Design for Testing and Reliability in the Presence of Transistor Self-Heating for Advanced Technologies

Milestone Report – GS-IMTR

Florian Klemme

October 10, 2022

Advisor / main examiner: Jun.-Prof. Dr.-Ing. Hussam Amrouch

Co-examiner: Prof. Dr. rer. nat. habil. Ilia Polian

Mentor: Dr. rer. nat. Martin Zimmermann, Advantest

1 Introduction

This report gives a comprehensive summary of the state of GS-IMTR project 9 (P9) "Design for Testing and Reliability in the Presence of Transistor Self-Heating for Advanced Technologies" as part of the Milestone Presentation. P9 focuses on investigating emerging reliability challenges in advanced technology nodes, with focus on (but not limited to) the transistor self-heating effect. With the lack of commercial tool support for many of these emerging challenges, P9 aims at developing novel approaches, including machine learning techniques, to overcome the limitations of state-of-the-art tool flows and elevate reliability challenges from individual transistors to the circuit level.

1.1 Transistor Self-Heating

Transistor self-heating is an ever-increasing reliability concern due to the continuous scaling in the semiconductor fabrication process. With more confined 3D structures and insulating materials in modern transistor designs, heat arising in the transistor's channel area cannot be dissipated easily and is hence "trapped" inside the transistor's channel [1, 2], as depicted in Figure 1. Due to the self-heating effect, the transistor effectively operates at a temperature higher than its surrounding material would reveal. This does not only impact the electrical characteristics of the transistor but, importantly, also its *reliability*. Under increased temperatures, existing reliability degradations are further exacerbated. In particular, transistor aging effects are considerably accelerated, ultimately leading to earlier failure of the chip. Therefore, to ensure reliability of the complete chip for its entire projected lifetime (e.g., 10 years), transistor self-heating must be accurately considered from the transistor level all the way up to the final chip (i.e., GDS-II level).

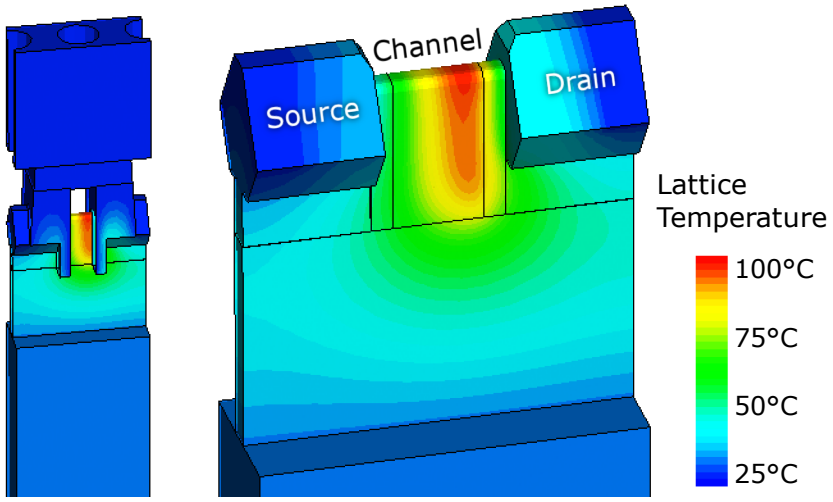

Figure 1: Transistor self-heating in 14 nm FinFET, modeled in Synopsys TCAD after calibration against industrial measurements. The insulating gate is removed visually, unveiling the temperature hot spot inside the channel. The surrounding material remains relatively cool, concealing the self-heating effect from on-chip thermal sensors.

While transistor self-heating is widely studied at the transistor level, research on its impact at the circuit level is still in its infancy and none of the existing Electronic Design Automation (EDA) tool flows can consider SHE during the chip design. Most research is currently limited to merely ring oscillators or simple logic cells using SPICE simulations [3, 4, 5]. With the lack of EDA support for SHE, circuit designers are forced to fall back to worst-case estimations or measurements obtained from the transistor level. To ensure reliability, they must assume the highest expected SHE-induced degradations to occur in every transistor across their entire chip. While it does guarantee reliability, such an approach leads to largely pessimistic safety margins, causing sub-optimal circuit designs and lost potential in the advanced technology.

