

UCLA

Symposium on Physical Design (ISPD)'26

Multi-Level Interconnect Planning for Signal-Power-Thermal Integrity in 2.5D/3D Integration

Siyuan Miao, Lingkang Zhu, Xiangqiao Meng, Wenkai Yang, Chengyu Zhu,
Chen Wu, Lei He

University of California, Los Angeles; BTD Technology Inc.

Contents

- Background and Motivation
 - Chiplets and package interconnects
 - Influence of interconnects on SI/PI/TI
- Our MIP-SPT framework
 - Hierarchical Variable Scheduling Strategy
 - Multi-phase Bayesian Optimization (MPBO)
- Experimental Results

Background: Chiplets and Package Interconnects

- Chiplet: a small chip that contains a well-defined subset of functionality
- Vertical interconnects (VIC):
 - Micro-bumps (μ bumps), C4 bumps and solder balls
 - Through-silicon vias (TSV), through-package vias (TPV) and vias in redistribution layer (RDL)
- Planar interconnects (PIC):
 - Wires in RDL on interposer
 - Planes in RDL on substrate

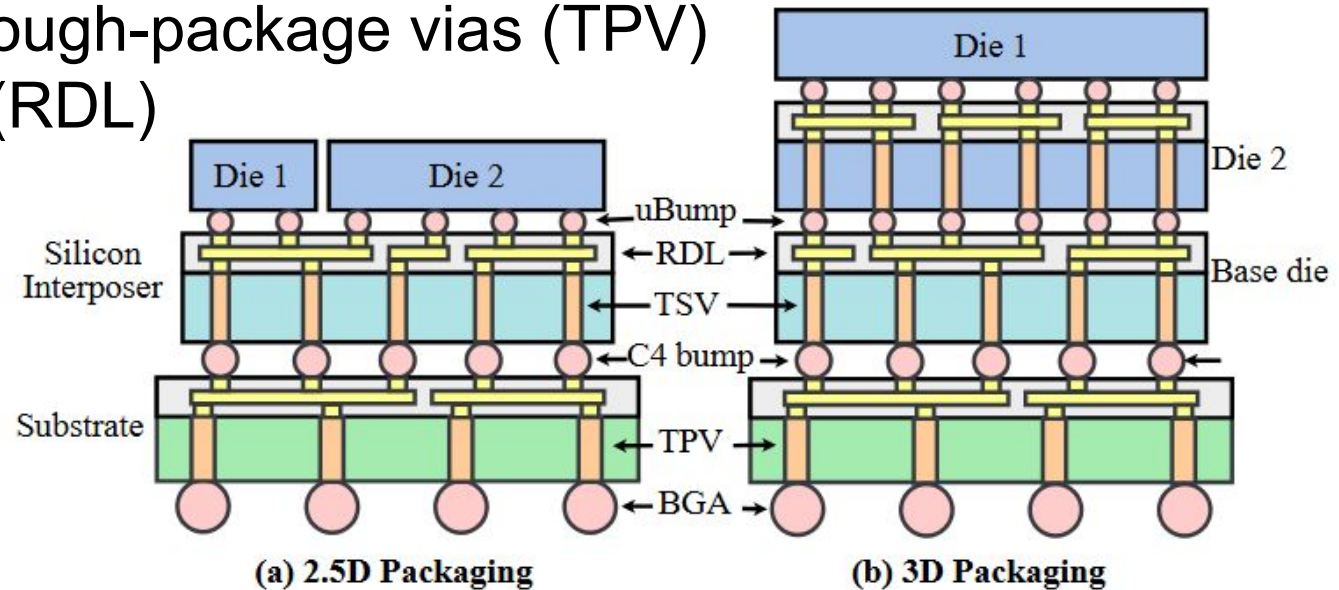


Figure 1: Structure of (a) 2.5D packaging (b) 3D packaging. Some VICs are aligned only for illustration; they can be unaligned when manufacturing.

Motivation: Interconnects Influence PI, SI and TI

- **Die to die SI:** Large number of **parallel** transmission lines -> add **ground guarded** tracks
- **DC PI:** Interconnects are considered as a **resistive network** -> determined by their density and geometric shape
- **TI:** Vertical interconnects form heat transfer channels and become hot spots




Figure 2: Examples of effects on (a) SI (b) PI and (c) TI

MIP-SPT: Challenges in interconnect planning

- The **design space** of interconnects in interposer and substrate is **too large**.
 - 15 dimensions in total, 25 billion combinations -> 25 billion physical simulations
 - Distinct design rules in interposer and substrate
- The multiphysics criteria (SI, PI, TI) and routability requirements are **coupled**.
 - Nonlinearities are introduced in the planning/optimization
 - Lack of quantitative SI/PI/TI analysis in previous works

MIP-SPT: Framework Overview

- We propose MIP-SPT, a multi-phase Bayesian optimization (MPBO) framework for multiphysics and route-driven interconnect planning

