Optimization of Modal Reflection Parameters of PCB Connector and Common-Mode Choke Using Gaussian Process Regression and Segmentation

Mikheil Kvizhinadze Nazim Talibzade Cheng Yang Matthias Mnich Christian Schuster
Hamburg University of Technology (TUHH)
Hamburg, Germany
mikheil.kvizhinadze@tuhh.de nazim.talibzade@tuhh.de cheng.yang@tuhh.de matthias.mnich@tuhh.de schuster@tuhh.de

***Abstract*—In this contribution the electrical properties of a high-speed signal transition from a connector onto a printed circuit board (PCB) followed by a common-mode choke (CMC) in a surface mount package in terms of its modal, i.e., differential- and common-mode, reflection parameters has been investigated. An implementation of segmentation, i.e. domain decomposition and subsequent composition of subdomains, in combination with Gaussian process regression (GPR) to optimize the reflections of the PCB connectors and CMC integration on the board is presented and analyzed. GPR is used with Bayesian optimization to find an optimal set of design parameters for a 3D electromagnetic model developed in a full-wave electromagnetic field solver up to 20 GHz. The results indicate that an engineering-informed domain decomposition strategy can reduce the computational burden of GPR-based optimization by lowering the per-evaluation full-wave cost. This surrogate-based optimization approach reduces the number of required full-wave simulations, enabling more efficient use of computational resources.**

***Index Terms*—Gaussian Process Regression, Bayesian Optimization, high-speed interconnects, PCB connectors, machine learning, common-mode choke packages, design optimization.**

I. Introduction

In high-speed printed circuit board (PCB) design finding design parameters that achieve optimal performance is a time-consuming process due to the large number of possible parameter combinations. Machine learning (ML) techniques can help reduce the need for costly electromagnetic (EM) simulations by efficiently selecting the most optimal next design point for evaluation, using methods such as Bayesian optimization (BO). Nonetheless the generation of data samples can cancel out computational advantages if the procedure is inefficient [1]. Bayesian learning approaches are becoming increasingly relevant for SI/ PI design in packaging, especially where uncertainty and insufficient data prevent direct exploration of expensive simulation models. Swaminathan *et al.* [2] explain how these tools can be used to solve SI/ PI problems in packaging and why probabilistic models are useful in these situations. Also, Bayesian framework for the modeling and optimization of high-speed interconnect channels is introduced by Torun *et al.* [3]. This framework demonstrates how surrogate-driven optimization can target link-level metrics. In their study of microwave devices, Garbuglia *et al.* [4] emphasize that the effectiveness of the GPR-BO framework depends on the design

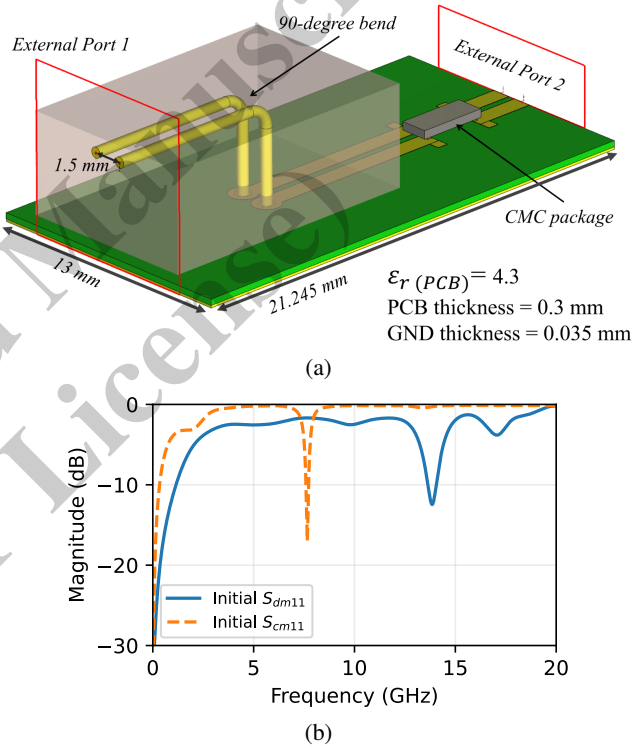

Fig. 1: (a) Parameterized 3D model used for full-wave simulation. (b) Scattering parameters of the initial design from full-wave simulation.