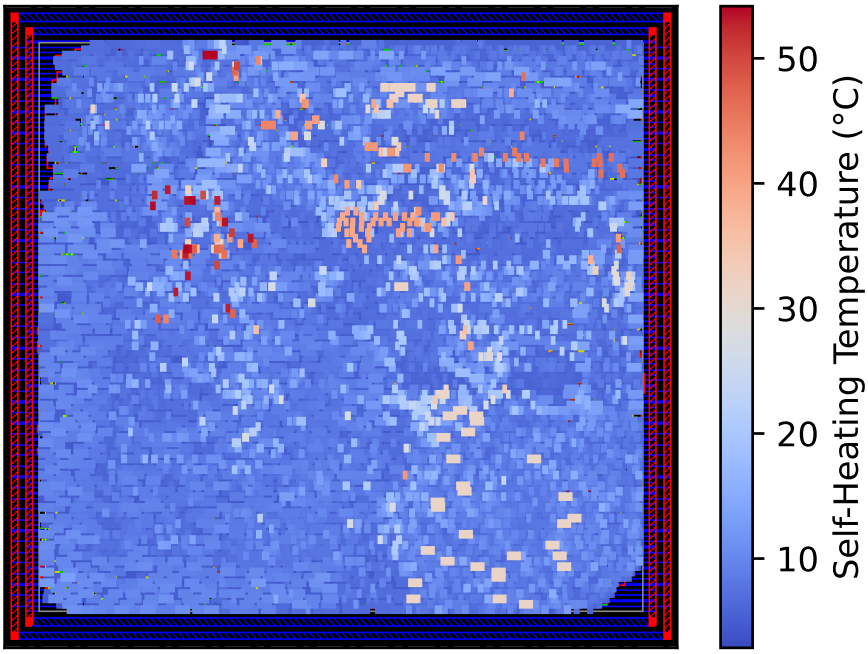

Figure 2: "Self-heating heat map", drawn on top of a RISC-V processor core [6]. Each cell in the RISC-V circuit is colorized according to its self-heating temperature (above chip) and drawn on top of the layout obtained from the place and route tool. Note that temperatures are not "spreading" to neighbor cells—Self-heating temperatures are trapped inside the individual transistors.

In our recent research [6], we propose a novel approach, elevating transistor-level self-heating effects to the circuit level, unveiling self-heating-induced reliability challenges on full circuit designs, as depicted in Figure 2. This step poses an important milestone in our GS-IMTR project, preceding our already published research on how to efficiently model heterogeneous standard cell degradations for large circuit designs. By combining our recent and already published research, a full reliability framework for transistor self-heating on the circuit level is only one step ahead.

1.2 Standard Cell Characterization using Machine Learning

In our research, we make use of standard cell characterization to convey transistor-level degradations (such as aging or process variation) to the circuit level. For individual transistors, we can express the impact of degradations as changes in the parameters of the SPICE transistor model. Importantly, the changed electrical characteristics of individual transistors within a circuit will also impact the performance of the entire circuit, as indicated in Figure 3. Standard cell characterization simulates the behavior of individual transistors on the SPICE level and measures their impact on delay, power, and more for any standard cell. The resulting standard cell libraries can be used in a wide range of EDA tools to reflect the performance of the underlying technology on higher abstraction levels.

In circuit design, the impact of different operating conditions and degradations on the circuit can be investigated by performing Static Timing Analysis (STA) on a range of PVT corners. Each corner is represented by a standard cell library capturing all gates in the circuit under certain conditions, e.g., maximum temperature or maximum aging-induced degradations at the end of the lifetime of the chip. Depending on the considered corner, the circuit's path delays and power consumption can be very different. Therefore, to account for reliable operation under all circumstances, the path delays under worst-case assumptions are used to define the additionally needed *timing guardband* and the maximum clock frequency that the circuit can be operated at. However, in practice, these theoretical worst-case

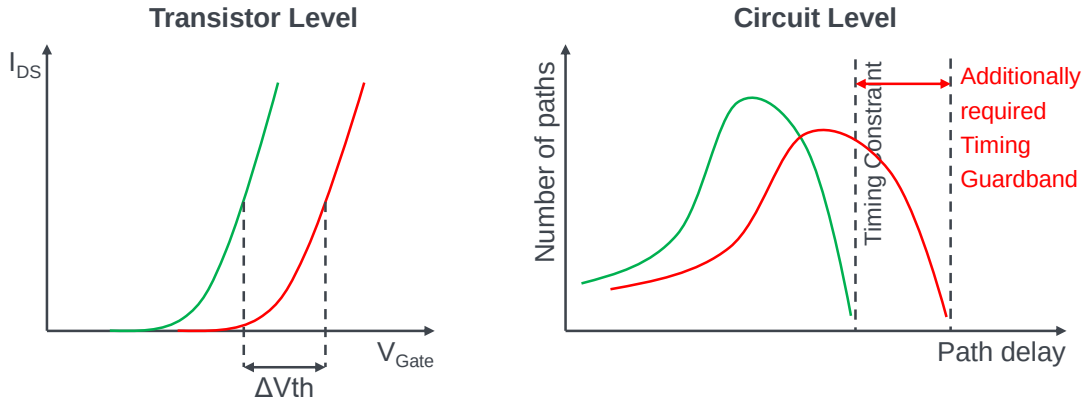

Figure 3: The two figurative graphs show how degradations on the transistor level will impact the path delays in a circuit. If these delays are not considered with an appropriate *guardband*, timing violations will occur, leading to the failure of the chip. One of the most prominent parameters to reflect degradation of the transistor is the *delta threshold voltage* (ΔV_{th}).

estimations are overly pessimistic. For example, the degradations induced by aging strongly depend on the workload of the chip, and thus, the experienced degradations in a real-world scenario will be significantly lower than under worst-case estimations [7]. As a result, parts of the timing guardband will have never been required, imposing an unnecessary reduction in the maximum clock frequency at design time and thus, a loss in performance.

