naturally generates manufacturable planes with integrated degassing holes. A novel bi-level Bayesian Optimization framework, accelerated by a fast evaluator, efficiently navigates the complex design space. Experimental results demonstrate that BLADE outperforms existing methods, achieving superior IR drop, enhanced scalability, and greater efficiency. This work provides a foundational step towards fully automated, reliability-aware package design.

8 Acknowledgement

This research was partially conducted by ACCESS-AI Chip Center for Emerging Smart Systems, supported by the InnoHK initiative of the Innovation and Technology Commission of the Hong Kong Special Administrative Region Government. This work is jointly supported by the Research Grants Council of Hong Kong SAR No. CUHK14211324 and No. T46-415/25-R.

References

[1] K. Liu, R. Frye, M. Hlaing, Y. Lee, H. Kim, G. Kim, S. Park, and B. Ahn, "High-speed packages with imperfect power and ground planes," in *The 63rd Electronic Components and Technology Conference (ECTC)*, 2013.

[2] G. E. Moore, "Cramming more components onto integrated circuits," *Proceedings of the IEEE*, vol. 86, no. 1, pp. 82–85, 1998.

[3] Z. Zhuang, B. Yu, K.-Y. Chao, and T.-Y. Ho, "Multi-package co-design for chiplet integration," in *The 41st IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2022.

[4] IEEE, "Heterogeneous Integration Roadmap (HIR)," https://eps.ieee.org/technology/heterogeneous-integrationroadmap/2021-edition.html., 2021.

[5] M.-H. Chung, J.-W. Chuang, and Y.-W. Chang, "Any-angle routing for redistribution layers in 2.5 d ic packages," in *The 60th ACM/IEEE Design Automation Conference (DAC)*, 2023.

[6] J. Mohamed, T. Michalka, S. Ozbayat, and G. R. Luevano, "Pdn design and sensitivity analysis using synthesized models in ddr si/pi co-simulations," in *2018 IEEE Electrical Design of Advanced Packaging and Systems Symposium (EDAPS)*, 2018.

[7] W. Shi and J. Fang, "New efficient method of modeling electronics packages with layered power/ground planes," *Transactions on advanced packaging*, vol. 25, no. 3, pp. 417–423, 2002.

[8] Cadence, "Power plane and ground plane in PCB design: Best practices," https://resources.pcb.cadence.com/blog/2024-power-plane-and-ground-pane-pcb-design-best-practices.

[9] T. Brunschwiler, T. Tick, M. Castriotta, G. Schlottig, D. Gschwend, K. Sato, T. Nakajima, S. Li, and S. Oggioni, "Thermal power plane enabling dual-side electrical interconnects for high-performance chip stacks: Implementation," in *The 5th Electronics System-integration Technology Conference (ESTC)*, 2014.

[10] W. Sun, L. Wang, N. Zhu, J. Xin, Y. Luo, X. Jiang, G. Fan, and M. Chen, "Characterization of packaging warpage, residual stress and their effects on the mechanical reliability of igbt power modules," *Engineering Failure Analysis*, vol. 152, p. 107517, 2023.

[11] M.-Y. Tsai, H.-Y. Chang, and M. Pecht, "Warpage analysis of flip-chip pbga packages subject to thermal loading," *Transactions on Device and Materials Reliability*, vol. 9, no. 3, pp. 419–424, 2009.

[12] S. Chen, S. Li, Z. Zhuang, S. Zheng, Z. Liang, T.-Y. Ho, B. Yu, and A. L. Sangiovanni-Vincentelli, "Floorplet: Performance-aware floorplan framework for chiplet integration," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 43, no. 6, pp. 1638–1649, 2023.

[13] S. Liang, Z. Zhuang, K.-Y. Chao, B. Yu, and T.-Y. Ho, "Multi-layer package power/ground planes synthesis with balanced dc ir drops: A game-theoretic optimization approach," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2025.

[14] Cadence, "Degassing with apd+." https://community.cadence.com/cadence_technology_forums/pcb-design/f/allegro-x-apd/52353/degassing-with-apd.

[15] Simense, "Three ways to meet manufacturing rules in advanced package designs." https://semiwiki.com/eda/siemens-eda/318083-three-ways-to-meet-manufacturing-rules-in-advanced-package-designs/.

[16] R. Bairamkulov, A. Roy, M. Nagarajan, V. Srinivas, and E. G. Friedman, "SPROUT—smart power routing tool for board-level exploration and prototyping," *Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, vol. 41, no. 7, pp. 2263–2275, 2021.

[17] C.-W. Lin, J.-Y. Weng, I.-T. Lin, H.-C. Hsu, C.-M. Liu, and M. P.-H. Lin, "Voronoi diagram-based multiple power plane generation on redistribution layers in 3d ics," in *The 61st ACM/IEEE Design Automation Conference (DAC)*, 2024.

[18] H. Liao, V. Patil, X. Dong, D. Shanbhag, E. Fallon, T. Hogan, M. Spasojevic, and L. Burak Kara, "Hierarchical automatic multilayer power plane generation with genetic optimization and multilayer perceptron," *Journal of Mechanical Design*, vol. 145, no. 10, 2023.