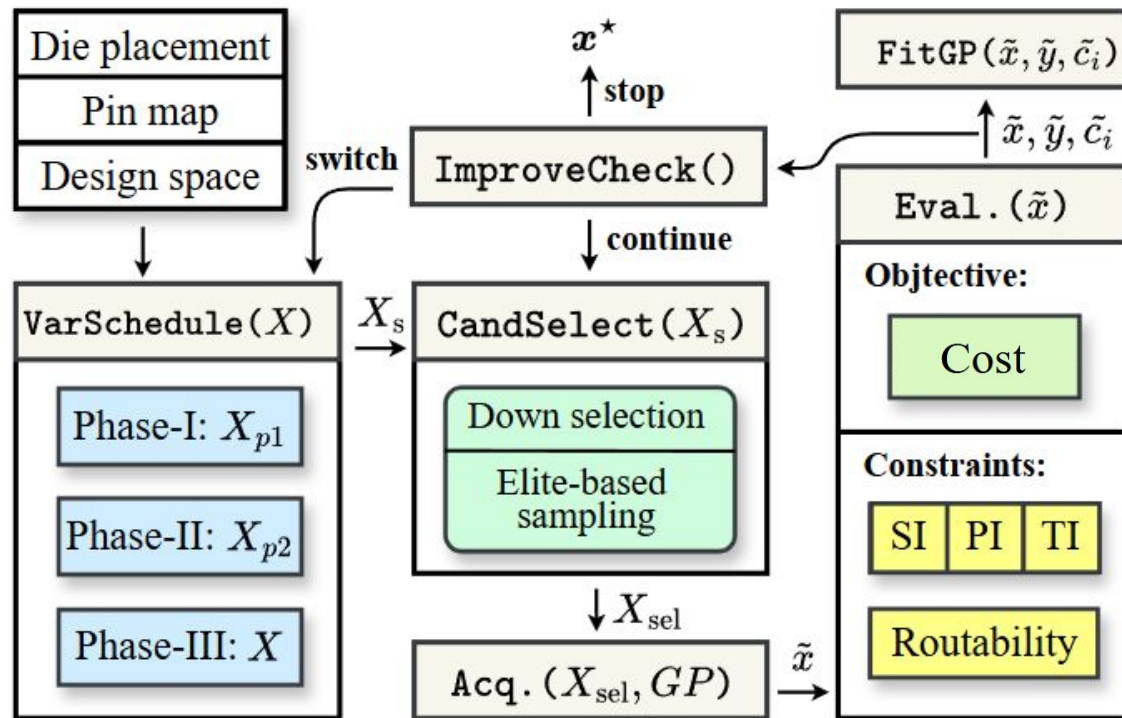


Figure 3. Overview of MIP-SPT framework. "Acq." denotes acquisition

MIP-SPT: Multi-phase Bayesian Optimization

- Objective and constraints

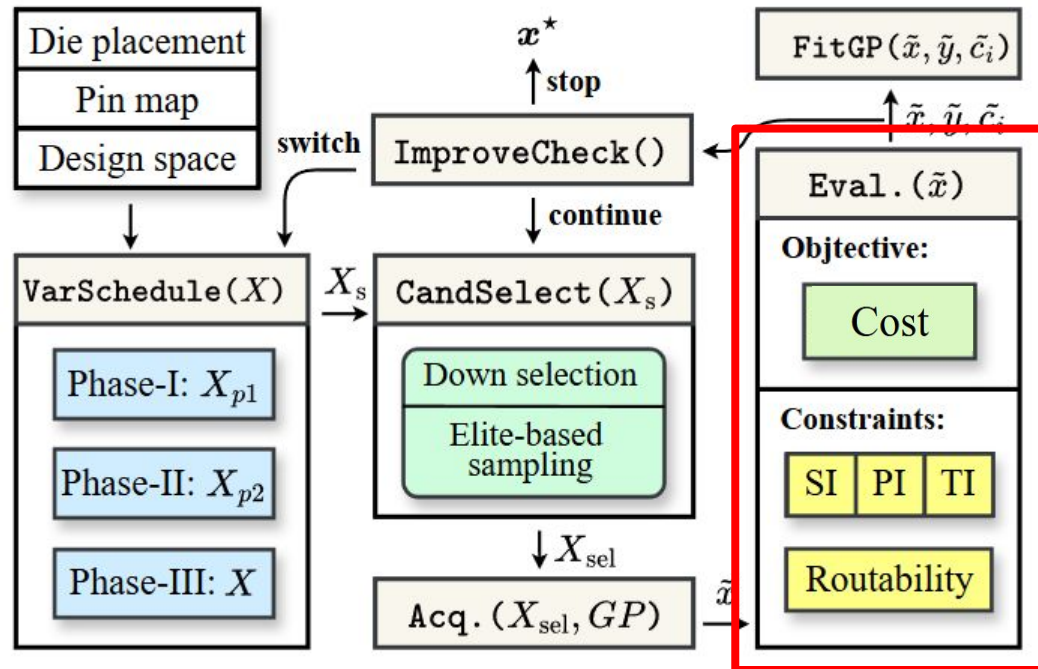


Figure 5. Overview of MIP-SPT framework

- Objective: interconnect cost

- Constraints

- **SI: Insertion loss** of inter-die connection below threshold
- **PI: Maximum IR drop** at die ports below threshold
- **TI: Maximum temperature** below threshold
- **Routability:** the chiplets are **routable** given the optimized design parameters

