

GNN4REL: Graph Neural Networks for Predicting Circuit Reliability Degradation

Lilas Alrahis, *Member, IEEE*, Johann Knechtel, *Member, IEEE*, Florian Klemme, *Member, IEEE*, Hussam Amrouch, *Member, IEEE*, and Ozgur Sinanoglu, *Senior Member, IEEE*

Abstract—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 this work, 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 $14nm$ 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.

Index Terms—Graph neural networks, Standard-cell libraries, Static timing analysis, Transistor aging, Reliability estimation

I. INTRODUCTION

The rapid semiconductor-technology scaling has been a primary driver for the industry and an essential factor for making high-performance electronics widely available. However, technology scaling imposes enormous challenges when it comes to ensuring the reliability of integrated circuits (ICs) over the projected device lifetime. Advanced technology nodes show significant increase in the manufacturing process variability (process variation), making it challenging to predict, for any given IC, the *timing guardband* that is required to protect the

TABLE I
COMPARISON OF GUARDBAND ESTIMATION METHODS

Method	Requirements			Capabilities	
	Cell Libraries	Transistor Models	Timing Analysis	Process Variation	Aging
Conventional STA	One ¹	No	Yes	No	No
Statistical STA	One (LVF) ²	No	Yes	Yes	No
Monte-Carlo STA [3]	Many	Yes ³	Yes	Yes	Yes
GNN4REL	No	No	No	Yes	Yes

¹For instance, worst-case corner library. ²Liberty variation format.

³To generate the required libraries under variation.

design against timing violations [1].¹ Furthermore, device aging (runtime variation), which significantly degrades circuit lifetime and performance, becomes more dominant in nanoscale nodes. Without accounting for the impact of variation on the delay of circuit *timing paths*,² designers cannot guarantee a reliable circuit operation over the desired lifespan [2].

A. State-of-the-Art (SOTA) and Their Limitations

Next we outline some general limitations for existing pre-silicon methods for guardband estimation (See Table I).

Performance Overheads: In the worst-case scenario, designers may add pessimistic guardbands (i.e., worst-case margins), which results in excessive performance overheads and does not allow the circuit to operate at its full potential.

Lack of Consideration for Variations in STA: Static timing analysis (STA) serves to obtain the longest path within a circuit, in terms of signal propagation delay, considering a constant delay value per cell. Such a *critical path* defines the maximum clock speed the circuit can be operated at without causing timing violations. However, conventional STA cannot account for variation due to multiple reasons. First, the impact of variation on the delay of each cell strongly varies, as demonstrated in Fig. 1. As a result, variation cannot be sufficiently modeled by simply increasing the critical-path delay by a fixed amount. Second, since the additional delay due to variation will impact each path differently, the critical path needs to be determined with variation in mind. Neglecting

Manuscript received April 07, 2022; revised June 11, 2022; accepted July 05, 2022. This article was presented in the International Conference on Compilers, Architectures, and Synthesis for Embedded Systems (CASES) 2022 and appears as part of the ESWEK-TCAD special issue.

Lilas Alrahis, Johann Knechtel and Ozgur Sinanoglu are with the Division of Engineering, New York University Abu Dhabi, Abu Dhabi 129188, UAE (e-mail: lma387@nyu.edu; johann@nyu.edu; ozgursin@nyu.edu).

Florian Klemme and Hussam Amrouch are with the Department of Computer Science, University of Stuttgart, Stuttgart 70174, Germany (e-mail: klemme, amrouch@iti.uni-stuttgart.de).

¹A timing guardband acts as *safety margin* to prevent timing violations. It imposes an additional slack on top of the critical path delay. Hence, the circuit is clocked at a lower frequency leading to performance “loss”. Note that guardbands are required, thus the related impact on performance should not be considered a loss a priori. Rather the value of the guardband is important here: one wants a margin as small as possible, yet sufficient. Obtaining such optimized value is a challenge especially under the presence of variations.

²A timing path is a path between a start point (i.e., a netlist’s input port or the clock pin of a sequential element) and an end point (i.e., an output port or a data input pin of a sequential element).




Fig. 1: Process variation causes different delay variations for different standard-cells, switching transition times, and net capacitances. This distribution is extracted from 1,000 cell libs, characterized for different instances of process variation. μ and σ denote the mean and standard deviation, respectively.

variation when analysing the delay of circuit imposes the risk of not correctly identifying the critical path, as the example in Fig. 2 illustrates. Since the aforementioned effects are not captured by traditional STA methods, variations cannot be accounted for, and thus, relying on STA standalone leads to estimating overly optimistic timing analysis. Hence, errors due to timing violations can appear during the circuit operation.

Computational Cost and Other Limitations of SSTA: Statistical STA (SSTA) handles process variation by considering statistical delay distributions rather than constant delay values [4], [5]. To enable SSTA, cell libraries (libs) containing variability information are required.³ Hence, SSTA can determine the critical path with respect to its sensitivity to variation. However, SSTA entails significant computational complexity, e.g., Monte-Carlo SPICE simulations to obtain the variability information for the process technology. Although SSTA is a significant step towards better timing guardbands, it still holds some limitations: (i) with only the average (μ) and standard deviation (σ) information stored in the lib, SSTA can *only* model Gaussian variability distributions; (ii) only the critical path and its variability is examined and then reported, although the actual critical path (in the presence of variability) is likely to change from a circuit instance to another. This will give a limited/skewed perspective when investigating the actual distribution of critical-path delays under process variation [3].

Need for Standard (Std)-Cell Libs and Transistor Models: To realize small, yet sufficient, practically relevant guardbands, note that both STA and SSTA require the generation of variation- and aging-aware std-cell libs. In other words, only once variation-aware libs are employed, then STA can evaluate the delay of paths under the effects that such degradations have. In turn, this requirement necessitates that designers have access to confidential transistor models from the foundry (i.e., physics-based aging models that well describe the impact of aging on the transistor parameters). Such access is not always attainable, especially not for advanced nodes.

B. Key Research Challenges

The above review (Sec. I-A) raises an important question: *how to predict process variation- and aging-induced degradation in an efficient manner without relying on std-cell libs,*

³Variability denotes the σ of the delay distribution of a given cell divided by μ . The liberty variation format (LVF) extends the industry-standard liberty format with corresponding σ for each contained timing information.

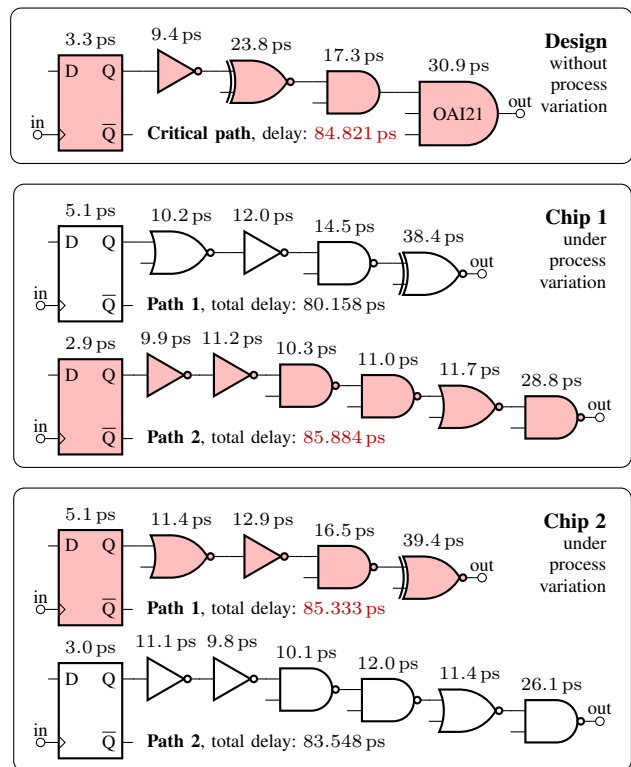


Fig. 2: Process variation and aging impact the location of the critical path. Even in a small circuit such as b01 from ITC-99, the critical path can change rapidly and differently for each chip. The delay analysis here was done using variation-aware libs characterized using SPICE simulations.

transistor models, or even STA? Developing such a versatile reliability assessment framework is an open research problem that poses the following research challenges.

