# Chapter 11

# PDNSim: An Open-source Static PDN Analysis Tool

Vidya A. Chhabria

PDN analysis is crucial to successful IC design closure, particularly for designs implemented in lower technology nodes that suffer from large wire parasitics and high power densities. As described in Chapter 3, PDN analysis involves IR drop estimation and EM current density analysis. These analyses are computationally expensive and traditionally take several hours to perform using traditional EDA tools.

While ML has found some success in addressing this problem [19, 22, 40, 81], the license-based system from today's commercial EDA tool vendors has created a scalability challenge for the adoption of ML techniques as they require tremendous amounts of ground-truth training data. With each PDN simulation taking several hours in industry-scale designs with large power grids with billions of nodes, creating a training dataset with hundreds of datapoints can quickly run into several days with a single license.

This chapter presents PDNSim [39], an open-source static PDN analyzer, developed as a part of the OpenROAD project which creates a fully autonomous, open-source toolchain for digital layout generation with a focus on the RTL-to-GDSII. Unlike commercial EDA tool counterparts that are closed source, PDNSim is open-source and integrated with the OpenROAD application [125] with a permissive BSD 3-clause license. The open-source nature of PDNSim addresses the scalability challenge for ground-truth training data generation as the permissive BSD-3 license allows running multiple simulations of PDN analysis in parallel. Further, while commercial EDA tools typically serve large design houses in taping out chips, PDNSim serves a different audience [189]. It provides accessible source code for PDN-related algorithmic to EDA researchers as it

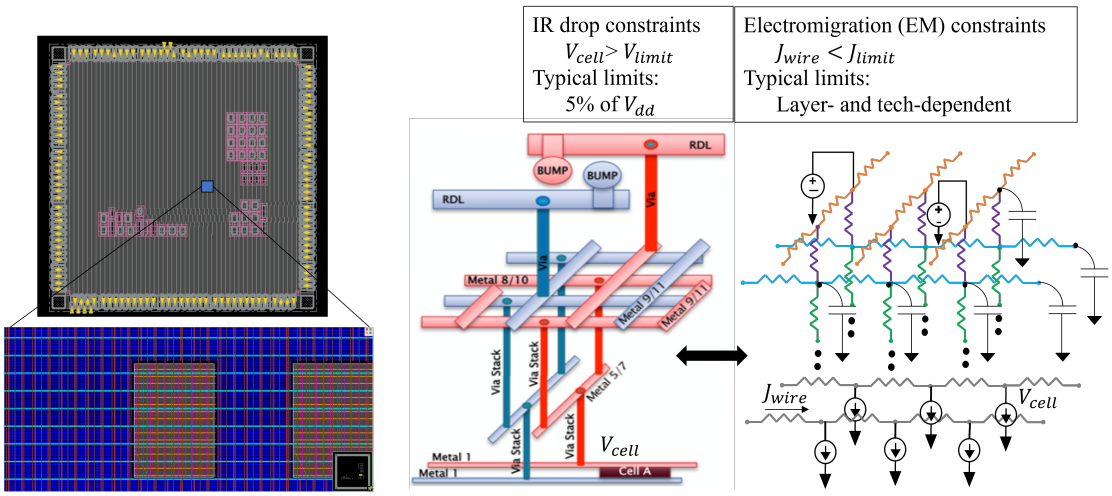


Figure 11.1: PDN analysis: (a) Floorplan of a 14nm RISC-V core; (b) multi-layer power grid that carries voltage from the C4 bumps down to the logic cells; (c) PDN modeled as a RC circuit with voltage sources and current sources.

exposes internal data structures and APIs which are otherwise unavailable in commercial EDA tools. It lowers barriers to entry for non-expert designers with simple push button flows set up as a part of OpenROAD-flow-scripts [125].

PDNSim is an integral part of the power integrity and PDN synthesis tasks of the OpenROAD project and has been used extensively as a part of OpenLane [41] to support the tape-outs done as a part of the Efabless MPW and ChipIgnite [42] programs with the SkyWater 130nm PDK. This chapter describes the software development, capabilities, and usage of PDNSim and also provides several example results across a variety of designs in different technology nodes.

## 11.1 Problem Statement

The goal of PDN analysis is to estimate the voltage drop between the C4 bumps or power I/O pins and every transistor on the chip. Fig. 11.1(a) shows the top-level layout view of the power grid, which consists of multiple metal layers that connect the C4 bump to the logic gates (Fig. 11.1). The PDN can be modeled as an RC circuit as shown in Fig. 11.1 where the C4 bumps are modeled as voltage sources, logic gates that draw current when they switch as current sources, and PDN wires as parasitics resistances

and capacitances. PDN analysis involves finding the voltage at every node ($V_{\text{cell}}$) on the network and finding the current through every resistor ($J_{\text{wire}}$) in the network.

For static analysis, the capacitances are shorted, and performing PDN analysis amounts to solving the resistor network for voltages at each node. As detailed in chapter 2, algorithmically this involves solving a linear system of equations of the form $GV = J$ where $G$ is the conductance matrix, $V$ is a vector of unknown voltages, and $J$ is a vector of current sources attached to the power grid. In the rest of this chapter, we describe the development of the $GV = J$ system of linear equations in PDNSim and highlight its capabilities as a part of the OpenROAD application.

## 11.2 PDNSim Software Framework

PDNSim is modularly developed with interfaces to OpenDB (OpenROAD database), and OpenSTA (OpenROAD static timing analysis engine) [190] as shown in Fig. 11.2. OpenDB provides the information on PDN topology including the metal layers in the PDN, the widths of the PDN wires in each layer, locations and dimensions of vias, per-unit resistances, location of instances, etc. and OpenSTA provides the total power per instance including leakage power, internal power, and switching power. PDNSim uses the information from OpenDB and OpenSTA to create the resistor network with voltage sources and current sources as shown in Fig. 11.1(c). It builds the system of equations $GV = J$ and solves it to estimate voltage at each node and current density in each branch of the circuit.
The inputs to PDNSim are the following:

1. A placed design (.def) that has locations of instances and a synthesized PDN
2. The technology library files including (.lef and .lib)
3. Constraints (.sdc) which specify design frequency
4. C4 bump or power pin location file (optional)

PDNSim generates the following outputs:

1. Worstcase and instance-level IR drop report
2. Worstcase and branch-level current density reports

3. PDN connectivity reports which indicate floating PDN stripes and instances and unconnected instances

4. SPICE netlist for the $GV = J$ system of linear equations representing the PDN.

The key routines within PDNSim are highlighted in Fig. 11.2. PDNSim begins by extracting relevant PDN topology information from OpenDB for the net being analyzed. Next, it discretizes the PDN stripes into nodes and creates a resistance network after performing resistance extraction. It then builds the $GV = J$ system of linear equation using modified nodal analysis (MNA), by creating connections between the nodes, inserting the voltage sources (either from the input file or the default bump locations) into the $G$ matrix, and inserting the current sources into the $J$ vector. It then leverages existing sparse matrix libraries to solve the system of equations and obtain the values in $V$. These values are then processed to obtain the branch current densities. Each of these steps are described below:

**Node discretization** PDNSim queries OpenDB for all wires of the net being analyzed. It iterates through the wires of the power grid and identifies the different metal layers that are a part of it. Next, for each wire, it creates nodes at all of its via locations. Each via is modeled as two nodes between the two metal layers it connects. In addition to via nodes, PDNSim creates nodes at a finer discretization at the bottom-most metal layer of the power grid (typically the M1/M2 power rails). Fig. 11.3 shows a picture of the OpenROAD GUI which highlighted PDN nodes. The nodes in green, yellow, and maroon (stacked one on top of the other) are the nodes at via locations and the nodes in red are M1 rail nodes.