of the objective function, as complex frequency-dependent behavior must be reduced to a scalar value for optimization. More recently, Manfredi [5] introduced conservative GP formulations and extended them to uncertainty quantification and BO settings for SI test cases in order to highlight the significance of reliable uncertainty estimates in GPR. In this paper, implementation of engineering-informed segmentation in combination with GP/ BO-driven optimization is proposed and associated speed–accuracy trade-offs are analyzed. By doing so, we hope to gain efficiency in the optimization process and reduction of computational burden compared to a brute-force approach [6]. Segmentation has been successfully employed as a diagnostic tool to locate resonance-induced issues in complex power delivery networks [7], this work

Accepted by International Symposium on Electromagnetic Compatibility (EMC Europe) 2026
August 31 2026 – September 4 2026
Prague Congress Centre, Prague, Czech Republic

extends the technique toward the active optimization of modal reflection parameters for high-speed signal transitions. 3D model shown in Fig. 1 (a) comprises a differential pair of 90-degree angled input tubes mounted on a PCB, serving as a signal launch site. The signal propagates through microstrip transmission lines and is terminated by CMC approximation with packaging. GPR is used as a surrogate modeling method to improve the performance of the described structure at high frequencies and the potential of engineering-informed segmentation and its efficiency in combination with GPR-based BO is investigated. The rest of the paper is organized as follows: Section II introduces the problem, Section III presents the analysis and discussion of the results, and Section IV summarizes the work.

II. Problem Statement and General Procedure

A. Signal Integrity Challenges and Computational Bottleneck

On PCBs, transitions and component interfaces act as impedance discontinuities, degrading signal-integrity performance. At high frequencies, the response is further influenced by parasitic inductive and capacitive effects, which can lead to additional reflections and unintended mode interactions [8]. In this work, we minimize differential-mode reflection while ensuring that common-mode suppression meets the required level across the frequency band by penalizing any violations in the scalar cost. Achieving these objectives requires repeated evaluation of the scattering parameters. However, direct global optimization of the complete 3D structure becomes computationally expensive. For instance, exploring just nine design parameters with three discrete levels per variable would need 3^9 full-wave simulations, which is unrealistic in typical design cycles. To address this issue, the structure is decomposed into three subdomains and each subdomain is optimized locally, enabling localized, data-efficient optimization and reducing simulation cost, while yielding an overall performance improvement.

B. Conceptual Framework: Full vs. Segmented Optimization

This research considers two approaches to apply GPR-based BO to a full-wave model. In the full approach, BO treats the interconnect as a single black-box function of all design parameters and proposes new candidates in the full joint design space. Therefore, the complete 3D structure must be simulated at each iteration. In the segmented approach, the same BO loop is applied to smaller substructures, each with a reduced set of parameters. The locally optimized substructures are then reassembled into a complete design. This engineering-informed decomposition trades some interactions between substructures for lower computational cost, making GPR-based optimization practical when the full approach is expensive. This contrasts with the full approach, which captures all cross-domain electromagnetic interactions and prevents discontinuities that might be introduced after the optimized subdomains are reassembled.

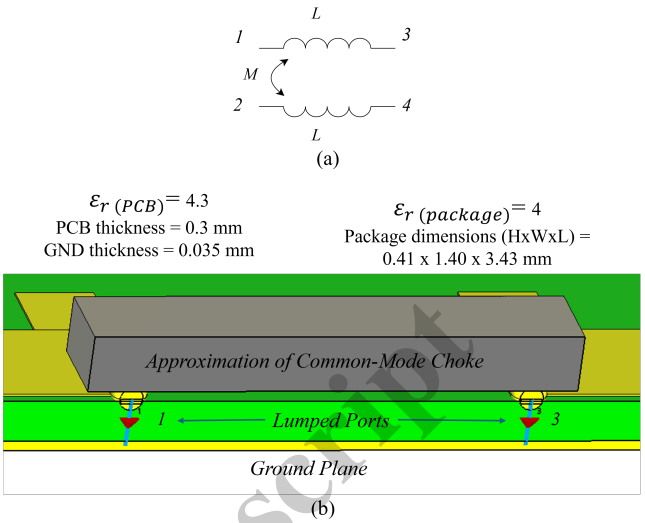

Fig. 2: (a) Equivalent circuit model for the CMC. (b) 3D model used for simulation of the CMC.

C. Mixed-Mode Scattering Parameters

The performance of a structure is evaluated using mixed-mode scattering parameters, which separate the response of a two-conductor differential pair into differential-mode (DM) and common-mode (CM) components [9]. In DM excitation, the two conductors carry equal-magnitude currents with opposite polarity, which is the intended signaling mode. In CM excitation, both conductors are driven in-phase with respect to the reference conductor, which is here a ground plane underneath. CM behavior is closely related to parasitic mode conversion and radiated emissions.