- 1) *Extensive full-circuit analysis:* The delay variation for each cell is impacted, among others, by its driving cell and load capacitances, i.e., by its location within the netlist. Thus, delay variation is not easy to account for, e.g., neither using a fixed delay increase per cell nor with a narrow analysis of a *single* critical path. Fig. 2 depicts how the critical path of a circuit can change due to variations, as in featuring a path that was originally critical and then becomes non-critical whereas another path becomes critical instead.
- 2) *Handling different types of degradation:* Depending on the type of considered degradation, different cell libs capturing the corresponding effects need to be generated first. The methods of characterizing process variation and aging-induced degradations are very different, challenging to set up, and expensive in terms of required simulation runtime. Thus, a generic framework that can handle various types of degradation, without requiring the designer to generate specific cell libs beforehand, is desired.

C. Our Novel Concept and Contributions within this Work

To address the challenges outlined above, we propose GNN4REL, a generic circuit-reliability assessment framework based on graph learning, which is the first of its kind. Graph learning and related graph neural networks (GNNs) are particularly promising here, as this general approach

leverages a naturally matching representation of circuits, unlike other traditional approaches for machine learning. It is noteworthy that graph learning and GNNs have shown remarkable achievements for other tasks related to circuits, e.g., reverse engineering of unknown designs [6] and evaluation of different design-for-trust techniques [7]–[11].

The key goal of this work is to train a GNN to predict delay degradations induced by process variation as well as aging for any given timing path within a circuit. We formulate the problem of estimating delay degradations as a regression problem and solve the problem using a GNN, as outlined in Fig. 3. This novel framework and its important contributions are supported by the following technical contributions:

- 1) We develop a framework for the path-to-graph conversion (Sec. III-A), which extracts timing paths from any gate-level netlist and represents each path as a subgraph. These subgraphs capture the Boolean functionalities of gates and their directed connectivity within the path. As a result, our platform captures and accounts for the driving cells, load capacitances, and the position/integration of gates in the netlist when making model predictions.
- 2) We build a GNN-based regression model that learns on subgraphs extracted around timing paths (Sec. III-B), which automatically extracts the relevant features of paths that help in estimating the delay degradation. The model can be trained by the foundry itself, thereby accounting for the manufacturing procedure and aging parameters, and then be shared with the designers to predict the performance degradations of their circuits. Thus, our platform simplifies the reliability assessment procedure for the designers (i.e., eliminating the need to build std-cell libs and run STA), while protecting confidential foundry information.

We demonstrate the effectiveness of GNN4REL on selected ITC-99 and EPFL benchmarks as well as RISC-V processors. Without loss of generality, our GNN framework is trained on a FinFET technology model that is carefully calibrated against $14nm$ measurement data from Intel obtained from [12]. Note that the calibrations were done for both transistor characteristics as well as variation. Our extensive experiments considering different variation sources and various dataset scenarios show that GNN4REL achieves excellent prediction performance when predicting delay degradations, reporting a mean absolute error (MAE) down to 0.01 percentage points (average of 0.76). We further release GNN4REL [13].

The remainder of the paper is organized as follows. Sec. II presents the background information, while Sec. III presents the concept and implementation details of GNN4REL. In Sec. IV, the technology calibration and std-cell libs generation (required for training GNN4REL) are presented. Sec. V discusses the experiments and results, while Sec. VI covers the related work. Finally, Sec. VII presents the conclusions.

II. BACKGROUND

In this section, we present the necessary background information about process variation, device aging, and GNNs.




Fig. 3: Predicting path delay degradation using graph neural networks (GNNs).

A. Process Variation and Device Aging

Process variation occurs due to imperfections in the manufacturing process of the chip. It is a time-independent source of variation that differs for each fabricated chip (and even within the chip itself). Random dopant fluctuation, metal gate granularity, and line-edge roughness are among the typical sources of variation that are considered for state-of-the-art FinFET devices [14]. An accurate estimation of process variation is prerequisite to ensure high yield at high performance. If the impact of process variation is underestimated, many fabricated chips will not pass chip testing and the yield is reduced.

In addition to process variation, transistors will also suffer from *aging-induced degradation*. Transistor aging is a time-dependent source of variation that depends on many conditions such as temperature, workload, and projected lifetime. Physical effects that lead to transistor aging include hot-carrier injection (HCI), and importantly, bias temperature instability (BTI) [2]. BTI is one of the dominant contributors to aging-induced degradations. Over the lifetime of the chip, electrical charges get trapped in the gate oxide of transistors, resulting in increased threshold voltages. In addition, interface traps can be generated at the $Si-SiO_2$ layer resulting in more undesired charges and hence further increase in the threshold voltage. Consequently, transistor switching times and circuit path delays increase.

We model and include both types of variation (i.e., process variation and aging) in the standard electronic design automation (EDA) tool flows, as later described in Sec. IV.

B. Graph Neural Networks (GNNs)

The GNN formalism is a dominant paradigm for deep learning with graph structured data. GNNs take a graph as an input and generate an *embedding* (1D vector) for each node in the graph through *neighborhood aggregation/message passing*. The key concept is that the generated embeddings depend on (i) the structure of the graph and (ii) any feature information associated with the nodes. Similar nodes in the graph should be close in terms of distance in the embedding space.

Let $G(V, E)$ be a directed attributed graph; V denotes the set of nodes, and E denotes the set of edges. Each node in the graph $v \in V$ is associated with a feature vector (x_v) that captures its properties. During neighborhood aggregation (*aggregate*), each node receives information (*messages*) from its neighboring nodes $N(v)$, and a new embedding is computed for each node by applying a learnable update function (*update*) on the node's current embedding and the aggregated messages. After L aggregation rounds, each node in the graph is represented by an embedding that captures its properties, the properties of its neighborhood, and its position/integration within the graph. The final node embeddings are then utilized to perform the desired




Fig. 4: The different steps of the proposed GNN4REL framework.

Algorithm 1 Pseudo-code for the proposed GNN4REL platform

Require: Netlist N , hop-size h
Ensure: Delay Degradation (D)

```

1:  $P \leftarrow \text{PathExtraction}(N)$ 
2:  $D \leftarrow \{\emptyset\}$ 
3: for  $p \in P$  do
4:    $S_p \leftarrow (v \in p)$ 
5:    $G_{(S_p, h)} \leftarrow \text{SAMPLE}(G, h, S_p)$ 
6:    $D.append(\text{GNN}(G_{(S_p, h)}))$ 
7: end for
8: return  $D$ 
9: procedure  $\text{GNN}(G)$ 
10:   $Z^{(L)} \leftarrow \text{PNA}(G)$ 
11:   $y_G \leftarrow \text{READOUT}(Z^{(L)})$ 
12:   $\hat{y} \leftarrow \text{MLP}(y_G)$ 
13:  return  $\hat{y}$ 
14: end procedure

```

task, such as node classification, graph classification, etc. The neighborhood aggregation procedure is abstracted as follows, where $z_v^{(l)} \in \mathbb{R}^f$ indicates the embedding of node v at the l -th aggregation round, and $\mathbf{a}_v^{(l)}$ denotes the aggregated information from $N(v)$ at the l -th layer. $Z^{(L)} \in \mathbb{R}^{n \times f}$ indicates the final 2D node embedding matrix, where $n = |V|$.

$$\mathbf{a}_v^{(l)} = \text{aggregate}^{(l)} \left(\left\{ z_u^{(l-1)} : u \in N(v) \right\} \right) \quad (1)$$

$$z_v^{(l)} = \text{update}^{(l)} \left(z_v^{(l-1)}, \mathbf{a}_v^{(l)} \right) \quad (2)$$

For graph-level tasks, a graph embedding y_G that captures the underlying properties of the graph is obtained by applying a *readout* function to the node embeddings. The readout function is typically an order-invariant function such as summing up the node embeddings (i.e., row-wise additions on $Z^{(L)}$).