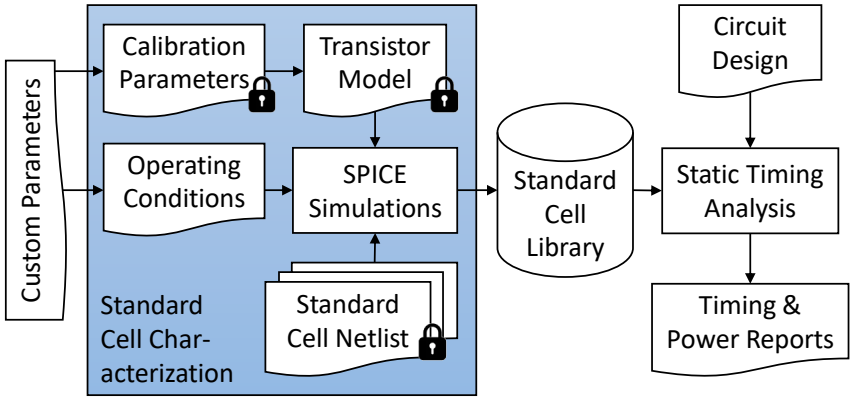

Figure 4: High-level flow showing how parameter changes to the transistors and operating conditions (e.g., temperature, supply voltage) impacts the timing and power reported on the circuit level. Within our research, we replace the blue box with a machine learning approach to accelerate standard cell characterization and eliminate the need of confidential data in the flow. Figure taken from [8].

To overcome the pessimism in worst-case estimations and achieve circuit designs closer to the edge, appropriate standard cell libraries for actually expected corners are needed. However, obtaining libraries for very specific conditions can be quite challenging. First and foremost, there is a wide area of possible applications for circuits, entailing different operating conditions, expected lifetimes, workloads, and so on. To cover all possible scenarios, the foundry would be required to provide hundreds of different standard cell libraries, posing a large characterization effort. Furthermore, for very specific libraries, e.g., including workload-specific aging-induced degradations, the characterization effort would be simply infeasible [7]. On the other hand, the task of cell characterization using SPICE simulations cannot be easily delegated to the designer since it requires confidential information about the foundry technology and standard cell designs. Necessary contracts and NDAs between the parties would exclude the vast amount of circuit designers from this option [9].

Both of the two major challenges in standard cell characterization for close-to-the-edge design can be tackled with machine learning (ML) techniques. First, time-consuming SPICE simulations can be replaced with fast ML inferences to overcome the infeasible computation effort of individual standard cell libraries. Second, many ML models offer implicit secrecy by abstraction and obfuscation of original design parameters, allowing for an easier sharing of the ML-based characterization model. Both points can be further strengthened by removing constant and unnecessary parameters from the characterization flow and thus, reducing the required complexity of the ML models. Looking at the traditional workflow shown in Figure 4, the entire characterization process (in blue) is replaced with our ML approach. From the hundreds of parameters needed to perform SPICE simulations, only a few "custom" parameters are selected to remain as *features* (inputs) to the ML (e.g., supply voltage, temperature, projected lifetime (for aging), etc). At the same time, we only require a subset of all the measurements generated by the SPICE simulations as our *labels* (outputs) of the ML: While SPICE is a powerful analog circuit simulator that reports detailed electrical responses of each transistor, only higher-level results such as cell delay or cell power consumption are needed to compile the cell library. This drastically reduces the required complexity of the ML models and also the amount of needed training data. Eventually, simple ML regression models (such as polynomial regression or k-nearest neighbors regression) already proved themselves capable of learning the dependency between selected design parameters and fully characterized libraries sufficiently [7]. While complex models might be able to further improve the accuracy of predictions, the low demand for training data is an important benefit of simple models, since the generations of training data can be very expensive in terms of computational effort and required software licenses.

With rapid standard cell characterization at hand, the possible applications are manifold: From investigation of process or aging variation [10], to accurate corners for workload-specific aging [7], to design technology co-optimization [9], or even hardware trojan detection [11]. More details on the development of our machine learning-driven standard cell characterization and the different applications are given in section 2.1.

2 Published Research, 15 works in Total (7 Journals, 8 Conferences)

With the support of the Graduate School Intelligent Methods for Test and Reliability (GS-IMTR) [12] at the University of Stuttgart, funded by Advantest, we published research at a variety of workshops, conferences, and journals. In publications in which Florian Klemme is the first author, research is usually heavily linked to the topic of P9. However, we also had the opportunity to contribute to publications of collaborations, which are usually more loosely related to the topic of P9. Authorship of Florian Klemme in cited publications is also highlighted in the references section at the end of this report, including 7 publications in journals, 8 publications at conferences, and 1 book chapter [13]. Since the start of P9, Florian Klemme has contributed to 11 publications (5 journals and 6 conferences), as depicted in Figure 5.

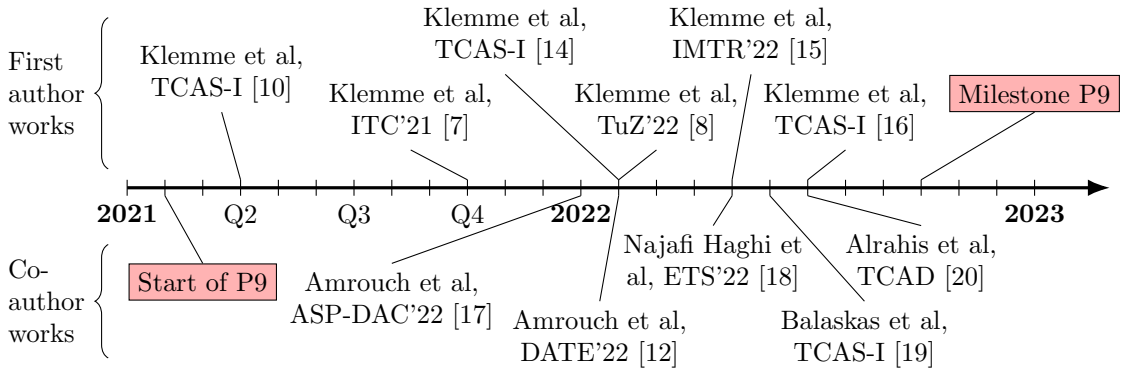

Figure 5: Timeline of GS-IMTR project 9, showing all first author publications (top) and contributions to other works (bottom) from the start of the project up to the milestone presentation date.