D. Simulation Setup and Port Configuration

Simulations were performed using full-wave simulator to extract the mixed-mode scattering parameters of the structure over 0-20 GHz [10]. The model was excited using two waveguide ports located at the external interfaces of the structure. In addition, four lumped ports were created internally at the CMC pins to provide terminals for the choke pins. Those lumped ports do not represent excitation sources, but rather, they serve as connections between the corresponding pin pads to allow equivalent circuit model integration. Fig. 2 shows the cut plane of the CMC package approximation domain pins 1 and 3 of the choke. To obtain a mixed-mode matrix, a co-simulation workflow was used: First, the 3D EM model was simulated using the frequency-domain solver which uses Finite

TABLE I: Design-space parameters and bounds.

Parameter	Symbol	Range
Signal tube radius (mm)	r_{sig}	[0.15, 0.43]
Cutout length (mm)	ℓ_{cut}	[2, 4]
Trace thickness (mm)	t_{tr}	[0.015, 0.095]
Trace width (mm)	w_{tr}	[0.1, 1]
Pad radius (mm)	r_{pad}	[0.1, 0.6]
Cutout width (mm)	w_{cut}	[1, 4]
CMC Pad width (mm)	$w_{\mathrm{cmc\_pad}}$	[0.1, 1.2]
Relative permittivity (PCB)	ε_r	4.3
PCB thickness (mm)	t_{pcb}	0.3
Coupling Coefficient (k)	k	[0.5, 1]
Inductance (nH)	L	[2, 10]

Accepted by International Symposium on Electromagnetic Compatibility (EMC Europe) 2026
August 31 2026 – September 4 2026
Prague Congress Centre, Prague, Czech Republic

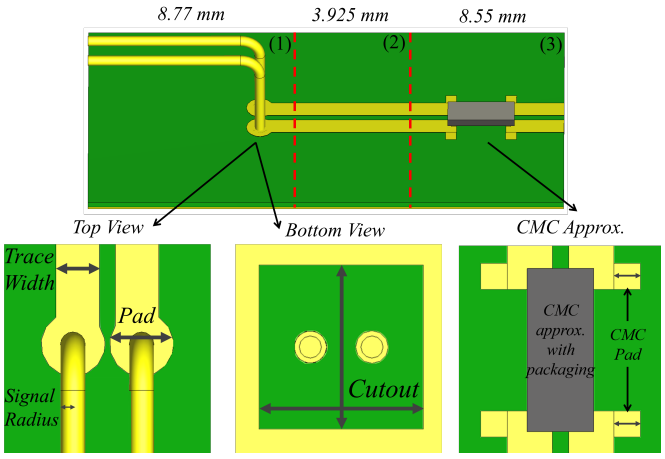

Fig. 3: Segmentation of the full interconnect into sub-domains used for simulation and optimization.

Integration Technique to obtain the full S-parameter description of the model. Next, the extracted S-parameter block was imported into the full-wave simulator's schematic environment and connected to the simplified coupled-inductor network between the corresponding discrete-port terminals. The DM and CM impedances are normalized to $Z_{DM} = 100\ \Omega$ and $Z_{CM} = 25\ \Omega$ and final mixed mode results are presented in Section III.

E. Decomposition Strategy and Design Space Definition

3D model shown in Fig. 3 is divided into three separate segments to make full-wave analysis faster and accelerate design exploration. The right-angle transition domain is the first region. It includes the 90-degree bend as well as the first part of PCB trace segments. The second segment is the transmission line area. The third segment is CMC package, which has both the connection and the transition for the CMC package approximation. Table I summarizes the parameters and their ranges used to define the design space, where r_{sig}, ℓ_{cut}, w_{cut}, and r_{pad} are local to the right-angle transition domain (Part 1). The trace dimensions w_{tr} and t_{tr} are shared across all three segments. The CMC is parameterized by the inductance L and coupling coefficient k. Finally, $\varepsilon_r(pcb)$, $\varepsilon_r(package)$, $\varepsilon_r(house)$ and t_{pcb} are kept fixed in all experiments. To validate the decomposition, each subdomain was simulated individually, and their S-parameter blocks were cascaded. The cascaded response shows strong broadband agreement. The largest discrepancy is a 0.26 GHz frequency shift at mid-frequency resonances, resulting in a localized magnitude deviation of 2.7 dB.

F. Parameterized Inductive Coupling Model

The behavior of CMC has been approximated by two coupled inductors, as shown in Fig. 2, to avoid the computational overhead, with

$$M = k\sqrt{L_1 L_2}\ . \tag{1}$$

If $L_1 = L_2 = L$, then $M = kL$, where $k \in [0, 1]$ is the coupling coefficient and we vary k and L in this simulation. Parasitic effects such as winding self-capacitance or inter-winding capacitance are not included in this study, since the focus is on reflection behavior.