GNN models mainly differ based on their aggregation and update functions, with the *mean*, *sum*, and *maximum* aggregation functions being the most adopted in the state-of-the-art architectures [15]–[18]. Different aggregation functions perform better on different tasks [15]. Therefore, in our work, we employ the state-of-the-art principal neighborhood aggregation (PNA) model that implements multiple aggregation strategies instead of a single aggregation function to improve the performance of the GNN model [19]. More details regarding the employed PNA architecture are given in Sec. III-B.

III. OUR PROPOSED GNN4REL FRAMEWORK

In this section, we provide an overview of the GNN4REL framework (summarized in Fig. 4 and Algorithm 1).




Fig. 5: Proposed path-to-subgraph transformation.

A. Path-to-Subgraph Transformation

We represent each gate-level netlist as a directed graph $G = (V, E)$, where V represents the set of nodes (i.e., gates, primary inputs (PIs), and primary outputs (POs)), while E represents the set of edges (i.e., interconnects). We chose a directed representation to capture the direction of the timing paths (i.e., start point, end point, and order of gates). Each node in the graph $v \in V$ is initialized with a feature vector x_v that captures its properties (more on that in Sec. III-A2). $X \in \mathbb{Z}^{n \times k}$ is the 2D matrix containing node features, where k denotes the length of the feature vector.

1) *Subgraph Extraction*: First, timing paths are extracted from the gate-level netlist and grouped into set P (line 1 in Algorithm 1). Then, the nodes forming a specific path $p \in P$ are grouped into set S_p (i.e., $\cup\{v \in p\}$) (line 4). We refer to the path nodes as *target nodes*.

Given the netlist graph G , an h -hop enclosing subgraph $G_{(S_p, h)}$ is extracted around the target nodes (line 5). Let $d(i, j)$ represent the shortest path distance between nodes i and j , then $G_{(S_p, h)}$ is extracted from G by $\cup_{j \in S_p} \{i \mid d(i, j) \leq h\}$. We extract an h -hop subgraph around the target path to capture the position of the path within the netlist and collect information regarding the driving cell and load capacitance of the gates in the path, which all impacts delay.

An example of 1-hop subgraph extraction is illustrated in Fig. 5. The extracted timing path highlighted with a dashed green line includes the gates $\{G2, G3, G4\}$, which form the set S_p (see 1). To extract a 1-hop subgraph around the target path, all the 1-hop neighbors of the gates in S_p are included in the subgraph, alongside the target gates (see 2). The list of 1-hop neighbors includes $\{DFF1, DFF2, G5, G1\}$.

2) *Feature Vector*: x_v is a one-hot encoded vector that represents the node's Boolean functionality, or, indicates, if applicable, that it is a PI or a PO. The length of the feature vector k depends on the number of Boolean functions available in the target std-cell lib. The feature vector example in Fig. 5 indicates that node $G4$ is a NAND gate.

B. Employed Graph Neural Network Model

We employ the state-of-the-art PNA GNN model to perform graph-level regression [19] (line 6 in Algorithm 1). In this context, a graph represents the extracted subgraph around a target timing path. The PNA employs four statistical aggregators, i.e., μ , *maximum* (max), *minimum* (min), and σ , so that each node is aware of the distribution of its incoming messages.

The aggregation functions are listed below. $Z^{(l)}$ are the nodes' embeddings at layer l . *ReLU* is the rectified linear unit




Fig. 6: The principal neighborhood aggregation (PNA) architecture employs four static aggregators and three degree scalars [19].

used to avoid negative values caused by numerical errors and ϵ is a small positive number to ensure σ is differentiable.

$$\mu_v(Z^{(l)}) = \frac{1}{|N(v)|} \sum_{u \in N(v)} z_u^{(l)} \quad (3)$$

$$\max_v(Z^{(l)}) = \max_{u \in N(v)} z_u^{(l)} \quad (4)$$

$$\min_v(Z^{(l)}) = \min_{u \in N(v)} z_u^{(l)} \quad (5)$$

$$\sigma_v(Z^{(l)}) = \sqrt{\text{ReLU}(\mu_v(Z^{(l)2}) - \mu_v(Z^{(l)})^2) + \epsilon} \quad (6)$$

Degree scalars allow the network to attenuate or amplify signals based on the degree of each node, i.e., the number of messages being aggregated. PNA uses the logarithmic scaler $S(d, \alpha)$ presented below, where α is a variable parameter that is positive for amplification, negative for attenuation, or zero for no scaling. δ is a normalization parameter computed over the training set, and d denotes the degree of the target node.

$$S(d, \alpha) = \left(\frac{\log(d+1)}{\delta} \right)^\alpha, \quad d > 0, \quad -1 \leq \alpha \leq 1 \quad (7)$$

Note that the PNA model employs a logarithmic scaler instead of a linear scaler since the latter would cause an exponential amplification of both the aggregated information and the gradients after multiple GNN layers. Such an exponential amplification, in turn, would reduce the ability of a GNN to generalize to unseen, possibly larger graphs [19].

The degree scalars are combined with the aggregator functions as follows, where \otimes represents the tensor product and I represents the identity matrix (i.e., no scaling).

$$\oplus = \underbrace{\begin{bmatrix} I \\ S(D, \alpha = 1) \\ S(D, \alpha = -1) \end{bmatrix}}_{\text{scalers}} \otimes \underbrace{\begin{bmatrix} \mu \\ \sigma \\ \max \\ \min \end{bmatrix}}_{\text{aggregators}} \quad (8)$$

A PNA layer used in GNN4REL can be abstracted as follows, where M and U are linear layers.

$$z_v^{(l+1)} = U \left(z_v^{(l)}, \bigoplus_{(u,v) \in E} M(z_v^{(l)}, z_u^{(l)}) \right) \quad (9)$$

A diagram for the PNA layer is illustrated in Fig. 6, where MLP represent a multi-layer perceptron. After L PNA layers (line 10 in Algorithm 1), a readout layer is added to obtain a graph-level embedding y_G (line 11), which gets passed to an MLP to generate the prediction \hat{y} (line 12). Details regarding the number and dimension of layers are included in Sec. V-A3.

C. Dataset Generation

We consider three scenarios of dataset generation to demonstrate the generic nature of our proposed platform. In all scenarios, the setup requires the generation of a dataset from which a training and a validation set are extracted. A dataset contains a list of timing paths extracted from gate-level netlists. The training and validation sets include labeled timing paths (i.e., known delay-degradation percentages), while the testing set includes unlabelled timing paths.

1) *Self-Referencing*: The timing paths of a single design are split based on an 81:10:9 training:validation:testing ratio. We expect GNN4REL to achieve the best prediction performance in this setup as it captures the design characteristics during training (without seeing the exact testing paths).

2) *Single-Design*: The timing paths of a specific design, referred to as X , are used for training and validation based on a 90:10 training:validation ratio, and the timing paths of another design, referred to as Y , are used for testing. The goal is to show that GNN4REL can generalize to unseen designs. For example, GNN4REL can be trained on the b14 benchmark and then evaluated on the rest of the ITC-99 benchmarks.

3) *Design Dataset*: We further consider a design dataset (i.e., a collection of gate-level netlists) for training GNN4REL. The design dataset does not include the target design but includes designs with a similar design structure. For example, when predicting the degradation of the b17 benchmark from ITC-99, only the timing paths of b14, b15, b20, b21, and b22 are used for training:validation based on a 90:10 ratio.