A finer discretization on the bottom-most layer allows accurate insertion of the current sources based on the location of the instances drawing current. For example,

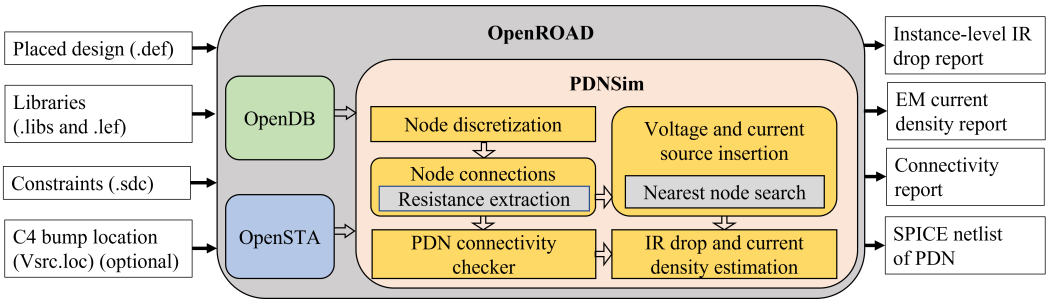


Figure 11.2: PDNSim software framework in the OpenROAD app with interfaces to OpenDB and OpenSTA highlighting the inputs and outputs of PDNSim.

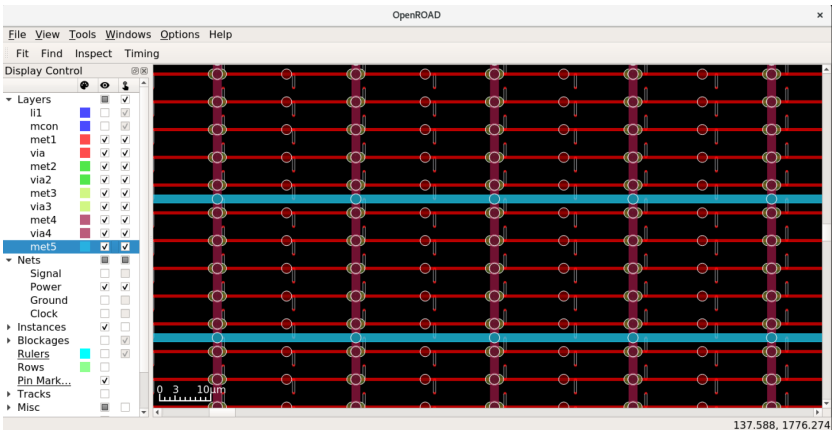


Figure 11.3: PDN nodes visualized in OpenROAD GUI in debug mode highlighting nodes at via locations in the maroon and green stacks and the nodes on the M1 rail at a higher density in red.

Fig. 11.4 shows a region of chameleon design implemented in SkyWater 130nm technology. The picture of the power maps and IR drop maps at two different M1 PDN node pitches. The figures in the top row show the power and IR drop map for a node pitch of 27$\mu$m and the bottom row show the power and IR drop map for a node pitch of 3$\mu$m. The total number of nodes in the power grid determines the size of the $G$ matrix which in turn determines runtimes and accuracy of the tool, where the larger discretizations provide faster but less accurate solutions as shown in the figure where the IR drop map is obtained at a much higher resolution in the bottom row when compared to the top row (white boxes). The high-resolution IR drop map allows for accurate instance-based IR drop annotation.

**Node connections and resistance extraction** Once the nodes have been created, PDNSim builds connections between them which involves creating an adjacency matrix-like representation of the power grid model where the nodes of the circuit are a graph and the edges are resistances. As the connections are being made, PDNSim populates the $G$ matrix with the conductance values between the nodes by performing a simple resistance extraction based on the distance between the nodes, metal layer, width of the power wire, and the obtained per-unit and via resistance from OpenDB. The extracted resistance value is then filled into the appropriate indices of the $G$ as per MNA.

**Voltage and current source insertion** The C4 bumps are modeled as voltage sources as shown in Fig. 11.1(c). These VDD and VSS voltage sources are inserted at C4 bump

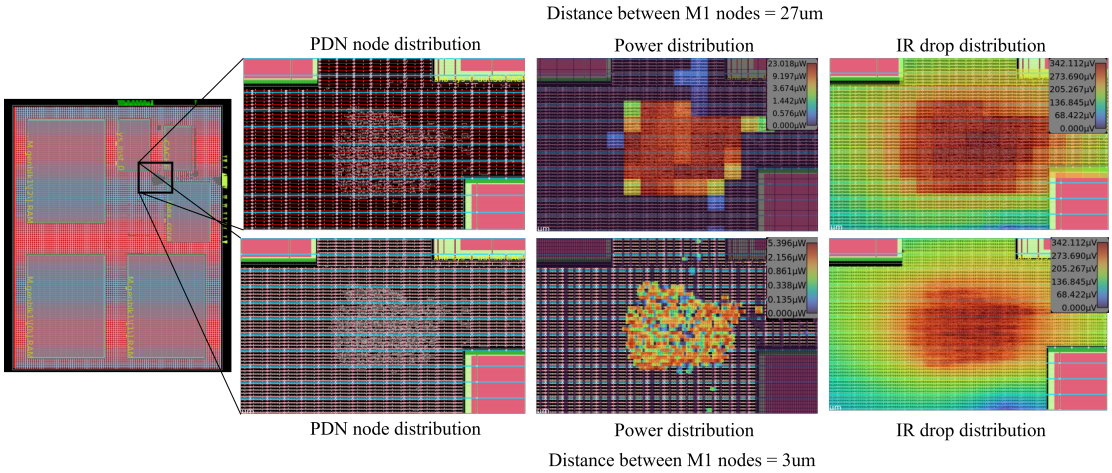


Figure 11.4: Impact of different PDN node discretizations on current source (power map) distributions and IR drop distributions. The top row shows the distributions of current and IR drop for a sparse PDN node pitch (27$\mu$m) and the bottom row is for a dense PDN node pitch (3$\mu$m) in the bottommost metal layer M1.

locations that are either user-defined in a vsrc file or they default to a checkerboard pattern as highlighted in Fig. 5.10. In either case, the C4 bump is modeled with a default RDL layer that connects the bump to the topmost metal layer of the power grid as shown in the figure. The voltage sources are attached to the nodes on the topmost metal layer of the PDN that is closest to the C4 bump locations. Based on MNA, PDNSim adds an additional row and column in the $G$ matrix for each voltage source and an additional element to the $J$ vector with the value of the voltage source.