MIP-SPT: Framework Overview

- Hierarchical Variable Scheduling Strategy

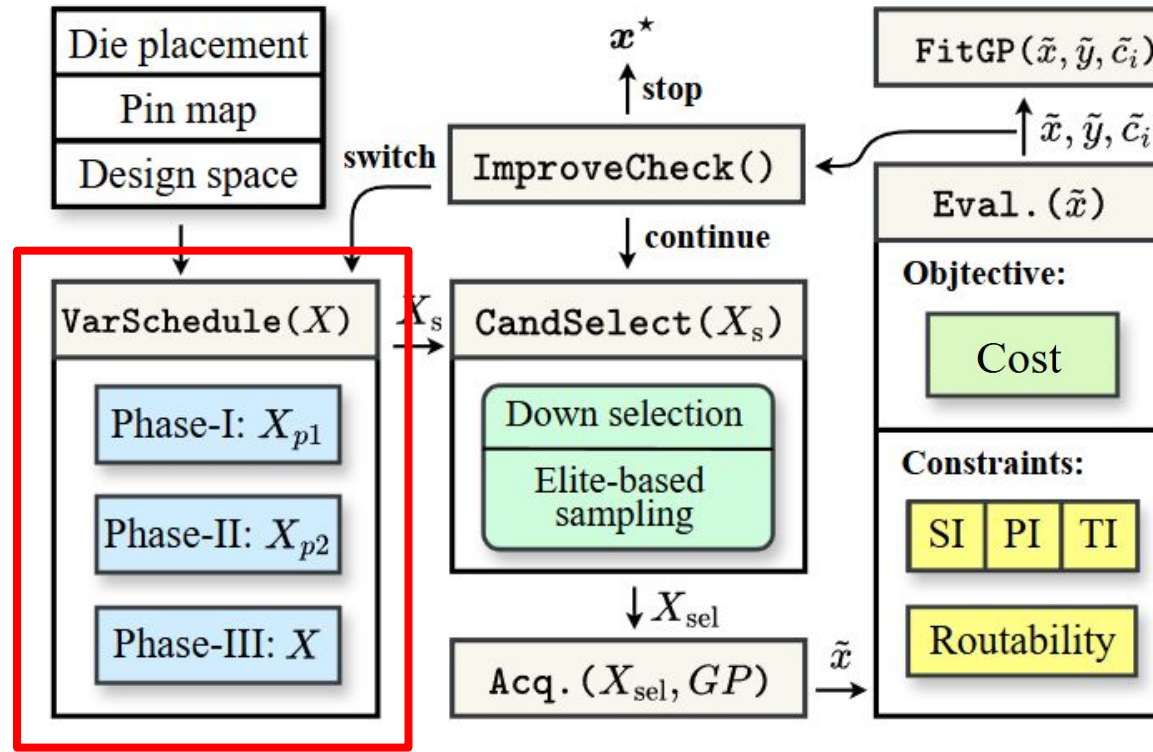


Figure 3. Overview of MIP-SPT framework

MIP-SPT: Hierarchical Variable Scheduling Strategy

- We decouple the design variables across the interposer/substrate based on their physical correlations -> optimize them in separate phases

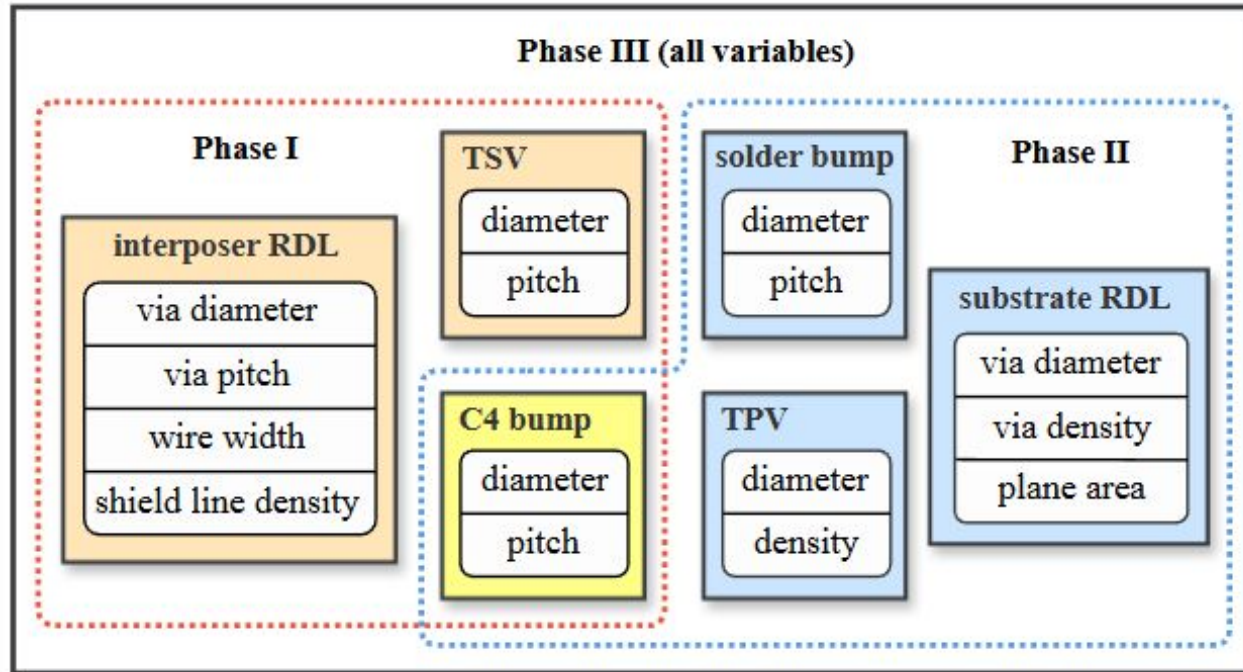


Figure 4: Variables scheduled in different phases

- Phase I: **interposer** and **C4 bump**
- Phase II: **C4 bump** and **substrate**
- Phase III: **all variables**, optimized based on the information from phase I and phase II

MIP-SPT: Hierarchical Variable Scheduling Strategy

- Physical logics behind the hierarchical scheduling
 - **SI**: die-to-die signal connections are routed on the interposer and are irrelevant to the substrate.
 - **PI**: the package PDN can be decomposed into to cascaded sub-networks. The total IR drop is roughly the summation of that in the interposer and substrate.
 - **TI**: similar to PI, the overall thermal resistance can be approximated as a series combination of the interposer and substrate components.