IV. TECHNOLOGY CALIBRATION AND STD-CELL LIBRARIES CREATION

We utilize mature 14 nm FinFET technology and commercial std-cell characterization tool flows to ensure the generation of accurate training data for realistic circuit timing analysis. This requires careful calibration of the underlying transistor model against silicon measurements, as well as the characterization a wide variety of std-cell libs to capture the technology for different operating conditions. As a result, we generate std-cell libs for the technology under typical conditions, under the impact of process variation, and under the impact of transistor aging as exhibited at the end of the lifetime of a chip.

A. Technology Calibration

For our analog SPICE simulations, the industry-standard compact model for FinFET technology (BSIM-CMG) [21] is used as underlying transistor model. All parameters of the compact model are carefully calibrated to reproduce measurements of the Intel production-quality 14 nm manufacturing process. Transistor measurements for validation are extracted from [12]. Fig. 7 demonstrates the excellent agreement between SPICE simulation results obtained with our calibrated transistor model and Intel measurement data. As can be seen from the figure, validation is performed for both n-type and p-type FinFET transistors, as well as for multiple $I_{ds}-V_{gs}$ and $I_{ds}-V_{ds}$ biases, to ensure a holistic representation of the technology and accurate simulation for all required corner cases. In addition, the compact model is further calibrated to reflect




Fig. 7: The industry compact model for FinFET (BSIM-CMG) is carefully calibrated to reproduce Intel 14 nm measurement data extracted from [12]. As demonstrated in the plots, SPICE simulations (using our calibrated models) achieve an excellent agreement with the measurement data for both nFinFET and pFinFET devices. The top figure (a) shows the validation for the case of I_{ds} - V_{gs} at high and low V_{ds} biases. The arrows in the figure indicate which curves belong to which y-axis. Bottom figure (b) shows the validation of I_{ds} - V_{ds} at various V_{gs} biases. Details on our calibration available in [20].

the technology under process variation. The considered sources of variation include the gate length (L_g), fin thickness (T_{fin}), fin height (H_{fin}), SiO_2 equivalent gate dielectric thickness (EOT), and the work-function of the gate (ϕ_g). Fig. 8 demonstrates the variability calibration as an I_{on} - I_{off} plot with regression lines obtained from Monte-Carlo SPICE simulations. The Intel reference regression line is once again extracted from [12]. As shown, our calibrated variability parameters are in excellent agreement with measurements of 14 nm FinFET variability.

B. Std-Cell Library Generation

With an accurately calibrated transistor model, we can characterize a full std-cell lib using accurate SPICE simulations. To this end, the commercial *HSPICE* from *Synopsys* tool flows was used. The required std-cell netlists are obtained from the NanGate 15 nm open-source cell lib [22]. All std-cell netlists are also annotated with post-layout parasitic resistances and capacitances. A commercial characterization tool flow [23] is then employed to instruct extensive SPICE simulations, determining all characteristics of the std-cells, including signal




Fig. 8: Variability calibration of our FinFET compact model against Intel 14 nm measurement data [12]. The regression line obtained from Monte-Carlo SPICE simulations on the transistor model is in good agreement with the data from the variability measurements. Details on our calibration available in [20].




Fig. 9: Cell library generation workflow for process variation and aging.

propagation delays, transition times, pin capacitances, switching power, and leakage power. The resulting std-cell lib is stored in the non-linear delay model (NLDM)-format, in which each data point is represented by a 7×7 -matrix to account for the different input signal slews and output load capacitances an std-cell can experience, depending on its position in the circuit. Without any further adjustments to the transistor model or cell netlists, the resulting cell lib reflects the manufacturing process under typical conditions and is suitable for circuit synthesis.

To generate std-cells for different instances of process variation, the transistor model and the std-cell netlists are prepared to accept parameter overwrites for each individual transistor instance. Afterward, each transistor in each cell netlist is annotated with a random set of variability parameters (L_g , T_{fin} , H_{fin} , EOT , ϕ_g) following the expected distribution of the parameters under process variation. With the annotated cell netlists, std-cell lib characterization is performed as usual. The obtained cell lib reflects each std-cell under the impact of one instance of process variation. Afterward, the aforementioned steps are repeated in a Monte-Carlo fashion, applying new random variability parameters to the std-cell netlists in each iteration. The resulting set of cell libs form a collection of std-cells, characterized under different instances of process variation. With such a collection, accurate STA for entire circuits under process variation can be achieved.

The gate-level netlist of a circuit is obtained by performing logic synthesis with the baseline std-cell lib (i.e., the reference

TABLE II
PROPERTIES FOR GRAPH REPRESENTATION OF EACH NETLIST

Dataset	ITC-99						EPFL					RISC-V		
Benchmark	b14	b15	b20	b21	b22	b17	adder	max	bar	square	multiplier	divisor	zero-riscy	RISCY
# Edges (# Interconnects)	9,630	13,627	21,442	21,268	32,711	39,774	3,126	4,530	5,028	23,739	52,123	143,261	35,718	83,652
# Nodes (# Gates + Pls + POs)	4,569	5,790	10,090	9,921	15,236	16,296	1,846	2,730	1,748	13,687	25,265	72,823	15,032	35,648
Timing Constraint (ns)	0.46	0.54	0.53	0.53	0.51	0.7	0.42	1	0.35	1.3	3	25	1.2	1.2




Fig. 10: The overall architecture of GNN4REL for reliability-degradation prediction. We utilize 4 PNA graph convolution layers [19], separated by batch normalization and ReLU. A set of linear layers translates the hidden PNA output to a single value that resembles the predicted delay degradation.

lib in which neither process variation nor transistor aging is applied). Afterward, each cell instance in the gate-level netlists is replaced by a random version of that cell under process variation. All characterized cells under process variation are suffixed with an individual index and merged into one large cell lib. The entire workflow is also outlined in Fig. 9.

For aging-aware std-cell libs, the approach is very similar to the above. However, instead of the variability parameters, the threshold voltage (V_{th}) of each transistor is increased, to reflect the impact of aging. In the industry-standard compact model for FinFET transistors (BSIM-CMG), the $dvtshift$ parameter is used to model the major impact of aging. In addition, we also adjust the C_{it} parameter to control capacitance change due to interface traps. This is, in fact, needed to also reflect the aging-induced degradation in the sub-threshold swing (SS) of the transistor. With all parameters annotated in the cell netlists, cell lib characterization is performed to generate the corresponding aging-aware cell libs. These cells can be inserted into gate-level netlists to enable accurate STA for circuits under the impact of aging.

V. EVALUATION AND COMPARISONS

A. Experimental Setup

Next, we describe the experimental setup in detail.

1) *Benchmark Designs*: We evaluate GNN4REL on selected ITC-99 and EPFL benchmarks, alongside RISC-V processors [24], [25]. For the latter, we consider two different configurations of an in-order central processing unit (CPU) as obtained using the open-source RISC-V core generator: a baseline CPU *RISCY* as well as a lighter version *zero-riscy*. *RISCY* implements a 32bit, 4-stage CPU, while *zero-riscy* features a 32bit, 2-stage CPU.

2) *Synthesis and Generation of Datasets*: Benchmarks are synthesized using *Synopsys Design Compiler (DC)* considering the “fresh std-cell lib” (i.e., the lib that has been characterized in the absence of variation). Such synthesis setup provides us with “fresh netlist” of every benchmark. We employ std-cell libs for the 14nm FinFET technology node calibrated with Intel 14nm FinFET measurements [12] (see Sec. IV). During synthesis, highest efforts for delay minimization have been targeted by using “compile_ultra” for *Synopsys DC*. The properties for

the graph representations of the synthesized gate-level netlists, alongside the timing constraints, are listed in Table II. *Synopsys PrimeTime* is used to perform STA to generate the required datasets for GNN training. The aging-aware and variation-aware std-cell libs are generated as discussed in Sec. IV.