Instances are modeled as current sources as shown in Fig. 11.1(c). PDNSim runs OpenSTA [190] under the hood to extract the power per instance in the design. It uses the power values as static current estimates for each instance and adds a current source of the same value to the PDN node that is closest to it. One or more instances may share the same node in which case the currents are accumulated. Therefore, each current source represents the sum of power values of several instances in the vicinity of the node to which it is connected. As previously highlighted in Fig 11.4 the power map resolution depends on the node pitch in M1. Further, large pitches cause several instances to share the PDN node and inaccurately share the same annotated IR drop.

**Nearest node search** Both the voltage source and current source insertions require finding the nearest PDN node based on the specified C4 bump locations and instance

location respectively. To find the nearest nodes PDNSim uses a runtime-efficient nearest node routine that leverages a near constant-time bounding box-based node query on a hash map-based data structure that stores node pointers. Therefore, the nearest node routine uses the location of the C4 bump or the instance and creates a bounding box of fixed size around it. Based on the bounding box, the query to the hash map data structure returns all nodes within it very rapidly. For example, Fig. 11.5 highlights the nearest node search for the instance current source insertion on the left and the C4 bump voltage source insertion on the right. For the former, the size bounding box is determined by the node pitch in the standard cell rails and for the latter, the bounding box is determined by the PDN stripe pitch in the topmost metal layer. The bounding box query returns the set of nodes highlighted in the white and red bounding boxes respectively. The routine iterates through these nodes, to return the closest node to either the C4 bump or the instance to attach the voltage source or current source. The chosen node is highlighted by the arrows in the figure.

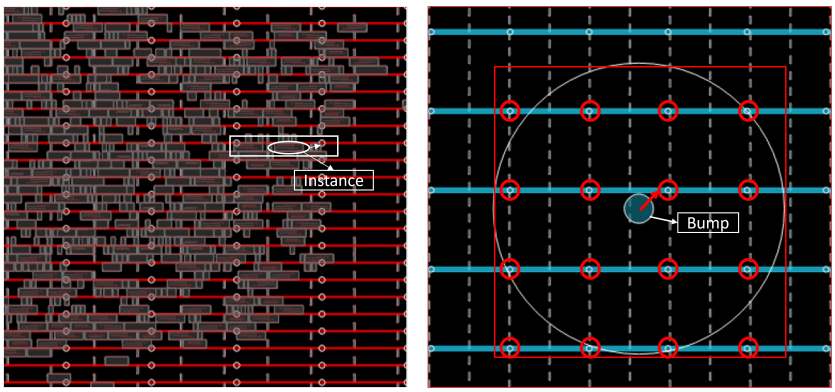


Figure 11.5: Bounding box-based query to find the nearest node to instances (left) and C4 bumps (right).

**IR drop and current density estimation** Once the $GV = J$ sparse system of linear equations is created using MNA, PDNSim leverages the Eigen library [191] to solve it. The vector $V$ contains the simulated voltages at all the nodes which are annotated back into the node class. Using the location of each node and the annotated voltage drop, PDNSim creates layer-wise IR drop heatmaps. Further, it also outputs instance-level IR drop where all instances connected to the same node share identical IR drop values. Therefore, its important to use finer node discretization for higher instance-level IR drop accuracy (See Fig. 11.4). PDNSim reports the worstcase IR drop and its location

which is useful for correct PDN design.

Using the voltage at every node, PDNSim also performs current density estimation where it uses the resistance value of the branch and the annotated voltage values at the two nodes it connects to find the current through the branch. It reports worstcase current and average current across all wires in the power grid for EM checks. These current estimates can be compared with EM limits to check the reliability of the PDN.

**Power grid connectivity checker** In addition to estimating the voltage at each node and the current density in each branch of the PDN, PDNSim also checks the connectivity of the power grid. The connectivity checker is based on a simple graph-like traversal of the adjacency matrix starting from one of the PDN nodes that is connected to a voltage source. The traversal checks if all other nodes are reachable through a recursive search. At the end of the traversal, the presence of unreachable nodes implies that there are floating power grid wires or even instances that are unconnected to the power grid. PDNSim flags and reports the locations of floating power wires and instance names that are not connected to the power grid. The connectivity checker allows for catching power supply-related LVS errors early in the physical design cycle.

The connectivity checker is vital to correct PDN synthesis as it flags several instances of missing power stripes within macro channels, missing connections between the macro power grid and the standard cell power grid, and missing connections between the power grid and the VDD and VSS I/O pins. The PDNSim connectivity checker has been more useful in designs that do not have uniform power grids, such as those with macros. Macros typically come with their own power grids that have been independently designed and therefore they act as blockages to standard cell power grid wires in lower metal layers. However, if the macros are on the same power domain as the rest of the standard cells there must be connections between these two independently designed grids through vias/pins in higher metal layers.

Due to blockages and independently designed power grids there may be several instances where regions that lie between blockages (macro channels) fail to have a power stripe for the standard cells that are placed between macros or there may be several instances where the power grid stripes of the macro in the upper metal layers fail to connect to the power grid of the standard cell due to differences in pitches (misalignment). Fig. 11.6 shows an instance of this issues in chameleon design in SkyWater 130nm technology where the power grid checker correctly identified instances (macros) that receive no voltage and locations where the macro power grid is not connected to the standard cell power grid in the upper metal layer.

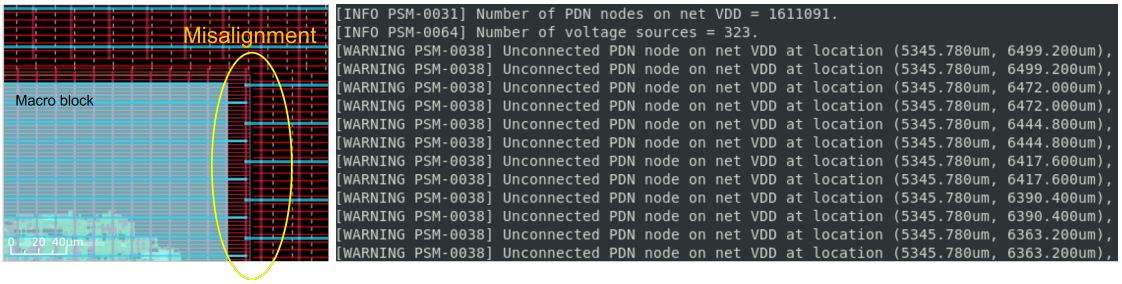


Figure 11.6: An output of PDNSim, flagging misalignment between the higher metal layers of the macro power grid and standard cell power grid. Misalignment on all four sides of the macro results in the macro receiving no voltage.