G. Cost Function Definition from S-Parameters

The optimizer needs one scalar score per design to compare several candidates. In our application, not all frequency bands are equivalent. For the purpose of this study, the lower-frequency region is prioritized, however extending performance to higher frequencies is still valuable, since it increases usable bandwidth and provides additional margin. Therefore, we decided to define a scalar objective that penalizes *reflection-loss* violations in two frequency bands for S_{dm11} and additionally penalizes the deviation of S_{cm11} from a 0 dB target across 0-20 GHz band. We evaluate all S-parameters in dB over $f \in [0.4, 20]$ GHz, using a split at 10 GHz, on a discrete frequency grid $\{f_k\}_{k=1}^{N_f}$, where N_f is the total number of sampled frequency points in this band.

1) S_{dm11} Reflection-Loss Penalty (minimize): For S_{dm11} we use a two-band specification with a stricter low-band requirement. In the low band (0.4–10 GHz) we set the reflection threshold to -35 dB and apply weight $w = 2$, while in the high band (10–20 GHz) we use a threshold of -30 dB with weight $w = 1$. On the discrete grid $\{f_k\}_{k=1}^{N_f}$, we penalize only threshold violations via

$$v_k(x) = w_k \max\Big(S_{dm11}(f_k) - S_k^{\text{TH}},\, 0\Big), \tag{2}$$

where $(S_k^{\text{TH}}, w_k) = (-35\ \text{dB}, 2)$ for $0.4 \le f_k \le 10$ GHz and $(S_k^{\text{TH}}, w_k) = (-30\ \text{dB}, 1)$ for $10 < f_k \le 20$ GHz. An example of the penalty region is shown in Fig. 4. We aggregate violations using an RMS over each band $b \in \{\text{low}, \text{high}\}$:

$$c_b(x) = \sqrt{\frac{1}{N_{f,b}} \sum_{f_k \in \mathcal{F}_b} v_k(x)^2}, \tag{3}$$

with $\mathcal{F}_{\text{low}} = \{f_k : 0.4 \le f_k \le 10\ \text{GHz}\}$ and $\mathcal{F}_{\text{high}} = \{f_k : 10 < f_k \le 20\ \text{GHz}\}$, and $N_{f,b} = |\mathcal{F}_b|$.

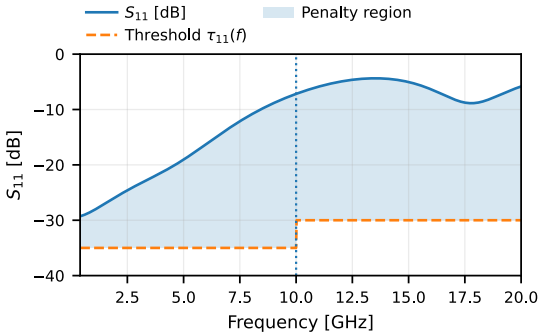

Fig. 4: Example of S_{dm11} cost function.

2) S_{cm11} Reflection-Loss Term (maximize): To suppress CM transmission, we maximize S_{cm11} by penalizing its deviation from a 0 dB target across the full band of interest:

$$c_{cm11}(x) = \sqrt{\frac{1}{N_f} \sum_{k=1}^{N_f} \Big(S_{cm11}(f_k; x) - 0\Big)^2}, \tag{4}$$

Accepted by International Symposium on Electromagnetic Compatibility (EMC Europe) 2026
August 31 2026 – September 4 2026
Prague Congress Centre, Prague, Czech Republic
Final Published Version © 2026 IEEE.

where the sum is taken over $0.4 \le f_k \le 20$ GHz.

3) Total Cost: BO needs one overall scalar objective to be minimized, and summing the penalty terms is the simplest smooth scalarization that reflects the overall level of violations across bands, which is expressed by

$$c_{\text{tot}}(x) = c_{\text{low}}(x) + c_{\text{high}}(x) + c_{cm11}(x) \ . \tag{5}$$

H. Gaussian Process Surrogate Model

BO requires a surrogate that maps a design vector $x \in \mathbb{R}^D$ (geometry / electrical parameters) to an objective value computed from full-wave simulated S-parameters, where D is the number of design parameters. In this work, we use Gaussian process regression (GPR) as the surrogate model because it provides a function model and an uncertainty that can be exploited by the acquisition function [11].