MIP-SPT: Multi-phase Bayesian Optimization

○ Candidate selection

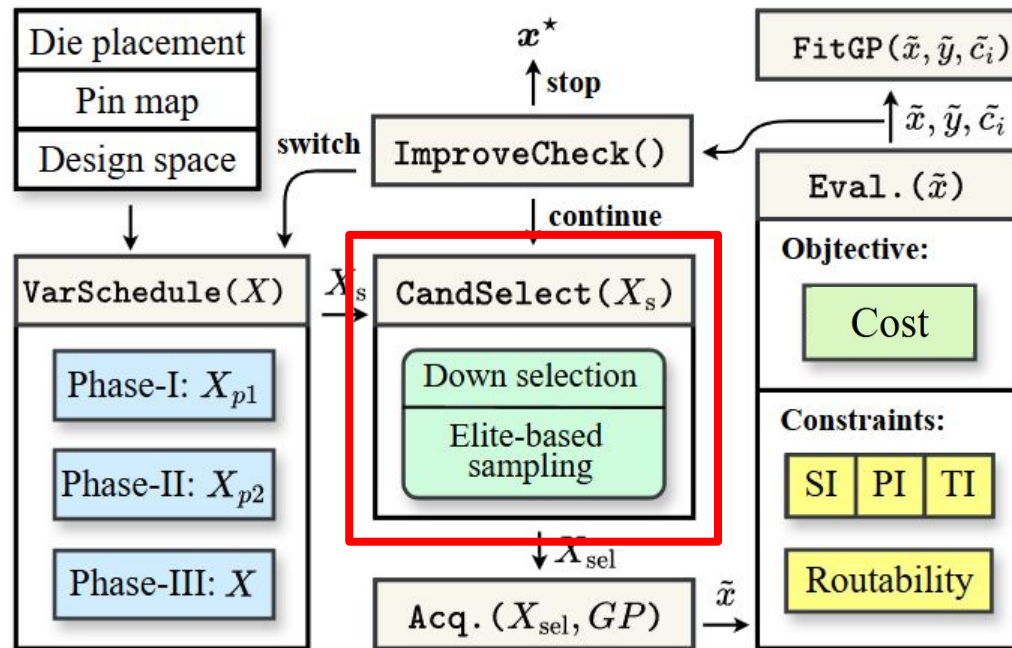


Figure 3. Overview of MIP-SPT framework

- Too many points in design space -> select most promising candidates for acquisition
- Techniques in different phases
 - Phase I and II: multi-score down selection
 - Phase III: down selection + elite-based sampling

MIP-SPT: MPBO - Multi-Score Down-Selection

○ Multi-score down-selection:

- **Exploitation score**: exploits the feasible regions for high objective value

$$S_E(x) = \mu_o(x) P_{\text{feas}}(x), \quad P_{\text{feas}}(x) = \prod_{j=1}^3 \Phi\left(\frac{-\mu_{c_j}(x)}{\sigma_{c_j}(x)}\right)$$

- **Exploration score**: explores uncertain yet promising regions

$$S_A(x) = \mu_o(x) + \beta \sigma_o(x)$$

- **Boundary-learning score**: selects points near the boundary of feasible region

$$S_{c_i}(x) = -\frac{|\mu_{c_i}(x)|}{\sigma_{c_i}(x)}, \quad i = 1, 2, 3$$

- Random selection as fallback

$\mu_o(x), \sigma_o(x), \mu_{c_i}(x), \sigma_{c_i}(x)$: objective and constraint posterior predictions

$\Phi(\cdot)$: cumulative distribution function (CDF) of the standard normal distribution

MIP-SPT: MPBO - Elite-based sampling

- Still too many points in **Phase III** (the entire design space)!
Fortunately, we have historically good feasible points (elites) from Phase I and II, so we **concentrate around the elites** and **sparsely sample the rest of the space**.

- **Elite-based sampling:**

- The distance to the elite set is:

$$D_{\mathbb{E}}(x) = \min_{e_i \in \mathbb{E}} \|x - e_i\|,$$

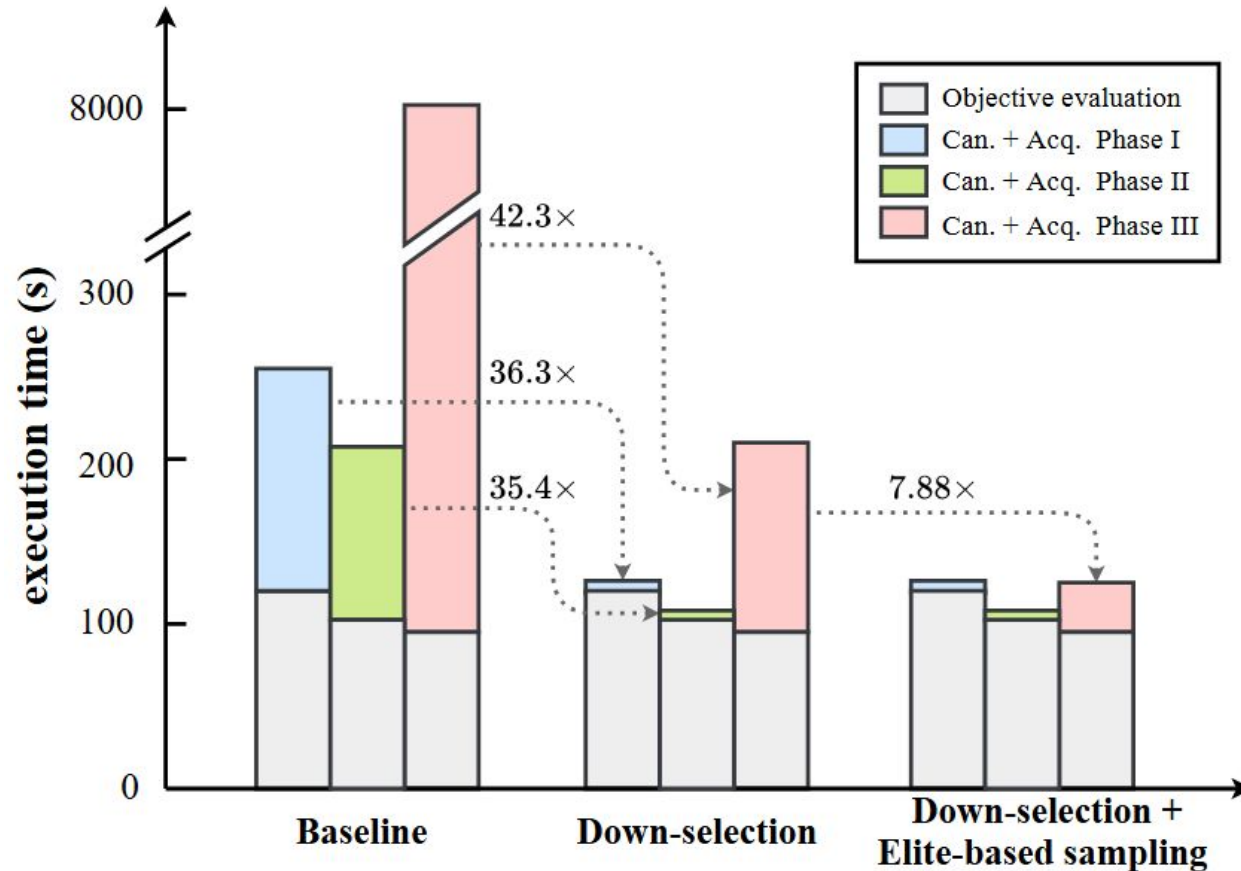
- For each point, the probability of being selected is:

$$q(x) = (1 - \alpha) \frac{\exp(-D_{\mathbb{E}}(x)/\tau)}{\sum_{x_i \in \mathbb{X}} \exp(-D_{\mathbb{E}}(x_i)/\tau)} + \alpha/N,$$

near elites

uniform

MIP-SPT: MPBO – Ablation Study of Candidate Selection



- Multi-score down selection: **38x** speedup on average in Phase I and II
- Elite-based sampling: additional **7.8x** speedup in Phase III

Figure 6: Ablation study of down-selection and elite-based sampling on execution time breakdown of MPBO phases. "Can." denotes candidate selection; "Acq." denotes acquisition.

MIP-SPT: Multi-phase Bayesian Optimization

- Acquisition function (abbreviated as Acq.)

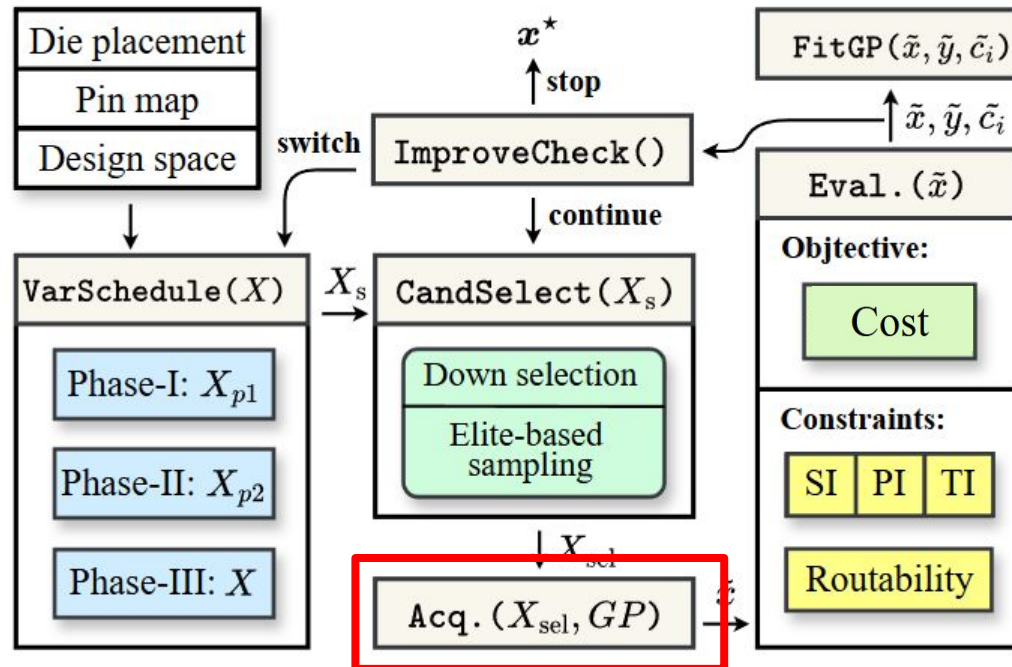


Figure 3. Overview of MIP-SPT framework

- At the beginning with few feasible points:

- probability of feasibility *LogPoF*

$$\text{LogPoF}(x) = \log P_{\text{feas}}(x) = \sum_{j=1}^3 \log \Phi(-\mu_{c_j}(x)/\sigma_{c_j}(x))$$