**PDN SPICE netlist generation** As a part of PDNSim, we have also developed a SPICE netlist writer. The SPICE netlist represents the $GV = J$ sparse system of linear equations and can directly be simulated using SPICE. This netlist is useful to check the correlation of the Eigen sparse matrix solver within PDNSim with a ground-truth SPICE solver. Further, the netlist is also useful for debugging purposes as it includes the different node names their locations, the resistance values between them, and the current and voltage sources connected to them.

## 11.3 PDNSim Results

PDNSim [39] is developed in C++17 and has a dependency on the Eigen library [191] for solving $GV = J$ sparse system of linear equations. In this section, we present the PDN analysis results from PDNSim across several designs in an open-source SkyWater 130nm [126] technology and a commercial FinFET 12nm technology node. Fig. 11.7 shows sample reports for the four different outputs of PDNSim, i.e., an IR drop report, a maximum current report, PDN connectivity report, and a snippet of a SPICE netlist. The reports state the worstcase and average IR drop and EM-related current, and throws a warning if the power grid is unconnected with a set of locations for the unconnected PDN nodes. The generated SPICE netlist encodes the node location within the node names, the corresponding resistance values, and associated current and voltage sources as shown in the figure. The commands to generate these reports using the TCL enablement of OpenROAD is presented in Appendix A.

We also compare PDN analysis results from PDNSim against "golden" commercial tool analysis in terms of the number of nodes, runtimes, IR drop distributions, and

Table 11.1: PDNSim (PSM) results compared with golden commercial tool results.

| Tech | Design | #Nodes | | Total run times (s) | | Worstcase IR drop (mV) | | Worstcase current (A) | |
|---|---|---|---|---|---|---|---|---|---|
| | | PSM | Comm. Tool | PSM | Comm. Tool | PSM | Comm. Tool | PSM | Comm. Tool |
| GF12 | gcd | 11,768 | 7,679 | 3 | 7 | 2.36 | 1.96 | 3.31E-4 | 1.76E-4 |
| | aes | 982,086 | 346,492 | 121 | 50 | 104.12 | 90.38 | 1.06E-2 | 0.73E-2 |
| | jpeg | 420,401 | 214,499 | 88 | 39 | 70.34 | 74.70 | 2.09E-2 | 0.98E-2 |
| | coyote | 5,716,123 | 2,241,225 | 1148 | 774 | 0.98 | 0.66 | 3.71E-4 | 1.68E-4 |
| | swerv_wrapper | 2,818,708 | 1,067,614 | 903 | 305 | 0.56 | 0.37 | 2.51E-4 | 1.13E-4 |
| | bp_single | 15,309,805 | 9,958,335 | 14,651 | 8,006 | 8.67 | 6.14 | 1.87E-2 | 1.04E-2 |
| Sky130HD | gcd | 7,601 | 27,179 | 3 | 17 | 0.89 | 0.43 | 1.47E-5 | 1.05E-5 |
| | aes | 115,700 | 532,994 | 37 | 62 | 1.75 | 1.07 | 6.19E-5 | 2.70E-5 |
| | jpeg | 944,237 | 3,379,623 | 129 | 788 | 0.74 | 0.49 | 2.18E-5 | 1.21E-5 |
| | ibex | 485,271 | 1,714,021 | 74 | 333 | 0.47 | 0.33 | 1.40E-5 | 9.70E-6 |
| | chameleon | 289,801 | 527,693 | 63 | 69 | 0.34 | 0.26 | 2.99E-4 | 1.76E-4 |
| | riscv32i | 11,274 | 31,438 | 2 | 19 | 1.94 | 1.47 | 2.65E-3 | 1.68E-3 |

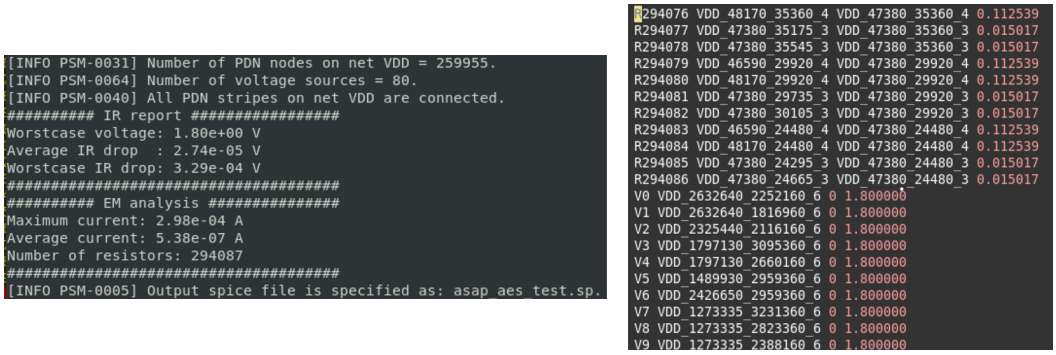


Figure 11.7: Sample IR drop, maximum current reports, and HSPICE netlist snippet.

current density distributions. Table 11.1 shows the results of PDNSim on several designs both with and without macros from OpenROAD-flow-scripts. The table lists the worstcase IR drop, worstcase current for EM checks, and compares the results with results from a commercial tool. It can be seen that PDNSim can perform static PDN analysis on large designs (14M nodes) in times that are comparable to commercial tool runtime. Fig. 11.8 compares the IR drop heatmaps in OpenROAD GUI against IR drop heatmaps from the commercial tools for the jpeg design in the commercial 14nm FinFET technology and open-source SkyWater 130nm. The PDNSim-generated IR drop heatmaps show high fidelity. The differences between commercial tools and PDNSim-reported voltage drops and currents can be attributed to simplified resistance extraction models within PDNSim and differences in OpenSTA-reported power numbers.

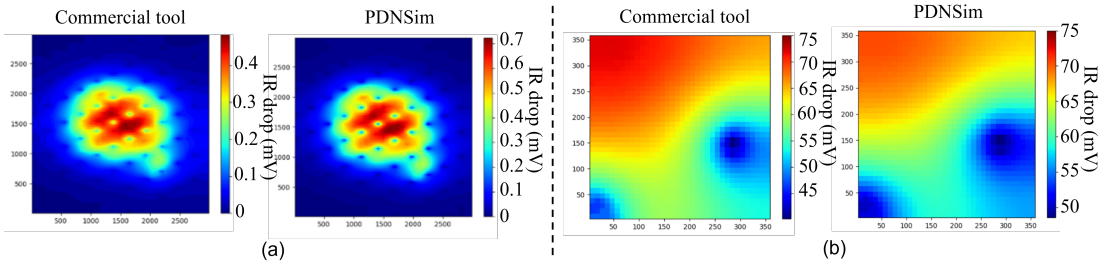


Figure 11.8: Comparison between commercial tool performance and PDNSim on jpeg implemented in: (a) SkyWater 130nm PDK and (b) commercial FinFET 14nm PDK.