3) *Subgraph Extraction and GNN Training*: We implement the netlist-to-subgraph conversion in *Perl* scripts and the subgraph extraction in *Python* scripts. The overall architecture of GNN4REL is illustrated in Fig. 10.

We use the *PyTorch Geometric* implementation of PNA for graph-level regression, using four PNA GNN layers (*PNAConv*) with an input/output channel size of 75 each. A batch normalization layer (*BatchNorm*) follows each *PNAConv* layer to standardize the embeddings to a mean of zero and a variance of one. A *ReLU* layer follows each *BatchNorm* layer. Following the sequence of graph convolution layers, the *global_add_pool Torch* function is used to generate a batch-wise graph-level representation by averaging node features across the node dimension. The graph representation is then passed to the following sequence of layers to obtain the final prediction: *Linear*(75, 50), *ReLU*, *Linear*(50, 25), *ReLU*, and *Linear*(25, 1). The *Linear*(*input_size*, *output_size*) layers apply a linear transformation to the incoming data. For the *PNAConv* layers, we set the aggregators to $\{\mu, \sigma, \max, \min\}$ and the scalers to $\{\textit{identity}, \textit{amplification}, \textit{attenuation}\}$.

To reduce the computational complexity of GNN4REL, the concept of *towers* is employed in PNA as in the message passing neural networks (MPNNs) [17]. Using towers, the f -dimensional node embedding $z_v^{(l)}$ is broken down into t different f/t -dimensional embeddings $z_v^{(l,f)}$. We set the number of towers to 5. We train GNN4REL on the 1-hop subgraphs for 500 epochs using Adam optimizer with a learning rate of 0.001 and batch size of 32. The model with the lowest MAE loss on the validation set is selected as the final model.

Hyperparameters: Since we propose a generic reliability assessment approach, we like to avoid fine-tuning/over-tuning the GNN model parameters such as the hidden dimensions, learning rate, etc., to obtain best performance for some given benchmarks. Instead, the robust results are obtained using the original parameter values of the PNA network proposed in [19]. Upon the release of our GNN4REL model, readers/users may




Fig. 11: Path-lengths distributions for the ITC-99 benchmarks. The dashed line represents the average value.




Fig. 12: Path-lengths distributions for the EPFL benchmarks. The dashed line represents the average value.

tune the parameters as interesting to them. Note that we study the effect of h on the performance of GNN4REL in Sec. V-E.

Fixed Architecture: The PNA GNN shows sound performance while analyzing the circuits (represented as graphs) and automatically extracting features that are suitable for the desired tasks. Therefore, we do not need to perform any manual feature engineering or modify the network model, when considering different regression tasks (i.e., aging-induced versus process-variation-induced degradation). Rather, we only change the output information (i.e., label) and let the GNN framework perform feature engineering. We argue that this simple approach highlights the generic nature of our proposed method; the GNN4REL model can be easily trained to solve different tasks.

4) *Dataset Generation and Evaluation for Prediction of Process Variation:* We consider selected ITC-99 and EPFL benchmarks (total of 12 netlists). Using STA, we obtain the required timings for 1,000 extracted timing paths from each synthesized netlist – which form the reference for delay-degradation computations. The timing paths are extracted based on the 1,000 endpoint flip-flops with the worst slack. Note that we limit the number of extracted timing paths to 1,000 for simplicity. However, more or less timing paths can be extracted from a design based on its size and the requirements of the designer. The distributions of path lengths (#gates) for the ITC-99 and EPFL benchmarks are shown in Fig. 11 and Fig. 12, respectively.

Using an in-house *Perl* script, we generate 100 versions of each netlist by replacing each std-cell in the netlist with random and uniform sampling across the corresponding (i.e., functionally equivalent) options in the variation-aware std-cell lib (see Sec. IV), resulting in 1,200 netlists and 120,000 timing paths in total. Next, we perform STA for the same 1,000 timing paths in the variation-affected netlists and compute the delay-degradation percentages (i.e., the relative delay increase) by comparing the required timings after variation with the baseline timings of the paths. The average (μ), standard deviation (σ), and the maximum (max) delay-degradation percentages are computed based on the STA results for a given netlist.

For an illustrative example, we randomly select six timing paths for the ITC b14 benchmark and plot the paths' delay distributions caused by process variation in Fig. 13. As can be




Fig. 13: Distributions of process-variation-induced delay degradations for selected paths of b14. The dashed line represents the average value.

observed, each path has a unique distribution, evidencing the need for degradation prediction per path basis.

We evaluate the prediction performance of GNN4REL based on the three dataset generation scenarios discussed in Sec. III-C. The related setup specifics are as follows:

- Regarding the self-referencing scenario (Sec. III-C1), the 1,000 timing paths of each design are split into 810 training paths, 100 validation paths, and 90 testing paths.
- For the training on a single-design scenario (Sec. III-C2), we consider the ITC-99 benchmarks and train GNN4REL using the b14 benchmark (without loss of generality). In this case, 900 paths are used for training and 100 paths are used for validation (both sets of paths extracted from b14). The 1,000 timing paths of the design under-evaluation (e.g., b15, b20, etc.) are used for testing. In this case, the b14 benchmark is excluded from the evaluation since GNN4REL is trained on all its extracted timing paths.
- For the design-dataset scenario (Sec. III-C3), one of the considered ITC-99 benchmarks will be kept for testing (i.e., 1,000 testing paths), while the rest of the considered ITC benchmarks are used for training/validation, resulting in 4,500 training paths and 500 validation paths.

5) *Dataset Generation and Evaluation for Prediction of Aging:* We consider the ITC-99, EPFL, and RISC-V benchmarks for this experiment. We extract 1,000 timing paths from each netlist along with the baseline delay values. We run STA considering the aging-aware std-cell lib for all the benchmarks and obtain the delay-degradation percentage for each timing path. As with the case of process variation, we evaluate the prediction performance of GNN4REL based on the three dataset generation scenarios discussed in Sec. III-C.

6) *Evaluation Metric:* For all experiments, the prediction performance of GNN4REL is reported using the MAE and mean absolute percentage error (MAPE) metrics, where $MAE = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|$ and $MAPE = \frac{1}{N} \sum_{i=1}^N \left| \frac{y_i - \hat{y}_i}{y_i} \right|$, where N (#testing paths) and y (true degradation%).

B. Prediction of Process Variation

The μ -degradation predictions by GNN4REL for every path per selected designs in the self-referencing scenario are shown in Fig. 14. The predictions are scatter-plotted versus the actual degradation percentages obtained from the accurate STA results. A line representing the outcome of an ideal regression model is plotted to visualize the MAE error. Here, the MAE error denotes the average distance between the prediction points to the regression line (i.e., average absolute residual error).

GNN4REL predicts the μ -degradation with an average MAE of 0.3, where the actual degradation percentages range from -0.5% to 2% . In the case of the ITC-99 benchmarks,




Fig. 14: Path-level regression for the average (μ) process variation degradation of selected ITC-99 and EPFL benchmarks under the self-referencing scenario. Note that the ideal-performance curve is obtained using the Monte-Carlo STA [3].