1) Single-output Gaussian Process Regression: Let $f(\cdot)$ denote the (unknown) mapping from design parameters to the scalar objective (e.g., $f(x) = c_{\text{tot}}(x)$ in (5)). We place a Gaussian process prior

$$f(\cdot) \sim \mathcal{GP}\big(m(\cdot),\, k(\cdot,\cdot)\big), \tag{6}$$

so that for training data $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ with $y_i = f(x_i)$, where N is the number of training samples, the posterior at a test point $x_\star$ is Gaussian with

$$\mu(x_\star) = k_\star^\top K^{-1} y, \tag{7}$$
$$\sigma^2(x_\star) = k(x_\star, x_\star) - k_\star^\top K^{-1} k_\star, \tag{8}$$

where $m(\cdot)$ is set to zero after standardization, $K_{ij} = k(x_i, x_j)$, $(k_\star)_i = k(x_i, x_\star)$, and $y = [y_1, \ldots, y_N]^\top$ [11]. The mean $\mu(x)$ serves as a surrogate model and $\sigma(x)$ quantifies predictive uncertainty.

2) Kernel Choice: Matérn with ARD: We use a Matérn kernel with automatic relevance determination (ARD), which is widely adopted in engineering because it offers controllable smoothness and typically generalizes well in low-dimensional design spaces [11]. For normalized inputs, the ARD Matérn-5/2 kernel is

$$k_\nu(x, x')\big|_{\nu=5/2} = \sigma_f^2 \left(1 + \sqrt{5}\, r + \frac{5}{3} r^2\right) \exp\left(-\sqrt{5}\, r\right), \tag{9}$$

with

$$r = \sqrt{\sum_{d=1}^{D} \frac{(x_d - x'_d)^2}{\ell_d^2}}, \tag{10}$$

where σ_f^2 is the signal variance and $\ell_d > 0$ are the ARD lengthscales. A smaller ℓ_d indicates that the surrogate changes more rapidly along parameter d, suggesting stronger sensitivity of the objective to that parameter, whereas larger ℓ_d indicates weaker dependence.

I. Bayesian Optimization and Acquisition Function

BO selects the next design point to evaluate by maximizing an acquisition function, and then evaluates it using a full-wave simulator [12]. At iteration t, the GP provides a predictive distribution for the scalar cost $c(x)$. Let $c_{\min} = \min_{i \le t} c(x_i)$ be the smallest observed cost. The next point is selected by maximizing expected improvement (EI):

$$\text{EI}(x) = \mathbb{E}\big[\max\big(c_{\min} - c(x) - \xi,\, 0\big)\big], \tag{11}$$

where $\xi \ge 0$ is an improvement margin that makes EI more exploratory.

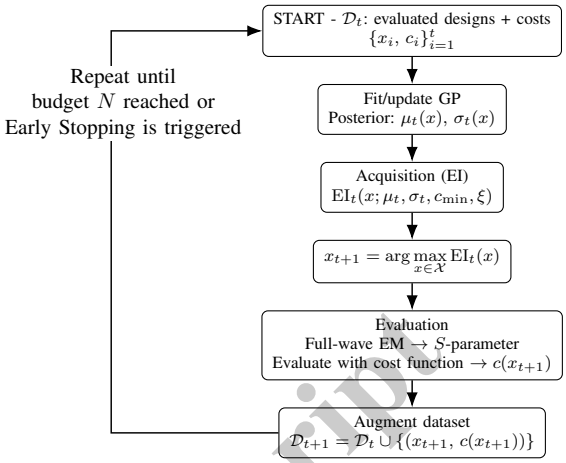

Fig. 5: Closed-loop Bayesian optimization.

J. Experimental Setup and Early Stopping

For each design case, we generate an initial set of designs using Latin hypercube sampling (LHS) [13] and evaluate each design with the full-wave simulator. For every evaluation, we extract the S-parameters, and compute the cost. Afterwards, we fit a GPR surrogate model to the compiled dataset. Then, we use BO to propose the next design by maximizing an acquisition function based on the current surrogate. Finally, we run the full-wave simulation for the proposed design, compute the cost of S-parameters, append the new sample to the dataset, and repeat the loop until the evaluation budget is exceeded or an early stopping rule is met.

We apply a plateau-based early stopping rule. Let $c_{\min}^{(t)}$ be the smallest cost observed up to iteration t. If

$$c_{\min}^{(t-1)} - c_{\min}^{(t)} < \varepsilon \tag{12}$$

holds for P consecutive iterations, we terminate early. We use $N_{\text{init}} = 20$ LHS evaluations, then BO with $N_{\max} = 50$, and set $\varepsilon = 10^{-4}$ and $P = 7$ for all configurations.