More results from PDNSim on different open-source technologies can be found with OpenROAD's nightly regression public gallery [192]. The primary users of PDNSim include the users of OpenLane [41] and OpenROAD-flow-scripts [125] as PDNSim is within the OpenROAD application that these repositories create wrappers around. The users span academia, industry, and government institutions. OpenROAD has had over 14000 clones, over 180 tapeouts in SkyWater 130nm PDK and a mixed signal SOC tapeout in a commercial 12nm FinFET technology with 500K instances and 53 macros.

# References


[1] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs trained by a two time-scale update rule converge to a local Nash equilibrium," in *Adv. NeurIPS*, 2017, pp. 6629–6640.

[2] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of GANs for improved quality, stability, and variation," in *Proc. ICLR*, 2018.

[3] C. Xu, S. K. Kolluri, K. Endo, and K. Banerjee, "Analytical thermal model for self-heating in advanced FinFET devices with implications for design and reliability," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 32, no. 7, pp. 1045–1058, 2013.

[4] E. Pop, S. Sinha, and K. E. Goodson, "Heat generation and transport in nanometer-scale transistors," *Proc. of the IEEE*, vol. 94, no. 8, pp. 1587–1601, 2006.

[5] Y. Ye, F. Liu, M. Chen, and Y. Cao, "Variability analysis under layout pattern-dependent rapid-thermal annealing process," in *Proc. DAC*, 2009, pp. 551–556.

[6] K. Rupp, "42 years of microprocessor trend data," https://www.karlrupp.net/2018/02/42-years-of-microprocessor-trend-data/.

[7] OpenAI, "AI and Compute," https://openai.com/blog/ai-and-compute/.

[8] A. Olofsson, "Silicon compilers – version 2.0," 2018, https://www.ispd.cc/slides/2018/k2.pdf.

[9] V. A. Chhabria, A. B. Kahng, M. Kim, U. Mallappa, S. S. Sapatnekar, and B. Xu, "Template-based PDN synthesis in floorplan and placement using classifier and CNN techniques," in *Proc. ASP-DAC*, 2020, pp. 44–49.

[10] Y.-C. Lu, S. Nath, V. Khandelwal, and S. K. Lim, "RL-Sizer: VLSI Gate Sizing for Timing Optimization using Deep Reinforcement Learning," in *Proc. DAC*, 2021, pp. 733–738.

[11] L. T. Clark, V. Vashishtha, L. Shifren, A. Gujja, S. Sinha, B. Cline, C. Ramamurthy, and G. Yeric, "ASAP7: A 7-nm FinFET predictive process design kit," *Microelectronics Journal*, vol. 53, pp. 105–115, Jul. 2016.

[12] E. Pop, R. Dutton, and K. Goodson, "Thermal analysis of ultra-thin body device scaling [SOI and FinFET devices]," in *Proc. IEDM*, 2003, pp. 36.6.1–36.6.4.

[13] E. Bilotti, O. Fenwick, B. C. Schroeder, M. Baxendale, P. Taroni-Junior, T. Degousée, and Z. Liu, "Organic thermoelectric composites materials," in *Comprehensive Composite Materials II*. Oxford: Elsevier, 2018, pp. 408–430.

[14] K. L. Knutson, S. Cea, M. Giles, P. Keys, P. Davids, C. Weber, L. Shifren, R. Kotlyar, J. Hwang, S. Talukdar, and M. Stettler, "Physical modeling of layout-dependent transistor performance," *ECS Transactions*, vol. 13, no. 1, p. 63, oct 2008.

[15] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, no. 6, pp. 84—90, May 2017.

[16] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. CVPR*, 2009, pp. 248–255.

[17] A. B. Kahng, "Scaling: More than Moore's law," *IEEE Des. Test Comput.*, vol. 27, no. 3, pp. 86–87, May 2010.

[18] D. Mallick, "An AI chip with unprecedented performance to do the unimaginable," https://www.cerebras.net/blog/an-ai-chip-with-unprecedented-performance-to-do-the-unimaginable/.

[19] V. A. Chhabria, V. Ahuja, A. Prabhu, N. Patil, P. Jain, and S. S. Sapatnekar, "Encoder-decoder networks for analyzing thermal and power delivery networks," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 28, no. 1, dec 2022.

[20] ——, "Thermal and IR drop analysis using convolutional encoder-decoder networks," in *Proc. ASP-DAC*, 2021, pp. 690–696.

[21] V. A. Chhabria, Y. Zhang, H. Ren, B. Keller, B. Khailany, and S. S. Sapatnekar, "MAVIREC: ML-aided vectored IR-drop estimation and classification," *arXiv:2012.10597 [cs.ar]*, 2020.

[22] ——, "MAVIREC: ML-aided vectored IR-drop estimation and classification," in *Proc. DATE*, 2021, pp. 1825–1828.

[23] V. A. Chhabria and S. S. Sapatnekar, "OpeNPDN," https://github.com/The-OpenROAD-Project/OpeNPDN.

[24] ——, "OpeNPDN: A neural-network-based framework for power delivery network synthesis," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 41, no. 10, pp. 3515–3528, 2022.

[25] V. A. Chhabria, K. Kunal, M. Zabihi, and S. S. Sapatnekar, "BeGAN: Power grid benchmark generation using a process-portable gan-based methodology," in *Proc. ICCAD*, 2021.

[26] ——, "BeGAN PDN Benchmarks," https://github.com/UMN-EDA/BeGAN-benchmarks.

[27] S. R. Nassif, "Power grid analysis benchmarks," in *Proc. ASP-DAC*, 2008, pp. 376–381.

[28] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Adv. NeurIPS*, 2014.

[29] V. A. Chhabria, B. Keller, Y. Zhang, S. Vollala, S. Pratty, H. Ren, and B. Khailany, "XT-PRAGGMA: Crosstalk pessimism reduction achieved with gpu gate-level simulations and machine learning," in *Proc. MLCAD*, 2022, pp. 63–69.

[30] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "PyTorch: An imperative style, high-performance deep learning library," in *Adv. NeurIPS*, 2019.

[31] M. Abadi, B. P, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray,

B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: A system for large-scale machine learning," in *Proc. OSDI*, 2016, pp. 265–283.

[32] Hugging Face, "Hugging Face," https://huggingface.co/.

[33] A. B. Kahng, "Looking into the mirror of open source," in *Proc. ICCAD*, 2019, pp. 1–8.

[34] T. Ajayi, V. A. Chhabria, M. Fogaça, S. Hashemi, A. Hosny, A. B. Kahng, M. Kim, J. Lee, U. Mallappa, M. Neseem, G. Pradipta, S. Reda, M. Saligane, S. S. Sapatnekar, C. Sechen, M. Shalan, W. Swartz, L. Wang, Z. Wang, M. Woo, and B. Xu, "Toward an open-source digital flow: First learnings from the OpenROAD project," in *Proc. DAC*, 2019.

[35] V. A. Chhabria and S. S. Sapatnekar, "TherMOS," https://github.com/UMN-EDA/TherMOS.

[36] ——, "Impact of self-heating on performance and reliability in finfet and gaafet designs," in *Proc. ISQED*, 2019, pp. 235–240.

