

FastRW: An Efficient Random Walk Method for Steady-State Thermal Analysis

Zixiao Wang, Tianshu Hou[†], Chenghan Wang, Zhen Zhuang, Tsung-Yi Ho, Farzan Farnia, Bei Yu[†]
The Chinese University of Hong Kong

Abstract—Thermal simulation is increasingly critical in modern IC design and manufacturing. Random walk methods based on the Feynman–Kac formula enable efficient local temperature estimation without computing the full temperature field. However, in practical scenarios without Dirichlet boundary conditions, these methods often require excessively long paths and heuristic truncation rules. In this work, we revisit Feynman–Kac sampling and derive an exact characterization of the truncation error: the expected residual contribution is a simple scalar multiple of the temperature at the truncation point. This insight leads to FastRW, a random-walk framework that safely applies aggressive truncation. FastRW uses a cheap, noisy prior temperature field to approximate the residual term and shorten individual paths, and further exploits cross-relations among query points through a Bayesian posterior update to reduce the number of required walks. Experiments on 3DIC steady-state thermal benchmarks show that FastRW achieves over 6× speedup over prior Feynman–Kac-based methods with better accuracy.

I. INTRODUCTION

Heterogeneous integration technologies such as 3D integrated circuits (3D-ICs) are essential for continued system-level scaling in advanced microelectronics [1]. However, the resulting increase in power density aggravates thermal issues, leading to performance degradation and reduced lifetime [2]. Steady-state heat conduction is governed by computationally expensive partial differential equations, and design-space exploration may require thousands of thermal simulations to evaluate different power maps and packaging options [3]–[5].

In many design tasks, engineers only need the temperatures at a small set of critical locations (e.g., hotspots), rather than the full temperature field [6]. Random-walk methods address this need by estimating temperatures at queried points through Monte Carlo simulation of stochastic particle paths. They parallelize naturally and have been successfully applied to thermal and other EDA problems under Dirichlet boundary conditions [7]–[9]. For mixed Neumann and Robin boundaries, Feynman–Kac-based formulations provide a principled way to handle boundary contributions by accumulating local time and weighting them with an exponential functional, as in [10]–[13]. While accurate, these methods often require very long paths in practical scenarios without Dirichlet boundaries, where path termination relies solely on driving the Feynman–Kac functional below a small heuristic threshold, with no quantitative control of truncation error [14].

[†]Corresponding authors.

This work is partially supported by Research Grants Council of Hong Kong SAR (No. RFS2425-4S02).

We consider steady-state 3D thermal conduction in a domain $\Omega \subset \mathbb{R}^3$ governed by a Poisson equation with mixed boundary conditions and focus on efficiently estimating temperatures at a set of query locations $\{x_q\} \subset \Omega$. Let X_t denote a standard reflecting Brownian motion (SRBM) in Ω , which is reflected at Neumann and Robin boundaries and absorbed at Dirichlet boundaries. The temperature field admits the Feynman–Kac representation [15]

$$T(x) = \mathbb{E} \left[\int_0^\infty \hat{e}_c(t) f(X_t) dt + \mathcal{B}(X_t, L(t)) \right], \quad (1)$$

where $\hat{e}_c(t) = \exp(-\int_0^t c(X_s) dL(s))$ is an exponential weight, f denotes the source term, $L(t)$ is the local time accumulated at Neumann and Robin boundaries, and $\mathcal{B}(X_t, L(t))$ compactly collects Dirichlet, Neumann, and Robin boundary contributions as in [11]–[13]. In practice, the infinite-time integrals in Equation (1) must be truncated, and the main challenge in the no-Dirichlet case is to control the error introduced by truncation.

In this work, we revisit Feynman–Kac-based random walks from the perspective of truncation error. Our contributions are:

- We provide a quantitative analysis of the truncation error in Feynman–Kac-based thermal estimation and show that the expected residual contribution is exactly a scalar multiple of the temperature at the truncation point.
- Based on this insight, we propose FastRW, a random-walk framework that safely applies aggressive truncation by combining a cheap, noisy prior with multi-point posterior fusion.
- Experiments on realistic 3D chip thermal benchmarks demonstrate that FastRW achieves over 6× overall speedup compared with prior Feynman–Kac-based random-walk solvers, without loss of accuracy.

II. TRUNCATION ERROR ANALYSIS

We focus on the practically important case without Dirichlet boundaries, where path termination in the Feynman–Kac representation Equation (1) is solely controlled by a threshold on the functional $\hat{e}_c(t)$. For a given threshold $\Lambda \in (0, 1)$, let

$$t_\Lambda \triangleq \inf\{t \geq 0 : \hat{e}_c(t) \leq \Lambda\} \quad (2)$$

be the first time the functional drops below Λ , and denote by X_{t_Λ} the truncation location. For a single random-walk path starting from x , we write the corresponding Monte Carlo estimate as

$$T_i(x) = P_0^{t_\Lambda} + P_{t_\Lambda}^\infty, \quad (3)$$

where $P_0^{t_\Lambda}$ collects the contribution from the interval $[0, t_\Lambda]$ and $P_{t_\Lambda}^\infty$ is the residual contribution from (t_Λ, ∞) . If the path is truncated at t_Λ , then $P_{t_\Lambda}^\infty$ is exactly the truncation error for this realization.