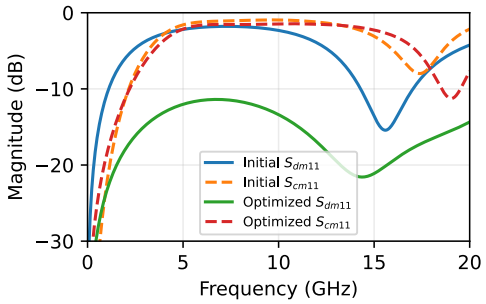

Fig. 6: S_{dm11} and S_{cm11} comparison of the 90 degree bend for the initial and optimized designs.

Accepted by International Symposium on Electromagnetic Compatibility (EMC Europe) 2026
August 31 2026 – September 4 2026
Prague Congress Centre, Prague, Czech Republic

III. Analysis and Discussion of Results

A. Evaluation of Results

The performance of the proposed approach is evaluated by comparing the reflections of the initial and optimized designs. The comparison results were obtained relative to a reference design: $r_{\text{sig}} = 0.15$mm, $\ell_{\text{cut}} = 4$mm, $t_{\text{tr}} = 0.015$mm, $w_{\text{tr}} = 0.55$mm, $r_{\text{pad}} = 0.6$mm, $w_{\text{cut}} = 4$mm, $w_{\text{cmc_pad}} = 1$mm, $\varepsilon_r = 4.3$, and $t_{\text{pcb}} = 0.3$mm and the remaining default CMC parameters are $k = 0.8$mm, $L = 2$ nH.

1) Segmented Optimization: In the segmented approach, each subdomain is optimized independently and the resulting geometries are finally reassembled to form a complete structure. Fig. 6 illustrates that for 90-degree bend, both S_{dm11} and S_{cm11} were successfully optimized. Fig. 8 (a) shows that the CMC approximation subdomain S_{dm11} is reduced, but S_{cm11} changed relatively little. Also, as seen from Fig. 7, local optimization can introduce step changes in the trace cross-section after reassembly. These boundaries create impedance discontinuities and hence local mismatch in characteristic impedance. Such discontinuities can be interpreted as localized parasitic reactances and are commonly modeled as microstrip step discontinuities whose effect becomes more pronounced at higher frequencies [14]. To control the mismatch, continuity constraints can be introduced at the reassembly planes by restricting the parameter boundaries to match within a specified tolerance e.g., $|w_2 - w_1| \leq \Delta w$. Finally, these regions can be used as a design regions, in which only boundary geometry is optimized. Fig. 8 (b) summarizes the intermediate and final optimization results.

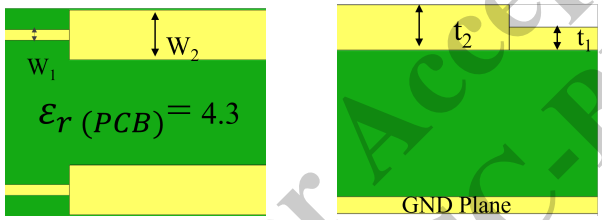

(a) Top view (b) Side view

Fig. 7: Trace Discontinuity in the 3D Model.

2) Full Optimization: In the full-structure optimization approach, the complete 3D model is optimized directly without segmentation. Fig. 8 (c) compares the initial design, the optimized design obtained from full-structure optimization, and design obtained from segmented approach. The latter improves the reflection relative to the initial design, while requiring a lower simulation cost per evaluation than the full-structure optimization case, as seen in Table II, which also illustrates the number of iterations at which the best objective was observed, and the number of iterations, at which the optimizer terminated.

B. ARD-Based Parameter Importance

We use ARD to quantify how strongly the GP surrogate varies with each design parameter in the normalized input space. The importance is based on inverse lengthscales $(1/\ell_d)$,