TABLE III

MAE OF PATH-LEVEL REGRESSION FOR PROCESS VARIATIONS ON ITC-99 BENCHMARKS, DIFFERENT TRAINING DATASETS. NA: NOT APPLICABLE

Benchmark	Self-Referencing			Training on b14			Design Dataset		
	μ	σ	max	μ	σ	max	μ	σ	max
b14	0.58	0.15	0.94	NA	NA	NA	0.65	0.30	1.34
b15	0.22	0.07	0.51	0.76	0.18	1.88	0.53	0.15	1.32
b20	0.56	0.09	0.93	0.71	0.29	1.36	0.61	0.14	1.13
b21	0.46	0.07	0.67	0.75	0.26	1.32	0.57	0.14	1.13
b22	0.50	0.07	0.59	0.85	0.19	1.60	0.63	0.15	1.02
b17	0.25	0.05	0.35	0.72	0.96	1.97	0.43	0.97	2.13

TABLE IV

MAE OF PATH-LEVEL REGRESSION FOR PROCESS VARIATIONS ON EPFL BENCHMARKS. SELF-REFERENCING SCENARIO

Benchmark	adder	multiplier	square	bar	max	divisor
μ	0.22	0.02	0.23	0.35	0.06	0.15
σ	0.07	0.01	0.05	0.08	0.04	0.04
max	0.59	0.13	0.55	0.23	0.19	0.33

GNN4REL performs particularly well for the cases of b15 and b17 benchmarks, reaching an MAE of 0.22 and 0.25, respectively, and showing a strong correlation between the predictions and the actual results. The MAE values of the μ -, σ -, and max-degradation predictions by GNN4REL for the ITC-99 benchmarks under all dataset generation scenarios are listed in Table III. Further, the MAE values of the μ -, σ -, and max-degradation predictions for the EPFL benchmarks under the self-referencing scenario are listed in Table IV.

Considering the ITC-99 benchmarks and the self-referencing scenario, GNN4REL achieves excellent prediction performance concerning σ -degradation, with an average MAE of 0.08, where the actual σ values range between 0.92 to 3.05 (average MAPE error of 4%). GNN4REL predicts the max-degradation with an average MAE of 0.66, where the actual max values range between 1.92 and 9.09 (average MAPE error of 12%).

Regarding the single-design scenario (b14 benchmark in this study, without loss of generality), GNN4REL achieves an average MAE of 0.75, 0.37, and 1.62 when predicting the μ , σ , and max values, respectively. This experiment shows that, even if the GNN does not utilize the under-evaluation design during training, it can make valuable predictions. Such capabilities are essential in practice, where training data for the design under-evaluation may not be available. However, considering only a single design for training does limit the type of samples the model sees and can lead to high variance.

Therefore, we further evaluate GNN4REL when using a design dataset for training; such approach is equally valid for scenarios where training data for the design under-evaluation itself are not available. For this case (training on a design

TABLE V

MAE PATH-LEVEL REGRESSION RESULTS FOR AGING-INDUCED DEGRADATION ON SELECTED ITC-99 BENCHMARKS, DIFFERENT TRAINING DATASETS. NA MEANS NOT APPLICABLE

Benchmark	Self-Referencing	Training on b14	Design Dataset
b14	0.39	NA	2.63
b15	0.55	3.38	1.46
b20	1.12	3.52	1.62
b21	1.15	3.52	1.47
b22	0.39	1.77	1.48
b17	0.31	8.10	4.29

TABLE VI

MAE PATH-LEVEL REGRESSION RESULTS FOR AGING-INDUCED DEGRADATION OF EPFL BENCHMARKS. SELF-REFERENCING SCENARIO

Benchmark	adder	multiplier	square	bar	max	divisor
Self-Referencing	0.37	0.03	0.51	1.71	0.18	0.61

dataset), GNN4REL achieves an average MAE of 0.56, 0.3, and 1.34 for predicting the μ , σ , and max values, respectively. The results show that this scenario achieves better prediction performance compared with the single-design scenario. In short, in case where training data for the design-under-evaluation are unavailable, the designer/foundry can (i) utilize a generic design dataset for training or (ii) track the properties of the timing paths in the testing set and accordingly train the model on similar paths, to enhance the prediction performance as needed.

C. Aging Prediction

The predictions for runtime-variation degradation for every path and selected designs in the self-referencing scenario are shown in Fig. 15. The predictions are again scatter-plotted versus the actual degradation percentages. The outcomes of an ideal regression model are also scatter-plotted, to visualize the MAE error. GNN4REL predicts the runtime-variation degradation percentages with an average MAE of 0.65, where the actual degradation percentages fall between 15% and 26%. The average MAPE value reported by GNN4REL is 3.17%, which indicates an excellent prediction performance.

The MAE values for the runtime-variation degradations predicted by GNN4REL under all dataset scenarios are listed in Table V. We can observe the same trend as with the process variation prediction: GNN4REL performs best in case of the self-referencing scenario (average MAE of 0.65) and worst in the single-design dataset scenario (average MAE of 4). Still, even in this relative worst-case scenario, GNN4REL achieves an average MAPE value of 20%, which indicates that it is performing “good forecasting” according to the study in [26].

For the design-dataset scenario, which “sits” between the other two scenarios in terms of prediction performance,




Fig. 15: Path-level regression for aging-induced degradation of selected ITC-99 and EPFL benchmarks under the self-referencing scenario. Note that the ideal-performance data points are obtained using the Monte-Carlo STA [3].

GNN4REL achieves an average MAE of 2.15 and an average MAPE of 9.7%, indicating on highly accurate forecasting.

We further evaluate GNN4REL under the self-referencing scenario on the EPFL benchmarks and report the MAE results in Table VI. GNN4REL achieves an average MAE of 0.64. This experiment shows that our platform can be generally applied to different types of designs.

For the RISC-V processors, we extract 1,000 timing paths from the RISC-V processor and train GNN4REL accordingly. Then, 1,000 timing paths are extracted from the zero-risky processor to estimate their delay degradation.⁴ GNN4REL reports an MAE of 3, demonstrating the scalability of the platform also for such more complex designs.

D. Scalability Analysis

We demonstrated the scalability of GNN4REL on complex designs such as RISC-V processors represented using graphs with up to 35,648 nodes and 3,652 edges. Moreover, we considered EPFL benchmarks represented using graphs with up to 72,823 nodes and 143,261 edges, as summarized in Table II. Recall that the number of nodes represents the total number of gates, PIs and POs in the corresponding design.

Increasing the size of the design does not undermine the performance of GNN4REL. Even on the contrary, larger designs can improve the prediction performance. E.g., predicting the average process-variation-induced delay degradation on the smallest considered ITC-99 benchmark, b14 with 9,630 nodes, GNN4REL achieves an MAE of 0.58, whereas when performing the same task on the largest considered ITC-99 benchmark, b17 with 39,774 ($4\times$ larger than b14), GNN4REL achieves the best performance with an MAE of 0.25.

Further, the training time incurred by GNN4REL grows only linearly (with factor < 1) with the size of the considered design. For example, training on b14 takes 01:42:18 (h:m:s) while training on b17 takes 02:36:42, i.e., a factor only $1.6\times$ training time although the circuit size increased to $4\times$. Additionally, in case longer training times are not desired, we have demonstrated how GNN4REL can also be trained on small designs such as b14 and perform inference on larger designs such as b17. In such as setup, GNN4REL reported an MAE of 0.72, which is $2.9\times$ higher than the MAE observed for self-reference training.

E. Effect of the Extracted Subgraph Size h

We study the effect of h -hop sampling on the performance of GNN4REL. We repeat the experiments for predicting the

⁴Thus, we consider the design-dataset scenario here, but specifically for RISC-V designs with a model trained separately from those used for the earlier experiments on ITC and EPFL benchmarks.