[37] ——, "RTA-Simulator," https://github.com/UMN-EDA/RTA-Simulator.

[38] ——, "Analysis of pattern-dependent rapid thermal annealing effects on SRAM design," in *Proc. ISQED*, 2023, under review.

[39] ——, "PDNSim," 2021, https://github.com/The-OpenROAD-Project/OpenROAD/tree/master/src/psm.

[40] Z. Xie, H. Ren, B. Khailany, Y. Sheng, S. Santosh, J. Hu, and Y. Chen, "PowerNet: Transferable dynamic IR drop estimation via maximum convolutional neural network," in *Proc. ASP-DAC*, 2020, pp. 13–18.

[41] M. Shalan and T. Edwards, "Building OpenLANE: A 130nm OpenROAD-based tapeout-proven flow," in *Proc. ICCAD*, 2020.

[42] Efabless, "efabless.com," https://efabless.com/.

[43] Y. Zhan, S. V. Kumar, and S. S. Sapatnekar, "Thermally-aware design," *Found. Trends Electron. Des. Automat.*, vol. 2, no. 3, pp. 255–370, March 2008.

[44] Y. Zhong and M. D. F. Wong, "Fast algorithms for IR drop analysis in large power grid," in *Proc. ICCAD*, 2005, pp. 351–357.

[45] J. N. Kozhaya, S. R. Nassif, and F. N. Najm, "A multigrid-like technique for power grid analysis," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 21, no. 10, pp. 1148–1160, Oct. 2002.

[46] K. Zhang, A. Guliani, S. Ogrenci-Memik, G. Memik, K. Yoshii, R. Sankaran, and P. Beckman, "Machine learning-based temperature prediction for runtime thermal management across system components," *IEEE Trans. Parallel Distrib. Syst.*, vol. 29, no. 2, pp. 405–419, Feb. 2018.

[47] D. Juan, H. Zhou, D. Marculescu, and X. Li, "A learning-based autoregressive model for fast transient thermal analysis of chip-multiprocessors," in *Proc. ASP-DAC*, 2012, pp. 597–602.

[48] S. Sadiqbatcha, H. Zhao, H. Amrouch, J. Henkel, and S. X. . Tan, "Hot spot identification and system parameterized thermal modeling for multi-core processors through infrared thermal imaging," in *Proc. DATE*, 2019, pp. 48–53.

[49] S.-Y. Lin, Y.-C. Fang, Y.-C. Li, Y. Liu, T. Yang, S.-C. Lin, C.-M. J. Li, and E. J.-W. Fang, "IR drop prediction of ECO-revised circuits using machine learning," in *Proc. VTS*, 2018.

[50] C.-T. Ho and A. B. Kahng, "IncPIRD: Fast learning-based prediction of incremental IR drop," in *Proc. ICCAD*, 2019.

[51] W. Jin, S. Sadiqbatcha, J. Zhang, and S. X.-D. Tan, "Full-chip thermal map estimation for commercial multi-core CPUs with generative adversarial learning," in *Proc. ICCAD*, 2020.

[52] J. Wen, S. Pan, N. Chang, W. Chuang, W. Xia, D. Zhu, A. Kumar, E. Yang, K. Srinivasan, and Y. Li, "DNN-based fast static on-chip thermal solver," in *Proc. IEEE Symp. Semicond. Therm. Meas., Model. Manage.*, 2020, pp. 65–75.

[53] S. Sadiqbatcha, J. Zhang, H. Amrouch, and S. X.-D. Tan, "Real-time full-chip thermal tracking: A post-silicon, machine learning perspective," *IEEE Trans. Comput.*, vol. 71, no. 6, pp. 1411–1424, 2022.

[54] E. Chiprout, "Fast flip-chip power grid analysis via locality and grid shells," in *Proc. ICCAD*, 2004, pp. 485–488.

[55] W. Jin, S. Sadiqbatcha, Z. Sun, H. Zhou, and S. X.-D. Tan, "EM-GAN: Data-driven fast stress analysis for multi-segment interconnects," in *Proc. ICCD*, 2020, pp. 296–303.

[56] E. Shelhamer, J. Long, and T. Darrell, "Fully convolutional networks for semantic segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 4, pp. 640–651, Apr. 2017.

[57] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervention*, 2015, pp. 234–241.

[58] X.-J. Mao, C. Shen, and Y.-B. Yang, "Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections," in *Adv. NeurIPS*, 2016.

[59] V. Badrinarayanan, A. Kendall, and R. Cipolla, "SegNet: A deep convolutional encoder-decoder architecture for image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 12, pp. 2481–2495, Dec. 2017.

[60] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735—1780, Nov. 1997.

[61] J. Singh and S. S. Sapatnekar, "Partition-based algorithm for power grid design using locality," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 25, no. 4, pp. 664–677, Apr. 2006.

[62] B. W. Amick, C. R. Gauthier, and D. Liu, "Macro-modeling concepts for the chip electrical interface," in *Proc. DAC*, 2002, pp. 391–394.

[63] F. Yazdani, *Foundations of Heterogeneous Integration: An Industry-Based, 2.5D/3D Pathfinding and Co-Design Approach.* Boston, MA, USA: Springer, 2018.

[64] J. Mänttäri, S. Broomé, J. Folkesson, and H. Kjellström, "Interpreting video features: A comparison of 3D convolutional networks and convolutional LSTM networks," in *Proc. ACCV*, 2021, pp. 411–426.

[65] J. You and J. Korhonen, "Deep neural networks for no-reference video quality assessment," in *Proc. Intl. Conf. on Image. Process.*, 2019, pp. 2349–2353.

[66] I. A. Blech, "Electromigration in thin aluminum films on titanium nitride," *J. Appl. Phys.*, vol. 47, no. 4, pp. 1203–1208, Apr. 1976.

[67] J. R. Black, "Electromigration failure modes in aluminum metallization for semiconductor devices," *Proc. IEEE*, vol. 57, no. 9, pp. 1587–1594, Sep. 1969.

[68] M. A. Korhonen, P. Borgesen, K. N. Tu, and C. Y. Li, "Stress evolution due to electromigration in confined metal lines," *J. Appl. Phys.*, vol. 73, no. 8, pp. 3790–3799, Aug. 1993.

[69] S. X.-D. Tan, M. Tahoori, T. Kim, S. Wang, Z. Sun, and S. Kiamehr, *VLSI systems long-term reliability – Modeling, simulation and optimization.* Boston, MA: Springer, 2019.

[70] M. A. A. Shohel, V. A. Chhabria, and S. S. Sapatnekar, "A new, computationally efficient "Blech criterion" for immortality in general interconnects," in *Proc. DAC*, 2021.

[71] M. A. A. Shohel, V. A. Chhabria, N. Evmorfopoulos, and S. S. Sapatnekar, "Analytical modeling of transient electromigration stress based on boundary reflections," in *Proc. ICCAD*, 2021.

[72] H. He and E. A. Garcia, "Learning from imbalanced data," *IEEE Trans. Knowl. Data Eng.*, vol. 21, no. 9, pp. 1263–1284, Sep. 2009.