2.1 First-Author Publications

The limitations of conventional standard cell characterization tool flows are a major obstacle to our circuit-level reliability analysis. The process of standard cell characterization is time-consuming, dependent on confidential models and calibration data, and, importantly, does not support heterogeneous PVT corners across circuit designs. However, these aforementioned challenges have been tackled and presented with appropriate approaches in our publications, as described in the following.

Our work *Machine Learning for On-the-Fly Reliability-Aware Cell Library Characterization* [10], published in *IEEE Transactions on Circuits and Systems I: Regular Papers (TCAS-I)* sets the cornerstone for many following publications. It is the first work to introduce rapid cell library characterization using machine learning techniques to overcome infeasible simulation times for heterogeneous PVT corners. In early works, we often focus on aging-induced degradation in general rather than self-heating effects specifically, However, aging defects are also severely accelerated by transistor self-heating. During the lifetime of a chip, aging effects such as Bias Temperature Instability (BTI) and Hot Carrier Injection (HCI) steadily degrade the performance of transistors under stress. The change in the transistors' electrical characteristic is commonly modeled as an increase in threshold voltage (V_{th}), resulting in slower switching times. Therefore, the degradation of our transistors can be expressed as *delta threshold voltage* (ΔV_{th}). Figure 6 show an example of how transistors age, depending on temperature and operating voltage. Importantly, transistor self-heating must be considered on top of the probed chip temperature for accurate estimation of aging-induced degradations.

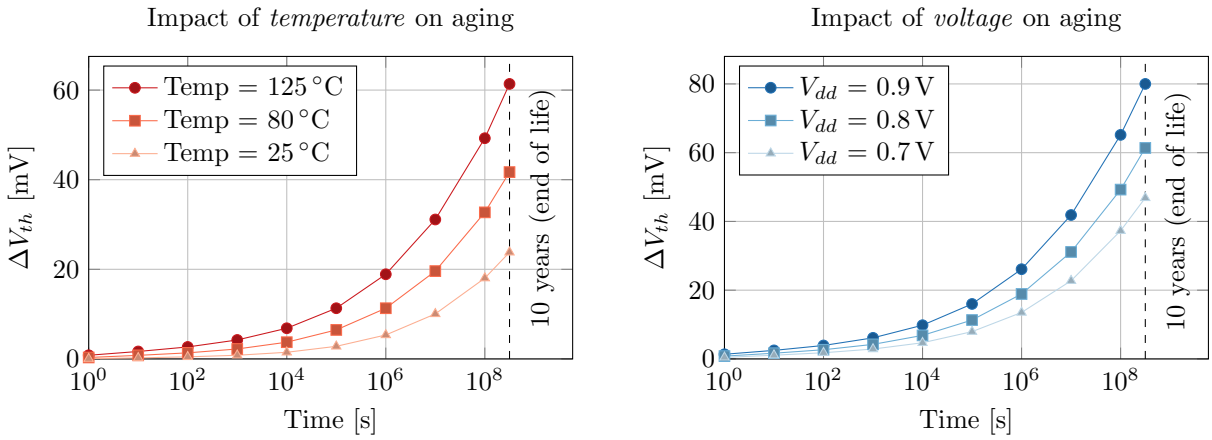

Figure 6: The depicted graphs show the impact of BTI and HCI-induced degradations on a 14 nm FinFET pMOS transistor. We see that ΔV_{th} is not only subject to time (aging effect) but also to V_{dd} and temperature. Therefore, a good selection of ΔV_{th} for guardband estimation is highly dependent operating conditions [24]. V_{dd} = 0.8 V and temperature = 125 °C are used in the top and in the bottom graph, respectively. Figure taken from [10].

Aging-induced degradation imposes a major challenge to the designer when estimating timing guardbands. This problem increases as traditional worst-case corners bring over-pessimism to designers, exacerbating competitive and close-to-the-edge designs. In [10], we present an accurate machine learning approach for aging-aware cell library characterization, enabling the designer to evaluate their circuit under the impact of precisely selected degradation. Unlike state-of-the-art, we bring cell library characterization to the designer, empowering their capability in exploring the impact of aging while protecting confidential information from the foundry at the same time. Furthermore, the fast inference of cell libraries makes it feasible, for the first time, to examine aging-induced variability analysis in a Monte-Carlo fashion. Finally, we show that the designer is able to select a less pessimistic timing guardband by choosing adequate delta threshold voltage (ΔV_{th}) for their design and their needs. The work completes the idea of employing machine-learning for standard cell characterization, first introduced in a previous work of ours [9].