- Normal cases:

- expected improvement *LogEI*

$$\log \text{EI}(x) = \log \left[(\mu(x) - f^* - \epsilon) \Phi(z(x)) + \sigma(x) \phi(z(x)) \right]$$

$$z(x) = \frac{\mu(x) - f^* - \epsilon}{\sigma(x)} \quad f^* = \max_{j=1, \dots, t} f(x_j)$$

$\Phi(\cdot)$: cumulative distribution function (CDF) of the standard normal

$\phi(\cdot)$: probability density function (PDF) of the standard normal distribution

MIP-SPT: MPBO – Acquisition Function

- Tuning constraint stringency in acquisition function
 - Constraints are considered as a multiplicative weight in acquisition function:

$$W_p = \prod_{j=1}^3 \Phi\left(\frac{-\mu_{c_j}(x)}{\eta\sigma_{c_j}(x)}\right)$$

η controls the softness

- η is determined by **recent feasible rate** r
 - large r -> too conservative -> increase η -> softer constraints
 - small r -> too adventurous -> decrease η -> harder constraints

MIP-SPT: MPBP - Phase Switching and Termination

○ Phase Switching and Termination Criteria

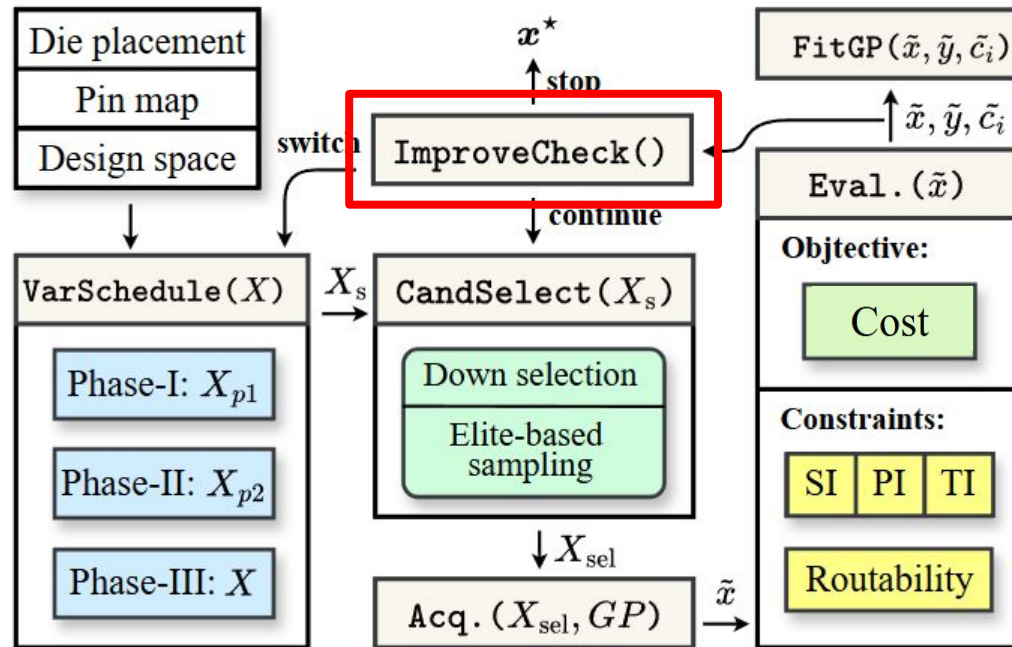


Figure 3. Overview of MIP-SPT framework

- Termination criteria for each phase:
 - Consecutive non-improvement trials $>$ threshold
- Phase switching:
 - Phase I \rightarrow Phase II
 - Phase II \rightarrow Phase III
 - Phase III \rightarrow Stop

MIP-SPT: Experiment setup

- Experiment setup

Architecture:	MCore-OPU ^[4]
Packaging:	Two 2.5D cases + one 3D case
Fixed structural parameters:	Listed below
SI constraints:	Insertion loss < 10dB ($S_{21} > -10$ dB)
PI constraints:	Maximum IR drop < 3%, 5%, 8%, 10%
TI constraints:	Maximum temperature < 343K, 348K

Table 1. Fixed structural parameters of 2.5D/3D cases

Cases	Die layer	μ Bumps		Interposer [†] RDL		TSVs	C4 bumps	Substrate RDL		TPVs	Solder Balls
	height	height	diameter	#layers	height	height	height	#layers	height	height	height
Case-1 (2.5D)	200 μ m	40 μ m	45 μ m	6 layers	120 μ m	200 μ m	100 μ m	4 layers	160 μ m	200 μ m	200 μ m
Case-2 (2.5D)	200 μ m	40 μ m	45 μ m	6 layers	120 μ m	200 μ m	100 μ m	4 layers	160 μ m	200 μ m	200 μ m
Case-3 (3D)	200+720 μ m [‡]	40 μ m	60 μ m	6 layers	120 μ m	200 μ m	120 μ m	4 layers	160 μ m	200 μ m	220 μ m

[†] Base-die in 3D integration shares the same parameters with interposer in 2.5D integration.

[‡] 200 μ m for ASIC, 720 μ m for HBM.