[73] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "Smote: Synthetic minority over-sampling technique," *J. Artif. Intell. Res.*, vol. 16, no. 1, pp. 321—-357, Jun. 2002.

[74] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, and S. Belongie, "Class-balanced loss based on effective number of samples," in *Proc. CVPR*, June 2019, pp. 9260–9269.

[75] V. Dumoulin and F. Visin, "A guide to convolution arithmetic for deep learning," *arXiv:1603.07285 [stat.ML]*, 2016.

[76] W. Luo, Y. Li, R. Urtasun, and R. Zemel, "Understanding the effective receptive field in deep convolutional neural networks," in *Adv. NeurIPS*, 2016.

[77] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-C. Woo, "Convolutional LSTM network: A machine learning approach for precipitation nowcasting," in *Adv. NeurIPS*, 2015.

[78] Ansys, "Icepak," 2018, https://www.ansys.com/products/electronics/ansys-icepak.

[79] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.

[80] N. Ahmed, M. Tehranipoor, and V. Jayaram, "Transition delay fault test pattern generation considering supply voltage noise in a SOC design," in *Proc. DAC*, 2007.

[81] Y.-C. Fang, H.-Y. Lin, M.-Y. Su, C.-M. Li, and E. J.-W. Fang, "Machine-learning-based dynamic IR drop prediction for ECO," in *Proc. ICCAD*, 2018.

[82] K. Acharya and N. Dhanwada, "Learning-based approach for early power grid analysis in high performance microprocessor designs," Presentation at DAC (User Track), 2020.

[83] P. Covington, J. Adams, and E. Sargin, "Deep neural networks for YouTube recommendations," in *Proc. ACM Conf. Recomm. Sys.*, 2016.

[84] X.-D. Tan, C. R. Shi, D. Lungeanu, J.-C. Lee, and L.-P. Yuan, "Reliability-constrained area optimization of VLSI power/ground networks via sequence of linear programmings," in *Proc. DAC*, 1999, pp. 78–83.

[85] X. Wu, X. Hong, Y. Cai, Z. Luo, C.-K. Cheng, J. Gu, and W. Dai, "Area minimization of power distribution network using efficient nonlinear programming techniques," *Proc. ICCAD*, vol. 23, no. 7, pp. 1086–1094, Jul. 2004.

[86] H. Su, J. Hu, S. S. Sapatnekar, and S. R. Nassif, "Congestion-driven codesign of power and signal networks," in *Proc. DAC*, 2002, pp. 64–69.

[87] J. N. Kozhaya, S. R. Nassif, and F. N. Najm, "Multigrid-like technique for power grid analysis," in *Proc. ICCAD*, Nov 2001, pp. 480–487.

[88] M. Zhao, R. V. Panda, S. S. Sapatnekar, and D. Blaauw, "Hierarchical analysis of power distribution networks," in *Proc. DAC*, 2000, pp. 150–155.

[89] P. Li, "Power grid simulation via efficient sampling-based sensitivity analysis and hierarchical symbolic relaxation," in *Proc. DAC*, 2005, pp. 664–669.

[90] H. Zhuang, W. Yu, S. Weng, I. Kang, J. Lin, X. Zhang, R. Coutts, and C. Cheng, "Simulation algorithms with exponential integration for time-domain analysis of large-scale power delivery networks," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 35, no. 10, pp. 1681–1694, Oct. 2016.

[91] S. Dey, S. Nandi, and G. Trivedi, "PowerPlanningDL: Reliability-aware framework for on-chip power grid design using deep learning," in *Proc. DATE*, 2020, pp. 1520–1525.

[92] W. Chang, C. Lin, S. Mu, L. Chen, C. Tsai, Y. Chiu, and M. C. . Chao, "Generating routing-driven power distribution networks with machine-learning technique," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 36, no. 8, pp. 1237–1250, Aug 2017.

[93] Cadence, "Cadence Innovus Implementation System," http://www.cadence.com.

[94] R. Jakushokas and E. G. Friedman, "Methodology for multi-layer interdigitated power and ground network design," in *Proc. ISCAS*, 2010, pp. 3208–3211.

[95] R. Bhooshan, "Novel and efficient IR-drop models for designing power distribution network for sub-100nm integrated circuits," in *Proc. ISQED*, 2007, pp. 287–292.

[96] S. Müller and L. Schüler, "GeoStat-Framework/GSTools: v1.3.0-rc1 'Pure Pink'," Jan. 2021.

[97] H. Hsu, M. Chen, H. Chen, H. Li, and S. Chen, "On effective flip-chip routing via pseudo single redistribution layer," in *Proc. DATE*, 2012, pp. 1597–1602.

[98] J. A. Snyman, *Practical Mathematical Optimization: An Introduction to Basic Optimization Theory and Classical and New Gradient-Based Algorithms.* Boston, MA, USA: Springer, 2005.

[99] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, Oct 2010.

[100] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, "Learning and transferring mid-level image representations using convolutional neural networks," in *Proc. CVPR*, 2014, pp. 1717–1724.

[101] M. Patil, R. T. Susarla, and S. B. Kota, "A multi-perspective approach to IC power grid development for 7nm based designs," in *"ACM/IEEE Design Automation Conference (DAC), Designer Track"*, 2019, (slides available at dac.com).

[102] A. B. Kahng, S. Kang, S. Kim, K. Samadi, and B. Xu, "Power delivery pathfinding for emerging die-to-wafer integration technology," in *Proc. DATE*, 2019, pp. 842–847.

[103] Y. Wei, C. Sze, N. Viswanathan, Z. Li, C. J. Alpert, L. Reddy, A. D. Huber, G. E. Tellez, D. Keller, and S. S. Sapatnekar, "GLARE: Global and local wiring aware routability evaluation," in *Proc. DAC*, 2012, pp. 768–773.

[104] D. Petrisko, F. Gilani, M. Wyse, D. C. Jung, S. Davidson, P. Gao, C. Zhao, Z. Azad, S. Canakci, B. Veluri, T. Guarino, A. Joshi, M. Oskin, and M. B. Taylor, "BlackParrot: An agile open-source RISC-V multicore for accelerator SoCs," *IEEE Micro*, vol. 40, no. 4, pp. 93–102, 2020.

[105] "SweRV," https://github.com/westerndigitalcorporation/swerv_eh1.

[106] Cadence, "Cadence Voltus IC Power Integrity Solution," http://www.cadence.com.

[107] F. Brglez and H. Fujiwara, "A neutral netlist of 10 combinational benchmark circuits and a targeted translator in FORTRAN," in *Proc. ISCAS*, 1985.

[108] F. Brglez, D. Bryan, and K. Kozminski, "Combinational profiles of sequential benchmark circuits," in *Proc. ISCAS*, 1989, pp. 1929–1934.

[109] F. Corno, M. Reorda, and G. Squillero, "RT-level ITC'99 benchmarks and first ATPG results," *IEEE Des. Test*, vol. 17, no. 3, pp. 44–53, 2000.