TABLE VII
GNN4REL PREDICTION PERFORMANCE IN TERMS OF MAE FOR DIFFERENT h -HOP NUMBERS UNDER THE SELF-REFERENCING SCENARIO

Benchmark	μ of Design-Time Degradation			Runtime Degradation		
	$h=0$	$h=1$	$h=2$	$h=0$	$h=1$	$h=2$
b14	0.60	0.58	0.59	0.40	0.39	0.75
b15	0.37	0.22	0.44	0.69	0.55	0.71
b20	0.56	0.56	0.58	1.11	1.12	1.18
b21	0.53	0.46	0.55	1.13	1.15	1.18
b22	0.54	0.50	0.66	0.32	0.39	0.79
b17	0.21	0.25	0.35	0.59	0.31	0.57
Average	0.47	0.43	0.53	0.71	0.65	0.87

TABLE VIII
TRAINING TIME OF THE PROPOSED GNN4REL PLATFORM

ITC Benchmark	b14	b15	b20	b21	b22	b17
Training Time	01:42:18	01:37:37	01:45:59	01:49:33	01:48:19	02:36:42
EPFL Benchmark	adder	multiplier	square	bar	max	divisor
Training Time	03:07:19	19:24:36	06:43:34	04:31:17	07:31:42	120:01:05

μ delay degradation (due to process variation) and predicting the end-of-life delay degradation (due to aging) for varying $h \in [0, 2]$ with a step size of 1. We consider the ITC-99 benchmarks for this experiment under the self-referencing scenario. See Table VII for the MAE results. A hop size of $h = 0$ indicates that only the gates within each timing path itself are represented in the extracted subgraphs. As can be observed from Table VII, the prediction performance of GNN4REL improves by increasing h from 0 to 1. We argue that more information regarding the fan-in and fan-out structures of the timing path gates get captured in the 1-hop subgraphs allowing for a better estimate of the delay degradation. However, moving to a hop size of $h = 2$, the prediction performance drops. With the increase in subgraph size, the properties of the timing path can get lost in the vast information captured by the graph, degrading the performance of the prediction.

F. Runtime Analysis

Std-Cell Lib Generation: Each cell in the std-cell lib is characterized under different input-signal slews and output-load capacitances settings, typically 7×7 . Further, each rise and fall condition for every input pin is considered. Hence, characterizing the entire std-cell lib involves a huge number of SPICE simulations, which is a very time-consuming process. This challenge is exacerbated when having to repeat the process for Monte-Carlo-like std-cell characterization under variations. For instance, characterizing 100 std-cell libs takes ≈ 48 hours on a modern high-capacity server (using one SPICE license).

STA: Run on the 1,000 paths extracted for the largest considered design (i.e., the *divisor* benchmark with 143,261 gates, PIs and POs) takes 80 seconds. Running STA to compute

the delay degradation (due to process variation) considering the 100 libs takes up to 2.2 hours for each design.

Training and Inference: We report the training time of GNN4REL on the ITC-99 and EPFL benchmarks (the details of the benchmarks are summarized in Table II) for predicting the delay degradation caused by aging under the self-referencing scenario in Table VIII – GNN4REL takes merely ≈ 2 hours to train. Recall that in the self-referencing scenario, the 1,000 timing paths of each design are split into 810 training paths, 100 validation paths, and 90 testing paths. Also recall that training GNN4REL is a one-time effort. The subgraph-extraction time is part of the total training time, which could be sped up using parallelism. Once GNN4REL is trained, it can be used to assess the reliability of any given design in the considered testing set. *The inference stage of GNN4REL, i.e., the actual reliability prediction, takes a few seconds.* The experiments are performed on an Intel(R) Xeon(R) CPU X5680 with 64GB of RAM.

VI. RELATED WORK

A. Learning-based Delay Degradation Prediction

Recently, different machine learning (ML)-based methods were developed to estimate delay degradation. However, these methods are limited in their capabilities, as we showcase next.

S. M. Ebrahimipour *et al.* [27] proposed an aging-aware delay model, termed *Aadam*, tailored for generic cell libs. In *Aadam*, a separate feed-forward, fully-connected neural network (FFNN) is trained for each cell in the lib, to capture the relation between a number of aging factors and the cell's delay degradation. During both the training and inference stages, *Aadam* first passes the gate-level netlist to a logic simulator to compute the signal probabilities for each transistor inside each gate. Then, the respective FFNN networks are invoked inside an STA tool, to infer the aging-induced delay of the circuit. Thus, *Aadam* eliminates the need for the generation of aging-aware std-cell libs. The main shortcomings of *Aadam* compared to GNN4REL are as follows. First, in GNN4REL, a single model is trained to predict degradation for different circuits containing various std-cells, unlike for *Aadam*, which requires training of as many networks as cells are in the lib. Second, GNN4REL eliminates the need for STA during inference, unlike *Aadam*. Third, *Aadam* requires invoking circuit simulations during inference, unlike for GNN4REL, which passes the netlist directly to the trained model, without requiring any simulations during inference. Fourth, process variations are not considered in [27].

F. Klemme *et al.* [3], [28] proposed ML-based cell-lib characterization methods, which can be invoked by STA tools, to obtain aging-induced and process-variation-induced degradation. Similar to *Aadam*, these methods do not eliminate the need for STA to compute the degradation. More recently, J. Guo *et al.* [29] proposed an ML-based platform for predicting path-delay variations. Their platform requires the user to first compute the nominal delay for each path in the netlist and then uses this data for an input feature. Thus, similar to the other SOTA methods above, their platform also requires conventional STA methods at inference time, unlike GNN4REL.

VII. CONCLUSION

We present GNN4REL, a machine learning-based generic platform for circuit-reliability assessment. GNN4REL empowers circuit designers to obtain fast and accurate estimations of the delay degradation imposed on their designs due to process variation and device aging. Further, GNN4REL takes the burden of generating variation-aware standard-cell libraries and running static timing analysis off the designer's shoulders, while protecting confidential foundry information.

Our experimental evaluation on selected ITC-99 and EPFL benchmarks, alongside RISC-V processors, shows that given a timing path, GNN4REL accurately predicts the delay degradation distributions' measures (i.e., mean, standard deviation, maximum) caused by process variation and device aging – with a mean absolute error down to 0.01 percentage points – within few seconds. Considering different dataset and training scenarios, we show that GNN4REL can operate under various setups based on the requirements of the designer. All in all, we believe that GNN4REL opens up new frontiers in advancing design-time reliability assessment methods. We will release GNN4REL as open-source framework to the community.

ACKNOWLEDGMENT

This work is supported in part by the Center for Cyber Security (CCS) at New York University Abu Dhabi (NYUAD). Besides, this work is also supported by Advantest as part of the Graduate School “Intelligent Methods for Test and Reliability” (GS-IMTR) at the University of Stuttgart. We would like to thank Sami Salamin for his valuable support the RISC-V processor experiments.

REFERENCES

- [1] K. J. Kuhn *et al.*, “Process technology variation,” *IEEE Transactions on Electron Devices*, vol. 58, no. 8, pp. 2197–2208, 2011.
- [2] H. Amrouch *et al.*, “Reliability-aware design to suppress aging,” in *ACM/IEEE Design Automation Conference (DAC)*, 2016, pp. 1–6.
- [3] F. Klemme and H. Amrouch, “Machine learning for on-the-fly reliability-aware cell library characterization,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 68, no. 6, pp. 2569–2579, 2021.
- [4] B. Li *et al.*, “On timing model extraction and hierarchical statistical timing analysis,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 32, no. 3, pp. 367–380, 2013.
- [5] V. Khandelwal and A. Srivastava, “A quadratic modeling-based framework for accurate statistical timing analysis considering correlations,” *IEEE transactions on very large scale integration (VLSI) systems*, vol. 15, no. 2, pp. 206–215, 2007.
- [6] L. Alrahis *et al.*, “GNN-RE: Graph neural networks for reverse engineering of gate-level netlists,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, pp. 1–1, 2021.
- [7] —, “MuxLink: Circumventing learning-resilient MUX-locking using graph neural network-based link prediction,” *IEEE Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2022.
- [8] —, “UNTANGLE: Unlocking routing and logic obfuscation using graph neural networks-based link prediction,” in *IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, 2021, pp. 1–9.
- [9] —, “OMLA: An oracle-less machine learning-based attack on logic locking,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, 2021.
- [10] —, “GNNUnlock+: A systematic methodology for designing graph neural networks-based oracle-less unlocking schemes for provably secure logic locking,” *IEEE Transactions on Emerging Topics in Computing*, no. 01, pp. 1–1, 2021.
- [11] L. Alrahis, S. Patnaik, F. Khalid, M. A. Hanif, H. Saleh, M. Shafique *et al.*, “Gnnunlock: Graph neural networks-based oracle-less unlocking scheme for provably secure logic locking,” in *IEEE Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2021, pp. 780–785.