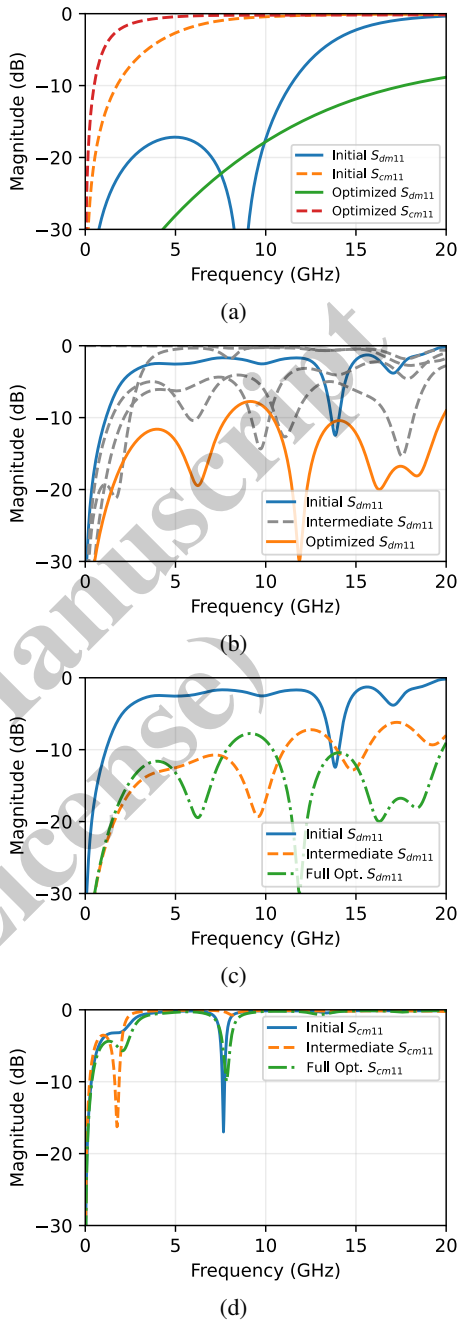

Fig. 8: (a) S_{dm11} and S_{cm11} comparison of CMC domain for the initial and optimized designs. (b) Evolution of $|S_{dm11}|$ during optimization of the full structure (initial, intermediate iterations, and final optimized). (c) $|S_{dm11}|$ and (d) $|S_{cm11}|$ comparison for the initial design, segmented-model optimized design, and full-3D optimized design over 0–20 GHz.

where larger values indicate higher sensitivity of the scalar objective. Since all parameters are scaled to $[0, 1]$, smaller ℓ_d implies stronger dependence within the sampled region. The importance ranking from the GP surrogate shows that the objective is most sensitive to $w_{\text{cmc\_pad}}$ $(1/\ell_{\text{cmc\_pad}} \approx 0.32)$, followed by k $(1/\ell_k \approx 0.26)$ and w_{tr} $(1/\ell_{w_{\text{tr}}} \approx 0.18)$. This indicates that parameters which shape impedance discontinuity and modal coupling at the CMC transition dominate the

Accepted by International Symposium on Electromagnetic Compatibility (EMC Europe) 2026
August 31 2026 – September 4 2026
Prague Congress Centre, Prague, Czech Republic
Final Published Version © 2026 IEEE.

TABLE II: Average next BO point proposal time and EM simulation duration per evaluation in seconds.

Segment	GPR+BO propose time	Sim/eval	Best Iteration	Stop Iteration
Input + 90° bend	1.48	123.00	5	12
Trace	0.62	91.31	1	8
CMC approximation	0.85	432.94	13	20
Whole structure	1.76	1018.24	13	20

reflection behavior across the band. The r_{sig} has a smaller influence ($1/\ell_{r_{\mathrm{sig}}} \approx 0.09$), while ℓ_{cut} ($1/\ell_{\ell_{\mathrm{cut}}} \approx 0.04$–$0.05$) and L ($1/\ell_L \approx 0.03$–0.04) contribute the least, suggesting a secondary effect within the explored parameter ranges. This ranking is a data-driven sensitivity indicator, not a universal physical law.

C. Computational Effort

We measure computational effort by tracking the full-wave solver time per simulation, T_s, and the number of BO iterations, as summarized in Table II. By relating T_s to the number of mesh cells, N_{mesh}, an approximate polynomial scaling behavior on a workstation equipped with an Intel Core i5-6500 CPU @ 3.2 GHz and 64 GB RAM was observed:

$$T_s \approx N_{\mathrm{mesh}}^{p}, \qquad p \approx 1.4 \text{ (within the tested range).} \quad (13)$$

This trend illustrates why full-structure optimization becomes difficult as N_{mesh} grows. Furthermore, exploring the design space suffers from the curse of dimensionality. As the dimensionality D increases, the volume of the search space grows exponentially. As a result, maintaining a comparable sampling density requires an exponentially larger number of design evaluations. In high-dimensional spaces, for any fixed N, the typical distance between samples increases with D, so each design has fewer nearby neighbors. Consequently, a GPR surrogate has limited local information to interpolate from, and the posterior uncertainty remains high over large portions of the domain unless N increases substantially with D. In addition, exact GP training scales cubically with the number of evaluated designs. Because the required sample size N must grow significantly as D increases to maintain predictive accuracy, the computational cost of the BO loop itself can become a bottleneck. Segmented approach used in this paper addresses both the curse of dimensionality and the solver burden by partitioning the original D-dimensional problem into lower-dimensional subproblems. This reduces surrogate-fitting cost and decreases cost per evaluation by simulating smaller subdomains. Finally, including the initial sampling budget, segmentation reduced the total runtime by 41.5%. In addition, the best parameter set was found in 51.3% fewer BO iterations, and the overall optimization finished in 33.3% fewer iterations.