Our key observation is that the expected truncation error admits a simple closed-form expression. Conditioning on the truncation state X_{t_Λ} and using the Markov property of the reflecting Brownian motion, the future evolution after t_Λ is statistically identical to a new path started from X_{t_Λ} with the Feynman–Kac functional initialized at $\hat{e}_c(t_\Lambda)$ [15]. Therefore the residual contribution satisfies

$$\mathbb{E}[P_{t_\Lambda}^\infty | X_{t_\Lambda}] = \hat{e}_c(t_\Lambda) T(X_{t_\Lambda}) = \Lambda T(X_{t_\Lambda}), \quad (4)$$

and hence $\mathbb{E}[P_{t_\Lambda}^\infty] = \Lambda \mathbb{E}[T(X_{t_\Lambda})]$. The identity in Equation (4) is the central theoretical insight underlying FastRW.

An intuitive interpretation of Equation (4) is that the residual integral $P_{t_\Lambda}^\infty$ behaves as if a virtual Dirichlet boundary with temperature $T(X_{t_\Lambda})$ were imposed at the truncation point. In contrast to heuristic choices of Λ in prior work, this identity shows that the expected truncation error scales linearly with both Λ and the temperature at the truncation location. Consequently, if a coarse prior $\tilde{T} \approx T$ is available at X_{t_Λ} , we can approximate the residual using $\Lambda \tilde{T}(X_{t_\Lambda})$ and safely employ much larger thresholds, which forms the basis of our FastRW framework in Section III.

III. FASTRW: ALGORITHM AND RESULTS

A. Algorithm Overview

Building on the truncation error characterization in Section II, FastRW consists of two components: single-point acceleration using a prior temperature field and multi-point acceleration using posterior fusion.

Single-point with priors. For a single query location x , conventional Feynman–Kac-based solvers terminate a path at time t when $\hat{e}_c(t) \leq \Lambda$ and use the partial integral P_0^t as the path contribution. In contrast, FastRW assumes a coarse prior

$$\tilde{T}(x) = T(x) + \epsilon, \quad x \in \Omega_s, \quad (5)$$

available in a subdomain $\Omega_s \subseteq \Omega$ (e.g., from a low-resolution FEM run). Whenever a path is truncated at time t inside Ω_s , we augment the contribution as

$$T_i(x) = P_0^t + \hat{e}_c(t) \tilde{T}(X_t). \quad (6)$$

Using the identity in Equation (4), the expected truncation error is reduced from $\hat{e}_c(t) T(X_t)$ to $\hat{e}_c(t) \epsilon$. Since the prior error ϵ is typically much smaller than the temperature scale, this reduction allows FastRW to use a larger truncation threshold Λ while keeping the truncation error within the same tolerance, thereby shortening path lengths and reducing variance.

Multi-point with posterior fusion. When multiple query locations $\{x_q\}_{q=1}^M$ are of interest, standard approaches estimate each $T(x_q)$ using only paths originating from x_q . However, a path that starts at x and reaches y at time t also encodes a linear relation between the two temperatures: the partial integral P_0^t can be viewed as a noisy observation of $T(x) - \alpha T(y)$ with

TABLE I Representative single-point results comparing PIRW and FastRW.

Case	Method	Error (K)	Steps	Paths	Speedup
1	PIRW	0.38	5.79e6	1000	1.0×
	FastRW	0.36	2.20e6	400	6.58×
2	PIRW	0.43	5.83e6	1000	1.0×
	FastRW	0.42	2.22e6	400	6.57×
3	PIRW	0.38	1.03e7	1000	1.0×
	FastRW	0.32	3.93e6	400	6.55×

weight $\alpha = \hat{e}_c(t)$. FastRW records such cross-point observations whenever a path starting from one query point visits another, stacks them into a collection of linear measurements, and fuses them with the single-point Monte Carlo estimates via a Gaussian prior and a linear-Gaussian posterior update. Intuitively, shared paths provide additional constraints across query points, increasing the *effective* number of samples per point without simulating extra paths.

B. Experimental Results and Conclusion

We evaluate FastRW on standard 3D steady-state chip thermal benchmarks following the setups in [11]–[13]. The test structure is a 2 cm × 2 cm three-layer stack with a middle heat-source layer sandwiched between two source-free layers (heat spreader and substrate) [11]. Robin boundary conditions are imposed on the top and bottom surfaces, and lateral surfaces are modeled as adiabatic Neumann boundaries. Golden references are obtained from a commercial FEM solver with fine discretization, whereas coarse FEM runs provide priors.