MIP-SPT: Experiment – Comparison with BO baseline

- Comparison between our **MPBO** and **baseline single-phase BO (SPBO)**

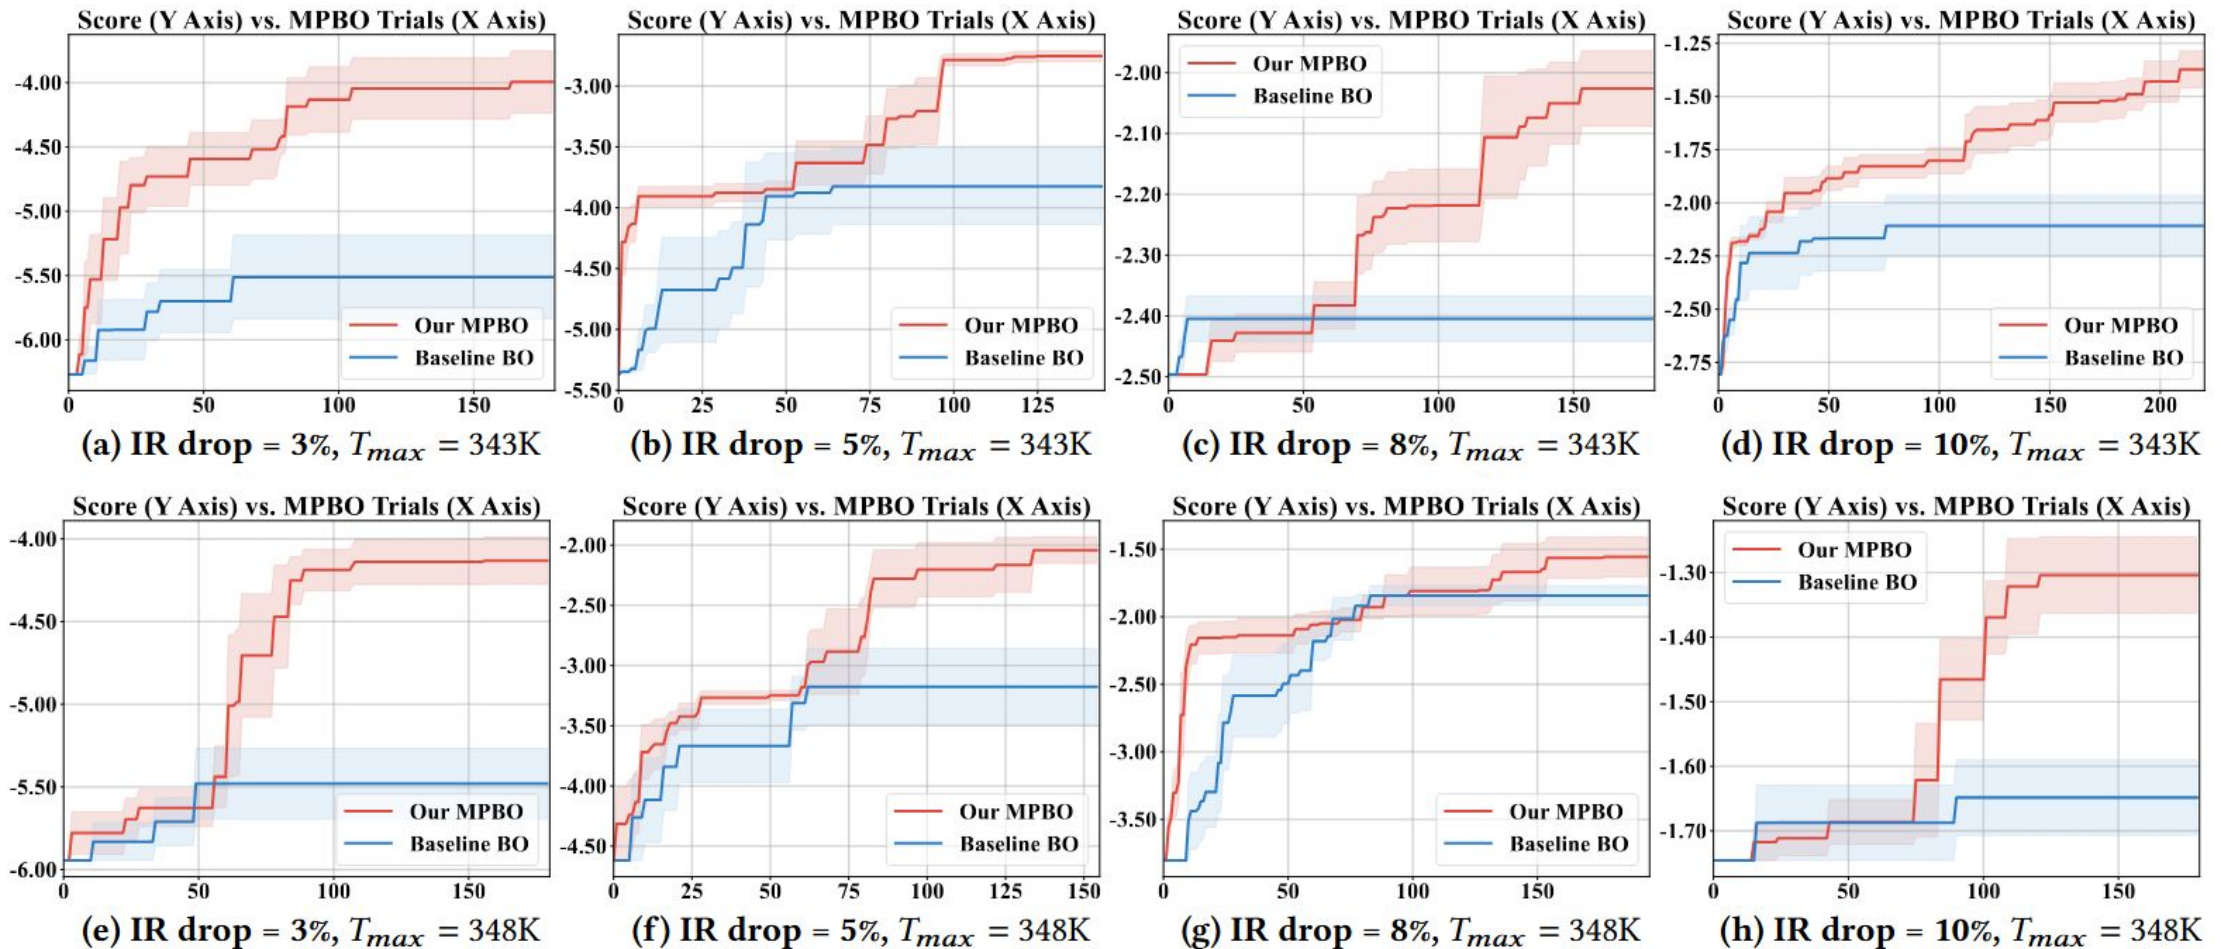


Figure 5. Comparison of Convergence between MIP-SPT and baseline BO on Case-1 (first row) and Case-3 (second row)

MIP-SPT: Experiment – Comparison with BO baseline

- Quantitative Comparison between our **MPBO** and **baseline SPBO**

Table 2. Mean and standard deviation of the best scores across eight experimental settings for MPBO and the baseline

Best score statistics		Case-1				Avg.	Imp.
		IR3	IR5	IR8	IR10		
Mean	Ours	-4.7442	-2.7550	-2.1227	-2.0260	-2.9120	24.12%
	Baseline	-6.2621	-3.8249	-2.8581	-2.4047	-3.8375	-
Std. Dev.	Ours	0.2423	0.0450	0.0877	0.0620	0.1092	0.53×
	Baseline	0.3272	0.3130	0.1441	0.0376	0.2055	-

Best score statistics		Case-3				Avg.	Imp.
		IR3	IR5	IR8	IR10		
Mean	Ours	-4.8816	-2.7909	-2.3065	-2.0539	-3.0082	20.57%
	Baseline	-6.2311	-3.9265	-2.5932	-2.3985	-3.7873	-
Std. Dev.	Ours	0.1435	0.1094	0.1478	0.0591	0.1149	0.69×
	Baseline	0.2151	0.3182	0.0742	0.0584	0.1665	-

[†]Imp. (%) for improvement: cost reduction and ratio of standard deviation

- We outperform the baseline on both final **converged value** and **standard deviation**

MIP-SPT: Experiment – Comparison with Previous Works

- We reproduce and compare with the previous works^{[6][7]}.
- On average, we reduce the manufacturing cost by 23.1% and 18.1% percent, respectively
- The advantage comes from the effective exploration of the design space and careful treatment of physical constraints.

Table 3. Comparison of our interconnect planning against previous works

Case	#RDL vias			#TSVs			Cost Reduction	
	[15]	[9]	Ours	[15]	[9]	Ours	v.s.[15]	v.s.[9]
Case1-2.5D	6423	5945	5430	280	264	190	-25.6%	-17.8%
Case2-2.5D	9887	9694	8338	432	420	293	-23.6%	-22.5%
Case3-3D	6619	6209	5585	549	528	440	-20.1%	-14.1%
Average	-	-	-	-	-	-	-23.1%	-18.1%

Conclusions