Our work *Machine Learning for Circuit Aging Estimation under Workload Dependency* [7] (IEEE International Test Conference, ITC'21) builds on top of [10] by introducing heterogeneity in the circuit design. Previous state-of-the-art techniques only allow for the consideration of uniformly applied cir-

cuit degradation, despite the fact that different workloads will lead to different degradations due to the different induced activities. This imposes over-pessimism in estimating the required timing guardbands, resulting in unnecessary losses of performance and efficiency. In [7], we propose an approach that takes real-world workload dependencies into account and generates workload-specific aging-aware standard cell libraries. This allows for accurate analysis of circuits under the actual effect of aging-induced degradation. Machine-learning-driven standard cell characterization is employed to overcome infeasible simulation times for individual transistor aging while sustaining high accuracy. In our evaluation on the PULP microprocessor, we achieve predictions of workload-dependent aging-aware standard cells with an average accuracy (R^2 score) of 94.7 %. Using the predicted cell libraries in Static Timing Analysis, timing guardbands are reported with an error of less than 0.1 %. Figure 7 demonstrates that timing guardband requirements can be reduced by up to 21 % by considering specific workloads over worst-case analysis as performed in state-of-the-art tool flows.

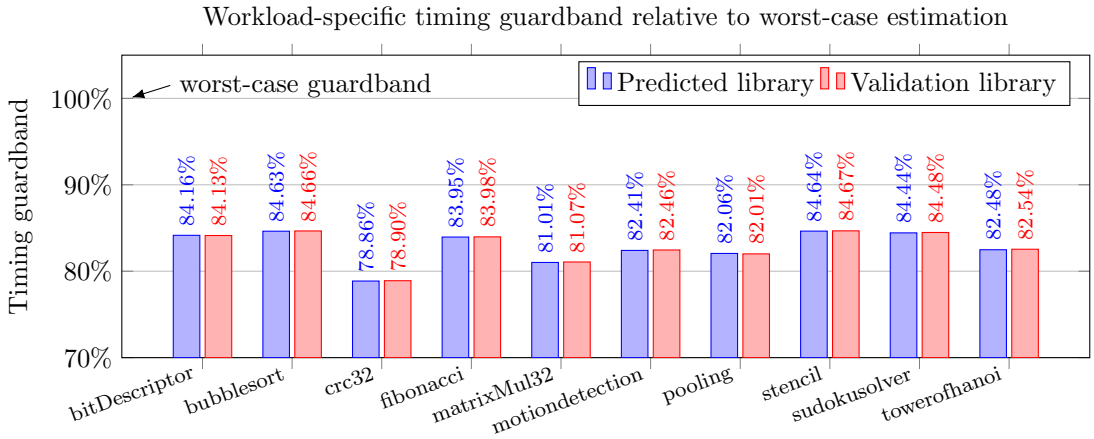

Figure 7: To ensure reliable operation, circuits must employ an additional timing guardband. The commonly employed worst-case guardband (100 %) can be safely reduced by considering accurate workload-induced degradations. Depending on the expected workload, the required timing guardband can be reduced by up to 21 %, enabling more efficient and performant circuit designs. Figure taken from [7].

In *Scalable Machine Learning to Estimate the Impact of Aging on Circuits Under Workload Dependency* [14] (TCAS-I), we extended upon [7] and set the focus on the scalability of our machine learning framework to previously unseen circuit designs. Only with ensured scalability, we can employ our framework for a wide range of circuit designs without repetitive re-training. To achieve scalability, we present multiple "coverage" metrics that can be applied to unseen circuits to estimate the prediction accuracy. Eventually, we show that our machine learning framework can predict cell libraries for unseen circuit design while sustaining an average error in timing guardband estimation below 1 %.

In *Efficient Learning Strategies for Machine Learning-Based Characterization of Aging-Aware Cell Libraries* [16] (TCAS-I), we tackle another performance problem of our machine learning framework. While we have been able to quickly predict standard cell libraries for a variety of scenarios, the benefits of our machine learning-based cell library characterization are strongly limited by its high demand in training data and the costly SPICE simulation required to generate the training samples. Therefore, efficient learning strategies are needed to minimize the required training data for machine learning models while still sustaining high prediction accuracy. In [16], we explore multiple active and passive learning strategies for machine learning-based cell library characterization with focus on aging-induced degradation. While random sampling and greedy sampling strategies operate with low computational overhead, active learning considers the performance of machine learning models to find the most valuable samples for training. We also introduce a hybrid approach of active learning and greedy sampling to optimize the trade-off between reduction in training samples and computational overhead. Our experiments demonstrate an achievable training data reduction of up to 77 % compared to the state of the art, depending on the targeted accuracy of the machine learning models, as shown in Figure 8.