IV. Conclusion

In this paper, the potential of segmentation, in combination with GPR for connector and CMC package integration on PCBs has been demonstrated. By dividing the structure into multiple domains, the problematic regions that dominate reflection are identified. The proposed method reduces computation time by minimizing the need to re-simulate the entire structure and enables faster optimization. Future work will extend the proposed approach to stochastic yield analysis, for example, considering manufacturing tolerances in trace geometries.

Acknowledgment

This project has received funding from the European Union's EU Framework Programme for Research and Innovation Europe Horizon (Grant Agreement No. 101169295, for Doctoral Candidates 1–15) and DMU (De Montfort University, for Doctoral Candidate 16). Project website: pattern-dn.eu/

References

[1] Y. Hassab, M. Schierholz and C. Schuster, "Application of Gaussian Process Regression for Data Efficient Prediction of PCB-Based Power Delivery Network Impedance Features," in *2024 IEEE 28th Workshop on Signal and Power Integrity (SPI)*, Lisbon, Portugal, 2024, pp. 1–4.

[2] M. Swaminathan, H. M. Torun, H. Yu, J. A. Hejase, and W. D. Becker, "Demystifying Machine Learning for Signal and Power Integrity Problems in Packaging," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 10, no. 8, pp. 1276–1295, Aug. 2020.

[3] H. M. Torun, M. Larbi, and M. Swaminathan, "A Bayesian Framework for Optimizing Interconnects in High-Speed Channels," in *Proc. 2018 IEEE MTT-S International Conference on Numerical Electromagnetic and Multiphysics Modeling and Optimization (NEMO)*, 2018

[4] F. Garbuglia, D. Spina, D. Deschrijver, I. Couckuyt, and T. Dhaene, "Bayesian optimization for microwave devices using deep GP spectral surrogate models," *IEEE Transactions on Microwave Theory and Techniques*, vol. 71, no. 6, pp. 2311–2318, 2022.

[5] P. Manfredi, "Conservative Gaussian Process Models for Uncertainty Quantification and Bayesian Optimization in Signal Integrity Applications," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 14, no. 7, pp. 1–12, 2024.

[6] M. Schierholz, Y. Hassab and C. Schuster, "Engineering-Informed Design Space Reduction for PCB-Based Power Delivery Networks," in IEEE Transactions on Components, Packaging and Manufacturing Technology, vol. 13, no. 10, Oct. 2023, pp. 1613-1623.

[7] Y. Chen, Z. Xu, J. Chen, B. Pu, J. Zhou, and X. Chu, "Problem Locating Method for PDN Structures with Segmentation Technique," in *2023 IEEE 7th International Symposium on Electromagnetic Compatibility (ISEMC)*, 2023.

[8] S. Wang, F. C. Lee and J. D. Van Wyk, "Design of inductor winding capacitance cancellation for EMI suppression," in *2006 37th IEEE Power Electronics Specialists Conference*, Jeju, Korea (South), 2006, pp. 1–7.

[9] D. E. Bockelman and W. R. Eisenstadt, "Combined differential and common-mode scattering parameters: theory and simulation," *IEEE Transactions on Microwave Theory and Techniques*, vol. 43, no. 7, pp. 1530–1539, July 1995.

[10] Dassault Systèmes, CST Studio Suite, version 2025. [Online]. Available: www.3ds.com

[11] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.

[12] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas, "Taking the human out of the loop: A review of Bayesian optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, 2016.

[13] M. D. Shields and J. Zhang, "The generalization of Latin hypercube sampling," *Reliability Engineering & System Safety*, vol. 148, pp. 96–108, 2016.

[14] N. H. L. Koster and R. H. Jansen, "The Microstrip Step Discontinuity: A Revised Description," *IEEE Transactions on Microwave Theory and Techniques*, vol. 34, no. 2, pp. 213–223, Feb. 1986.

Accepted by International Symposium on Electromagnetic Compatibility (EMC Europe) 2026
August 31 2026 – September 4 2026
Prague Congress Centre, Prague, Czech Republic