- We proposed a hierarchical variable scheduling strategy that decoupled interposer and substrate design variables and optimized them in separate phases
- To improve the convergence and efficiency of our MPBO, we employed score-based down-selection and elite-based sampling.
- Under the same SI/PI/TI constraints, our MPBO achieves a manufacturing cost of 22.4% lower than the baseline SPBO. Compared to two previous works, MIP-SPT reduces interconnect cost by 23.1% and 18.1%, respectively.

References

- [1] X. Meng, Y. Zhang, Y. Liu, P. Li, C. Wu and L. He, "Stackable Thermal Model for 3D Integration," 2025 International Symposium of Electronics Design Automation (ISED), Hong Kong, China, 2025, pp. 602-608
- [2] W. Liao et al., "Temperature and Supply Voltage Aware Performance and Power Modeling at Microarchitecture Level," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 24, no. 7, pp. 1042–1053, 2005.
- [3] R. Balasubramonian et al., "CACTI 7: New Tools for Interconnect Exploration in Innovative Off-Chip Memories," ACM Transactions on Architecture and Code Optimization (TACO), vol. 14, June 2017.
- [4] S. Lu et al., "An FPGA-based Multi-Core Overlay Processor for Transformer-based Models," in 2024 2nd International Symposium of Electronics Design Automation (ISED), pp. 697–702, 2024.
- [5] K. Wang et al., "Rethinking Thermal Via Planning with Timing-Power-Temperature Dependence for 3D ICs," in Proceedings of the 16th Asia and South Pacific Design Automation Conference, ASPDAC '11, p. 261–266, IEEE Press, 2011
- [6] H. Yu et al., "Thermal Via Allocation for 3D ICs Considering Temporarily and Spatially Variant Thermal Power," in ISLPED'06 Proceedings of the 2006 International Symposium on Low Power Electronics and Design, pp. 156–161, 2006.
- [7] Siyuan Miao, Ling kang Zhu, Wenkai Yang, Teng Lu, Yanze Zhou, Chen Wu, Zhiping Yu, Ting-Jung Lin, and Lei He. Electrothermal Simulation and Vertical Interconnect Planning for Integrated Chiplets. In 2025 International Symposium of Electronics Design Automation (ISED), pages 705–711, 2025.

Thank You