- [12] S. Natarajan, M. Agostinelli, S. Akbar, M. Bost, A. Bowonder, V. Chikarmane *et al.*, "A 14nm logic technology featuring 2 nd-generation finfet, air-gapped interconnects, self-aligned double patterning and a 0.0588 μm^2 sram cell size," in *IEEE International Electron Devices Meeting*, 2014, pp. 3–7.
- [13] L. Alrahis. (2022) GNN4REL datasets. [Online]. Available: <https://github.com/lilasrahis/GNN4REL>
- [14] Z. Zhang *et al.*, "Extraction of process variation parameters in finfet technology based on compact modeling and characterization," *IEEE Transactions on Electron Devices*, vol. 65, no. 3, pp. 847–854, 2018.
- [15] K. Xu *et al.*, "How powerful are graph neural networks?" *arXiv preprint arXiv:1810.00826*, 2018.
- [16] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.
- [17] J. Gilmer *et al.*, "Neural message passing for quantum chemistry," in *International Conference on Machine Learning-Volume*. JMLR. org, 2017, pp. 1263–1272.
- [18] P. Veličković *et al.*, "Neural execution of graph algorithms," *arXiv preprint arXiv:1910.10593*, 2019.
- [19] G. Corso *et al.*, "Principal neighbourhood aggregation for graph nets," *Advances in Neural Information Processing Systems*, vol. 33, pp. 13 260–13 271, 2020.
- [20] H. Amrouch, G. Pahwa, A. D. Gaidhane, C. K. Dabhi, F. Klemme, O. Prakash *et al.*, "Impact of variability on processor performance in negative capacitance finfet technology," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 67, no. 9, pp. 3127–3137, 2020.
- [21] S. Venugopalan *et al.*, "Bsim-cmg 110," 2016. [Online]. Available: <http://bsim.berkeley.edu/models/bsimcmg/>
- [22] Silvaco, Inc., "Silvaco and si2 release unique, free 15nm open-source digital cell library," 2019. [Online]. Available: https://www.silvaco.com/news/pressreleases/2019_05_30_01.html
- [23] Synopsys, Inc., "Primelib user guide," 2022.
- [24] "Chippyard's documentation." [Online]. Available: <https://chippyard.readthedocs.io/>
- [25] "PULP Platform." [Online]. Available: <https://pulp-platform.org/implementation.html>
- [26] C. D. Lewis, *Industrial and business forecasting methods: A practical guide to exponential smoothing and curve fitting*. Butterworth-Heinemann, 1982.
- [27] S. M. Ebrahimpour *et al.*, "Aadam: A fast, accurate, and versatile aging-aware cell library delay model using feed-forward neural network," in *IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, 2020, pp. 1–9.
- [28] F. Klemme, Y. Chauhan, J. Henkel, and H. Amrouch, "Cell library characterization using machine learning for design technology co-optimization," in *IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, 2020, pp. 1–9.
- [29] J. Guo *et al.*, "Novel prediction framework for path delay variation based on learning method," *Electronics*, vol. 9, no. 1, p. 157, 2020.



Lilas Alrahis is a Postdoctoral Associate at New York University Abu Dhabi. She received the M.Sc. degree and the Ph.D. degree in electrical and computer engineering from Khalifa University, UAE, in 2016 and 2021, respectively. Her research interests include Hardware Security, Design for Trust, Logic Locking, and Applied Machine Learning. She won the MWSCAS Myril B. Reed Best Paper Award in 2016 and the Best Paper Award at the Applied Research Competition held in conjunction with Cyber Security Awareness Week, in 2019. Dr. Alrahis is

currently serving as Associate Editor of the Integration, the VLSI Journal.



Johann Knechtel is a Research Scientist with New York University Abu Dhabi, United Arab Emirates. He received the M.Sc. degree in Information Systems Engineering (Dipl.-Ing.) and the Ph.D. degree in Computer Engineering (Dr.-Ing., summa cum laude) from TU Dresden, Germany, in 2010 and 2014, respectively. His research interests cover VLSI physical design automation, with particular focus on emerging technologies and hardware security. He was a Postdoctoral Researcher with the Masdar Institute of Science and Technology, Abu Dhabi, from 2015–2016. From 2010 to 2014, he was a Ph.D. Scholar with the DFG Graduate School on "Nano- and Biotechnologies for Packaging of Electronic Systems" at TU Dresden. In 2012, he was a Research Assistant with the Chinese University of Hong Kong, Hong Kong. In 2010, he was a Visiting Research Student with the University of Michigan at Ann Arbor, MI, USA.



Florian Klemme (M'20) is a Doctoral Researcher at the Chair of Semiconductor Test and Reliability (STAR), University of Stuttgart. He received the B.Sc. in System Integration from the University of Applied Sciences Bremerhaven, Germany, in 2014 and the M.Sc. in Computer Science from the Karlsruhe Institute of Technology, Germany, in 2018. He is currently working towards the Ph.D. degree at the Chair of Semiconductor Test and Reliability, University of Stuttgart. His research interests include cell library characterization and machine learning techniques in electronic design automation and computer-aided design. He is a member of the IEEE. ORCID 0000-0002-0148-0523.



Hussam Amrouch (S'11-M'15) is a Jun.-Professor heading the Chair of Semiconductor Test and Reliability (STAR) within the Computer Science, Electrical Engineering Faculty at the University of Stuttgart as well as a Research Group Leader at the Karlsruhe Institute of Technology (KIT), Germany. He currently serves as Editor at the Nature Scientific Reports Journal. He received his Ph.D. degree with the highest distinction (Summa cum laude) from KIT in 2015. His main research interests are design for reliability and testing from device physics to systems, machine learning for CAD, HW security, approximate computing, and emerging technologies with a special focus on ferroelectric devices. He holds eight HiPEAC Paper Awards and three best paper nominations at top EDA conferences: DAC'16, DAC'17 and DATE'17 for his work on reliability. He has served in the technical program committees of many major EDA conferences such as DAC, ASP-DAC, ICCAD, etc. and as a reviewer in many top journals like Nature Electronics, T-ED, TCAS-I, TVLSI, TCAD, TC, etc. He has around 185 publications (including 74 journals) in multidisciplinary research areas across the entire computing stack, starting from semiconductor physics to circuit design all the way up to computer-aided design and computer architecture. His research in HW security and reliability have been funded by the German Research Foundation (DFG), Advantest Corporation, and the U.S. Office of Naval Research (ONR).



Ozgur Sinanoglu is a professor of electrical and computer engineering at New York University Abu Dhabi. He obtained his Ph.D. in Computer Science and Engineering from University of California San Diego. He has industry experience at TI, IBM and Qualcomm, and has been with NYU Abu Dhabi since 2010. During his Ph.D. he won the IBM Ph.D. fellowship award twice. He is also the recipient of the best paper awards at IEEE VLSI Test Symposium 2011 and ACM Conference on Computer and Communication Security 2013. Prof. Sinanoglu's research interests include design-for-test, design-for-security and design-for-trust for VLSI circuits, where he has more than 200 conference and journal papers, and 20 issued and pending US Patents. Prof. Sinanoglu is the director of the Center for CyberSecurity at NYU Abu Dhabi. His recent research in hardware security and trust is being funded by US National Science Foundation, US Department of Defense, Semiconductor Research Corporation, Intel Corp, and Mubadala Technology.