Our main baseline is PIRW [11], a Feynman–Kac-based solver that truncates paths using a small threshold $\Lambda = 10^{-4}$. In FastRW, we keep all Feynman–Kac parameters identical but use a much larger threshold $\Lambda = 0.03$ enabled by the truncation error analysis and priors. TABLE I reports three representative cases from the full benchmark suite. In all cases, both methods achieve sub-kelvin absolute error, but FastRW consistently shortens paths and reduces the number of required walks, yielding an overall speedup of over 6×; for example, in Case 1 the runtime improvement is 6.58×.

We also examine the benefit of posterior fusion in multi-point settings. For a representative 3D stack, we uniformly select 16 query locations in the heat-source region and apply FastRW with a fixed number of paths per point. As random walks more frequently visit other query locations, the posterior fusion reduces the empirical standard deviation from 0.446 to 0.281, corresponding to an effective 2.5× increase in sample size per point at the same random-walk cost. The extra runtime of the posterior stage remains modest: even for all 16 points fused together, the combined overhead of prior and posterior computation is below 6% of the random-walk time.

In summary, FastRW accelerates steady-state thermal analysis by combining an exact characterization of Feynman–Kac truncation error with coarse prior estimation and multi-point posterior fusion.

REFERENCES

- [1] “Microelectronics and advanced packaging technologies roadmap,” Tech. Rep. [Online]. Available: <https://srcmapt.org>
- [2] K. Puttaswamy and G. H. Loh, “Thermal analysis of a 3d die-stacked high-performance microprocessor,” in *Proceedings of the 16th ACM Great Lakes symposium on VLSI*, 2006, pp. 19–24.
- [3] H. Sultan and S. R. Sarangi, “Varsim: A fast process variation-aware thermal modeling methodology using green’s functions,” *Microelectronics Journal*, vol. 142, p. 105995, 2023.
- [4] Y. Ma, L. Delshadtehrani, C. Demirkiran, J. L. Abellan, and A. Joshi, “Tap-2.5d: A thermally-aware chiplet placement methodology for 2.5d systems,” in *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2021, pp. 1246–1251.
- [5] D.-W. Wang, B.-W. Zhang, L.-T. Wang, P. Zhang, and W.-S. Zhao, “A proposal of fast thermal simulation method for 2.5-d advanced packaging to enable efficient thermal-aware placement optimization,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, pp. 1–1, 2025.
- [6] Y. Zhan, B. Goplen, and S. S. Sapatnekar, “Electrothermal analysis and optimization techniques for nanoscale integrated circuits,” in *Proceedings of the 2006 Asia and South Pacific Design Automation Conference*, 2006, pp. 219–222.
- [7] E. Wong and S. K. Lim, “3d floorplanning with thermal vias,” in *Proceedings of the Design Automation & Test in Europe Conference*, vol. 1, 2006, pp. 1–6.
- [8] H. Qian, S. Nassif, and S. Sapatnekar, “Power grid analysis using random walks,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 24, no. 8, pp. 1204–1224, 2005.
- [9] M. Visvardis, P. Liaskovitis, and E. Efstathiou, “Deep-learning-driven random walk method for capacitance extraction,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 42, no. 8, pp. 2643–2656, 2023.
- [10] Y. Liang, W. Yu, and H. Qian, “A hybrid random walk algorithm for 3-d thermal analysis of integrated circuits,” in *2014 19th Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2014, pp. 849–854.
- [11] L. Yang, C. Ding, C. Yan, D. Zhou, and X. Zeng, “A high-precision stochastic solver for steady-state thermal analysis with fourier heat transfer robin boundary conditions,” in *2022 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, 2022, pp. 1–9.
- [12] Y. Luo, Y. Wang, Z. Dong, C. Yan, Z. Bi, K. Zhu, S.-G. Wang, and X. Zeng, “Gpu acceleration of a high-precision stochastic solver for steady-state thermal analysis with robin boundary conditions,” in *2025 International Symposium of Electronics Design Automation (ISED)*. IEEE, 2025, pp. 287–292.
- [13] Z. Dong, L. Yang, C. Ding, C. Yan, Z. Bi, S.-G. Wang, D. Zhou, and X. Zeng, “ppirw: An efficient and accurate precalculation path integral random walk solver for steady-state thermal simulation with robin boundary conditions,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 44, no. 4, pp. 1489–1502, 2025.
- [14] R. Sawhney, B. Miller, I. Gkioulekas, and K. Crane, “Walk on stars: A grid-free monte carlo method for pdes with neumann boundary conditions,” *ACM Transactions on Graphics (TOG)*, vol. 42, no. 4, pp. 1–20, 2023.
- [15] B. Oksendal, *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.