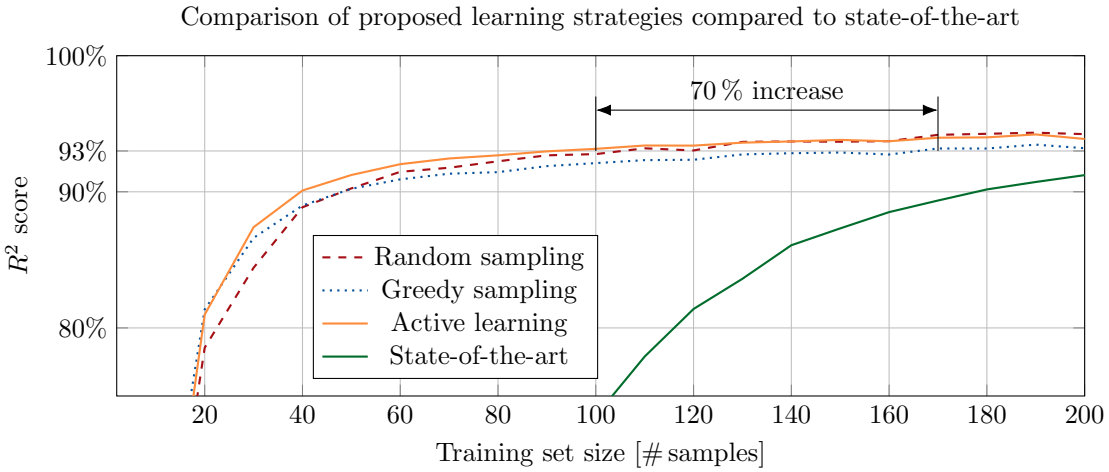

Figure 8: The required SPICE simulations to generate training data for our machine learning framework pose a large overhead. Therefore, we aim to learn as much as possible from as little data as possible. With our implementation and different proposed approaches, we can clearly outperform the approaches of our previous works (state-of-the-art). Even between the similar-looking approaches, there is a notable difference, as indicated in the figure. Figure taken from [16].

Besides the aforementioned conferences and journals, our work has been presented as part of an invited talk at the *27th Asia and South Pacific Design Automation Conference (ASP-DAC'22)* [17], at *The 34th Workshop on Test Methods and Reliability of Circuits and Systems (TuZ'22)* [8], and at the *Workshop on Intelligent Methods for Test and Reliability (IMTR'22)* [15].

2.2 Co-Authorship in Collaborations

The work *Design Close to the Edge for Advanced Technology using Machine Learning and Brain-Inspired Algorithms* [17] is a collaboration between P9 and P10 within the GS-IMTR and has been published as an invited special session paper at ASP-DAC'22. In advanced technology nodes, transistor performance is increasingly impacted by different types of design-time and run-time degradation. First, variation is inherent to the manufacturing process and is constant over the lifetime. Second, aging effects degrade the transistor over its whole life and can cause failures later on. Both effects impact the underlying electrical properties of which the threshold voltage is the most important. To estimate the degradation-induced changes in the transistor performance for a whole circuit, extensive SPICE simulations have to be performed. However, for large circuits, the computational effort of such simulations can become infeasible very quickly. Furthermore, the SPICE simulations cannot be delegated to circuit designers, since the required underlying transistor models cannot be shared due to their high confidentiality for the foundry. In [17], we tackle these challenges at multiple levels, ranging from transistor to memory to circuit level. We employ machine learning and brain-inspired algorithms to overcome computational infeasibility and confidentiality problems, paving the way towards design close to the edge. Our machine learning-based characterization flow, as one part of the work, is shown in Figure 9.

The work *On Extracting Reliability Information from Speed Binning* [18], published at the *IEEE European Test Symposium (ETS'22)* is a collaboration between P9 and PA1. The work utilizes standard cell library characterization to estimate the impact of small delay faults within larger circuit designs. Adaptive Voltage Frequency Scaling (AVFS) is an important means to overcome process-induced variability challenges for advanced high-performance circuits. AVFS requires and allows determining the maximum speed $F_{max}(V_{dd})$ reachable under a set of certain operation voltages V_{dd}. In [18], it is shown that the $F_{max}(V_{dd})$ measurements contain relevant data to identify some hidden defects in a chip that are reliability threats and can cause device failures, but pass the speed binning procedure within the given specifications. Static Timing Analysis (STA) is applied to a circuit designed by using standard cell libraries in which the underlying transistors along with process variations have been carefully calibrated

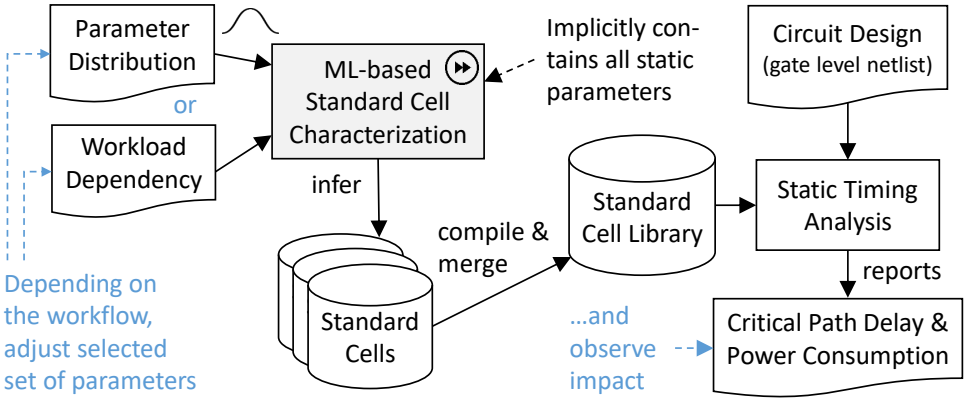

Figure 9: A very high-level perspective on the application of our machine learning-based characterization flow. Depending on the application, different parameters are swept in conventional SPICE simulation to generate the necessary training data for our framework. Afterward, it predicts individual cells, which are collected and merged into a full, compatible standard cell library. Finally, the generated standard cell library can be employed in conventional EDA tool flows to perform a variety of evaluations. Figure take from [17].

against industrial 14 nm FinFET measurement data, and instances with and without injected small resistive open defects are generated. From the slope of the function $F_{max}(V_{dd})$, a machine learning procedure can identify some defects with high precision and few false positives. These chips can be then discarded without any further need and cost for testing. It has to be noted that this reliability information comes for free from the data which is already generated and does not need any additional measurements.