[19] M. O. Hossen, B. Chava, G. Van der Plas, E. Beyne, and M. S. Bakir, "Power delivery network (PDN) modeling for backside-PDN configurations with buried power rails and μ tsvs," *Transactions on Electron Devices*, vol. 67, no. 1, pp. 11–17, 2019.

[20] G. Sisto, B. Chehab, B. Genneret, R. Baert, R. Chen, P. Weckx, J. Ryckaert, R. Chou, G. van Der Plas, E. Beyne *et al.*, "Ir-drop analysis of hybrid bonded 3d-ics with backside power delivery and μ-& n-tsvs," in *International Interconnect Technology Conference (IITC)*, 2021.

[21] M.-S. Zhang and H. Z. Tan, "Ir-drop modeling and reduction for high-performance printed circuit boards," *Electromagnetic Compatibility Magazine*, vol. 4, no. 4, pp. 90–101, 2015.

[22] P. Wilson, "Chapter 2 - printed circuits," in *The Circuit Designer's Companion (4th Edition)*, fourth edition ed., P. Wilson, Ed. Newnes, 2017, pp. 49–91.

[23] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 11 2005.

[24] F. Hutter, H. H. Hoos, and K. Leyton-Brown, "Sequential model-based optimization for general algorithm configuration," in *Learning and Intelligent Optimization: 5th International Conference, LION 5, Rome, Italy, January 17-21, 2011. Selected Papers 5*, 2011.

[25] B. Lei, T. Q. Kirk, A. Bhattacharya, D. Pati, X. Qian, R. Arroyave, and B. K. Mallick, "Bayesian optimization with adaptive surrogate models for automated experimental design," *Npj Computational Materials*, vol. 7, no. 1, p. 194, 2021.

[26] H. A. Chipman, E. I. George, and R. E. McCulloch, "Bart: Bayesian additive regression trees," 2010.

[27] D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient global optimization of expensive black-box functions," *Journal of Global optimization*, vol. 13, no. 4, pp. 455–492, 1998.

[28] S. Ament, S. Daulton, D. Eriksson, M. Balandat, and E. Bakshy, "Unexpected improvements to expected improvement for bayesian optimization," *Advances in Neural Information Processing Systems*, vol. 36, pp. 20 577–20 612, 2023.

[29] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger, "Gaussian process optimization in the bandit setting: No regret and experimental design," *arXiv preprint arXiv:0912.3995*, 2009.

[30] S. Liang, R. Fu, M. Li, T.-M. Tseng, U. Schlichtmann, and T.-Y. Ho, "Raber: Reliability-aware bayesian-optimization-based control layer escape routing for flow-based microfluidics," in *IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2024.

[31] R. C. Prim, "Shortest connection networks and some generalizations," *The Bell System Technical Journal*, vol. 36, no. 6, pp. 1389–1401, 1957.

[32] C. Y. Lee, "An algorithm for path connections and its applications," *IRE transactions on electronic computers*, no. 3, pp. 346–365, 1961.

[33] S. Liang, Y. Zhang, M. Li, T.-M. Tseng, U. Schlichtmann, and T.-Y. Ho, "AutoRE: Bayesian-optimization-based automatic reliability enhancement tool for flow-based microfluidic biochips," in *2025 62nd ACM/IEEE Design Automation Conference (DAC)*, 2025.

[34] M. G. Genton, "Classes of kernels for machine learning: a statistics perspective," *Journal of machine learning research*, vol. 2, no. Dec, pp. 299–312, 2001.

[35] P. Athavale, "Kirchoff's current law and kirchoff's voltage law," *Johns Hopkins University, Retrived*, vol. 6, 2018.

[36] D. M. Young, *Iterative solution of large linear systems*. Elsevier, 2014.

[37] Y. Chen, T. A. Davis, W. W. Hager, and S. Rajamanickam, "Algorithm 887: Cholmod, supernodal sparse cholesky factorization and update/downdate," *ACM Transactions on Mathematical Software (TOMS)*, vol. 35, no. 3, pp. 1–14, 2008.

[38] S.-Y. Hsiao, Y.-Y. Chang, and J.-T. Li, "Automatic die-to-die routing with shielding," in *Proceedings of the 2025 International Symposium on Physical Design*, 2025.

[39] M. Dittrich, A. Heinig, and F. Hopsch, "Electrical characterization of a high speed hbm interface for a low cost interposer," in *2018 IEEE 68th Electronic Components and Technology Conference (ECTC)*, 2018.

[40] H.-T. Wen, Y.-J. Cai, Y. Hsu, and Y.-W. Chang, "Via-based redistribution layer routing for info packages with irregular pad structures," *Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, vol. 41, no. 12, pp. 5554–5567, 2022.

[41] Y. Zhong and M. D. Wong, "Fast algorithms for ir drop analysis in large power grid," in *The 24th IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2005.