The work *Variability-Aware Approximate Circuit Synthesis via Genetic Optimization* [19] (TCAS-I) is another work that utilizes our machine learning-driven characterization framework to generate a large set of standard cells under process variation, similar to [10]. Process variations are one of the major barriers that CMOS devices face at the nanometer scale due to manufacturing imperfections. They severely inhibit the reliable operation of circuits, as the operational frequency at the nominal process corner is insufficient to suppress timing violations across the entire variability spectrum. To avoid variability-induced timing errors, previous efforts impose pessimistic and performance-degrading timing guardbands atop the operating frequency. In [19], we employ approximate computing principles and propose a circuit-agnostic automated framework for generating variability-aware approximate circuits that eliminate process-induced timing guardbands. Variability effects are accurately portrayed with the creation of variation-aware standard cell libraries, fully compatible with standard EDA tools. The underlying transistors are fully calibrated against industrial measurements from Intel 14 nm FinFET in which both electrical characteristics of transistors and variability effects are accurately captured. In [19], we explore the design space of approximate variability-aware designs to automatically generate circuits of reduced variability and increased performance without the need for timing guardbands. Experimental results show that by introducing negligible functional error of merely 5.3×10^{-3}, our variability-aware approximate circuits can be reliably operated under process variations without sacrificing the application performance.

Finally, the work *GNN4REL: Graph Neural Networks for Predicting Circuit Reliability Degradation* [20], published in *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)* is our latest publication from collaboration, again utilizing our standard cell characterization flow for the generation of large training data set to be used in graph neural networks. Process variations and device aging impose profound challenges for circuit designers. Without a precise understanding of the impact of variations on the delay of circuit paths, guardbands, which keep timing violations at bay, cannot be correctly estimated. This problem is exacerbated for advanced technology nodes, where transistor dimensions reach atomic levels and established margins are severely constrained. Hence, traditional worst-case analysis becomes impractical, resulting in intolerable performance overheads. Contrarily, process-variation/aging-aware static timing analysis (STA) equips designers with accurate statistical delay distributions. Timing guardbands that are small, yet sufficient, can then be effectively estimated.

However, such analysis is costly as it requires intensive Monte-Carlo simulations. Further, it necessitates access to confidential physics-based aging models to generate the standard-cell libraries required for STA. In [20], we employ graph neural networks (GNNs) to accurately estimate the impact of process variations and device aging on the delay of any path within a circuit. Our proposed GNN4REL framework empowers designers to perform rapid and accurate reliability estimations without accessing transistor models, standard-cell libraries, or even STA; these components are all incorporated into the GNN model via training by the foundry. Specifically, GNN4REL is trained on a FinFET technology model that is calibrated against industrial 14 nm measurement data. Through our extensive experiments on EPFL and ITC-99 benchmarks, as well as RISC-V processors, we successfully estimate delay degradations of all paths – notably within seconds – with a mean absolute error down to 0.01 percentage points.

3 Conclusion and Future Work

Within the first 20 months of P9, we set many important cornerstones to investigate reliability challenges in advanced technologies. Most importantly, we demonstrated in multiple publications [10, 7, 14] how machine learning-driven standard cell characterization can be employed to efficiently elevate reliability threats from the transistor level to the circuit level. With our latest research on transistor self-heating [6], we will soon be able to also investigate self-heating effects on the circuit level and propose novel counteractive approaches. Our research will serve as an important tool to close the gap between self-heating-induced reliability challenges on the transistor and the circuit level, which allows for less pessimistic safety margins and thus more efficient circuit designs while sustaining full reliability of the chip.

References

[1] Om Prakash, Chetan K Dabhi, Yogesh S Chauhan, and Hussam Amrouch. "Transistor self-heating: The rising challenge for semiconductor testing". In: *2021 IEEE 39th VLSI Test Symposium (VTS)*. IEEE. 2021, pp. 1–7.

[2] Yi Zhao and Yiming Qu. "Impact of self-heating effect on transistor characterization and reliability issues in sub-10 nm technology nodes". In: *IEEE Journal of the Electron Devices Society* 7 (2019), pp. 829–836.

[3] Hai Jiang, SangHoon Shin, Xiaoyan Liu, Xing Zhang, and Muhammad Ashraful Alam. "The impact of self-heating on HCI reliability in high-performance digital circuits". In: *IEEE Electron Device Letters* 38.4 (2017), pp. 430–433.

[4] Lining Zhang, Debin Song, Ying Xiao, Xinnan Lin, and Mansun Chan. "On the formulation of self-heating models for circuit simulation". In: *IEEE Journal of the Electron Devices Society* 6 (2018), pp. 291–297.

[5] Hussam Amrouch, Victor M van Santen, Om Prakash, Hammam Kattan, Sami Salamin, Simon Thomann, and Jörg Henkel. "Reliability challenges with self-heating and aging in finfet technology". In: *2019 IEEE 25th International Symposium on On-Line Testing and Robust System Design (IOLTS)*. IEEE. 2019, pp. 68–71.

[6] **Florian Klemme** and Hussam Amrouch. "Upheaving Self-Heating Effects from Transistor to Circuit Level using Conventional EDA Tool Flows". In: *Design, Automation & Test in Europe Conference (DATE'23)*. (under review). IEEE. 2023.

[7] **Florian Klemme** and Hussam Amrouch. "Machine learning for circuit aging estimation under workload dependency". In: *2021 IEEE International Test Conference (ITC)*. IEEE. 2021, pp. 37–46.

[8] **Florian Klemme** and Hussam Amrouch. "Machine Learning for Reliability-Aware, yet Confidential Standard Cell Characterization". In: *The 34th Workshop on Test Methods and Reliability of Circuits and Systems (TuZ'22)*. Feb. 2022.

[9] **Florian Klemme**, Yogesh Chauhan, Jörg Henkel, and Hussam Amrouch. "Cell library characterization using machine learning for design technology co-optimization". In: *2020 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*. IEEE. 2020, pp. 1–9.

[10] **Florian Klemme** and Hussam Amrouch. "Machine learning for on-the-fly reliability-aware cell library characterization". In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 68.6 (2021), pp. 2569–2579.

[11] Virinchi Roy Surabhi, Prashanth Krishnamurthy, Hussam Amrouch, Kanad Basu, Jörg Henkel, Ramesh Karri, and Farshad Khorrami. "Hardware trojan detection using controlled circuit aging". In: *IEEE Access* 8 (2020), pp. 77415–77434.

[12] Hussam Amrouch, Jens Anders, Steffen Becker, Maik Betka, Gerd Bleher, Peter Domanski, Nourhan Elhamawy, Thomas Ertl, Athanasios Gatzastras, Paul R. Genssler, Sebastian Hasler, Martin Heinrich, André van Hoorn, Hanieh Jafarzadeh, Ingmar Kallfass, **Florian Klemme**, Steffen Koch, Ralf Küsters, Andrés Lalama, Raphael Latty, Yiwen Liao, Natalia Lylina, Zahra Paria Najafi-Haghi, Dirk Pflüger, Ilia Polian, Jochen Rivoir, Matthias Sauer, Denis Schwachhofer, Steffen Templin, Christian Volmer, Stefan Wagner, Daniel Weiskopf, Hans-Joachim Wunderlich, Bin Yang, and Martin Zimmermann. "Intelligent Methods for Test and Reliability". In: *Design, Automation & Test in Europe Conference & Exhibition, DATE 2022*. Mar. 2022, pp. 1–6.

[13] Victor M van Santen, **Florian Klemme**, and Hussam Amrouch. "The Vital Role of Machine Learning in Developing Emerging Technologies". In: *Recent Findings in Boolean Techniques*. Springer, 2021, pp. 29–57.

[14] **Florian Klemme** and Hussam Amrouch. "Scalable Machine Learning to Estimate the Impact of Aging on Circuits Under Workload Dependency". In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 69.5 (2022), pp. 2142–2155.

[15] **Florian Klemme** and Hussam Amrouch. "Mitigating the Complexity of Chip Designs with ML-based Cell Library Characterization". In: *Workshop on Intelligent Methods for Test and Reliability (IMTR'22)*. May 2022.

[16] **Florian Klemme** and Hussam Amrouch. "Efficient Learning Strategies for Machine Learning-Based Characterization of Aging-Aware Cell Libraries". In: *IEEE Transactions on Circuits and Systems I: Regular Papers* (2022).

[17] Hussam Amrouch, **Florian Klemme**, and Paul R. Genssler. "Design Close to the Edge in Advanced Technology using Machine Learning and Brain-Inspired Algorithms". In: *27th Asia and South Pacific Design Automation Conference (ASP-DAC'22)*. Invited. Jan. 2022.

[18] Zahra Paria Najafi-Haghi, **Florian Klemme**, Hussam Amrouch, and Hans-Joachim Wunderlich. "On Extracting Reliability Information from Speed Binning". In: *2022 IEEE European Test Symposium (ETS)*. IEEE. 2022, pp. 1–4.

[19] Konstantinos Balaskas, **Florian Klemme**, Georgios Zervakis, Kostas Siozios, Hussam Amrouch, and Jörg Henkel. "Variability-Aware Approximate Circuit Synthesis via Genetic Optimization". In: *IEEE Transactions on Circuits and Systems I: Regular Papers* (2022).

[20] Lilas Alrahis, Johann Knechtel, **Florian Klemme**, Hussam Amrouch, and Ozgur Sinanoglu. "GNN4REL: Graph Neural Networks for Predicting Circuit Reliability Degradation". In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (2022).

[21] Sami Salamin, Georgios Zervakis, **Florian Klemme**, Hammam Kattan, Yogesh Chauhan, Jörg Henkel, and Hussam Amrouch. "Impact of NCFET technology on eliminating the cooling cost and boosting the efficiency of Google TPU". In: *IEEE Transactions on Computers* 71.4 (2021), pp. 906–918.

[22] **Florian Klemme**, Jannik Prinz, Victor M van Santen, Jörg Henkel, and Hussam Amrouch. "Modeling emerging technologies using machine learning: challenges and opportunities". In: *Proceedings of the 39th International Conference on Computer-Aided Design*. 2020, pp. 1–9.

[23] Hussam Amrouch, Girish Pahwa, Amol D Gaidhane, Chetan K Dabhi, **Florian Klemme**, Om Prakash, and Yogesh Singh Chauhan. "Impact of variability on processor performance in negative capacitance finfet technology". In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 67.9 (2020), pp. 3127–3137.

[24] Hussam Amrouch, Seyed Borna Ehsani, Andreas Gerstlauer, and Jörg Henkel. "On the efficiency of voltage overscaling under temperature and aging effects". In: *IEEE Transactions on Computers* 68.11 (2019), pp. 1647–1662.