[110] S. Yang, "Logic Synthesis and Optimization Benchmarks User Guide: Version 3.0," MCNC Technical Report, Tech. Rep., 1991.

[111] C. Albrecht, "IWLS 2005 benchmarks," 2005, iwls.org/iwls2005/benchmark_presentation.pdf.

[112] G.-J. Nam, C. J. Alpert, P. Villarrubia, B. Winter, and M. Yildiz, "The ISPD2005 placement contest and benchmark suite," in *Proc. ISPD*, 2005, pp. 216–220.

[113] G.-J. Nam, C. Sze, and M. Yildiz, "The ISPD global routing benchmark suite," in *Proc. ISPD*, 2008, pp. 156–159.

[114] W.-H. Liu, S. Mantik, W.-K. Chow, Y. Ding, A. Farshidi, and G. Posser, "ISPD 2019 initial detailed routing contest and benchmark with advanced routing rules," in *Proc. ISPD*, 2019, pp. 147–151.

[115] "IBM power grid benchmarks," web.ece.ucsb.edu/~lip/PGBenchmarks/ibmpgbench.html.

[116] Z. Li, R. Balasubramanian, F. Liu, and S. Nassif, "2012 TAU power grid simulation contest: Benchmark suite and results," in *Proc. ICCAD*, 2012, pp. 478–481.

[117] R. Aitken, "Time to retire our benchmarks," *IEEE Des. Test*, p. 88, 2010.

[118] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv:1511.06434 [cs.LG]*, 2016.

[119] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. CVPR*, 2019, pp. 4396–4405.

[120] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of stylegan," in *Proc. CVPR*, 2020, pp. 8107–8116.

[121] A. Noguchi and T. Harada, "Image generation from small datasets via batch statistics adaptation," in *Proc. ICCV*, 2019, pp. 2750–2758.

[122] M. Zhao, Y. Cong, and L. Carin, "On leveraging pretrained GANs for limited-data generation," in *Proc. ICML*, 2020, pp. 11 340–11 351.

[123] P.-W. Luo, C. Zhang, Y.-T. Chang, L.-C. Cheng, H.-H. Lee, B.-L. Sheu, Y.-S. Su, D.-M. Kwai, and Y. Shi, "Benchmarking for research in power delivery networks of three-dimensional integrated circuits," in *Proc. ISPD*, 2013, pp. 17–24.

[124] Y. Wang, C. Wu, L. Herranz, J. van de Weijer, A. Gonzalez-Garcia, and B. Raducanu, "Transferring GANs: generating images from limited data," in *Proc. ECCV*, 2018, pp. 220–236.

[125] The OpenROAD Project, "OpenROAD-flow-scripts," https://github.com/The-OpenROAD-Project/OpenROAD-flow-scripts.

[126] "SkyWater 130nm PDK," github.com/google/skywater-pdk.

[127] A. Brock, J. Donahue, and K. Simonyan, "Large scale GAN training for high fidelity natural image synthesis," in *Proc. ICLR*, 2019.

[128] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, "Spectral normalization for generative adversarial networks," in *Proc. ICLR*, 2018.

[129] P. Bojanowski, A. Joulin, D. Lopez-Pas, and A. Szlam, "Optimizing the latent space of generative networks," in *Proc. ICML*, 2018, pp. 600–609.

[130] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Proc. ECCV*, 2016, pp. 694–711.

[131] K. Hu, A. N. Nowroz, S. Reda, and F. Koushanfar, "High-sensitivity hardware trojan detection using multimodal characterization," in *Proc. DATE*, 2013, pp. 1271–1276.

[132] Z. C. Lipton and S. Tripathi, "Precise recovery of latent vectors from generative adversarial networks," in *ICLR Workshop*, 2017.

[133] Synopsys, "PrimeTime SI: Crosstalk Delay and Noise," 2022, https://www.synopsys.com/implementation-and-signoff/signoff/primetime.html.

[134] Cadence, "Tempus Timing Signoff Solution," 2022, https://www.cadence.com/en_US/home/tools/digital-design-and-signoff/silicon-signoff/tempus-timing-signoff-solution.html.

[135] B. Franzini, C. Forzan, D. Pandini, P. Scandolara, and A. Dal Fabbro, "Crosstalk aware static timing analysis: A two step approach," in *Proc. ISQED*, 2000, pp. 499–503.

[136] H. Fatemi and P. Tehrani, "Crosstalk timing windows overlap in statistical static timing analysis," in *Proc. ISQED*, 2013, pp. 245–251.

[137] M. Becer, V. Zolotov, R. Panda, A. Grinshpon, I. Algol, R. Levy, and C. Oh, "Pessimism reduction in crosstalk noise aware STA," in *Proc. ICCAD*, 2005, pp. 954–961.

[138] A. Glebov, S. Gavrilov, R. Soloviev, V. Zolotov, M. Becer, C. Oh, and R. Panda, "Delay noise pessimism reduction by logic correlations," in *Proc. ICCAD*, 2004, pp. 160–167.

[139] A. Glebov, S. Gavrilov, D. Blaauw, S. Sirichotiyakul, C. Oh, and V. Zolotov, "False-noise analysis using logic implications," in *Proc. ICCAD*, 2001, pp. 515–521.

[140] A. Krstic, Y.-M. Jiang, and K.-T. Cheng, "Pattern generation for delay testing and dynamic timing analysis considering power-supply noise effects," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 20, no. 3, pp. 416–425, 2001.

[141] B. Paul and K. Roy, "Testing cross-talk induced delay faults in static CMOS circuit through dynamic timing analysis," in *Proc. of Int. Test Conf.*, 2002, pp. 384–390.

[142] Y.-M. Jiang, A. Krstic, and K.-T. Cheng, "Dynamic timing analysis considering power supply noise effects," in *Proc. ISQED*, 2000, pp. 137–143.

[143] Y. Chen, A. B. Kahng, B. Liu, and W. Wang, "Crosstalk-aware signal probability-based dynamic statistical timing analysis," in *Proc. ISQED*, 2015, pp. 424–429.

[144] A. B. Kahng, M. Luo, and S. Nath, "SI for free: machine learning of interconnect coupling delay and transition effects," in *Proc. SLIP*, 2015, pp. 1–8.

[145] R. Liang, Z. Xie, J. Jung, V. Chauha, Y. Chen, J. Hu, H. Xiang, and G.-J. Nam, "Routing-free crosstalk prediction," in *Proc. ICCAD*, 2020.

[146] H. Ren, B. Khailany, M. Fojtik, and Y. Zhang, "Machine learning and algorithms: Let us team up for EDA," *IEEE Des. Test*, 2022.

[147] Y. Zhang, H. Ren, A. Sridharan, and B. Khailany, "GATSPI: GPU accelerated gate-level simulation for power improvement," in *Proc. DAC*, 2022.

[148] M. Wang, D. Zheng, Z. Ye, Q. Gan, M. Li, X. Song, J. Zhou, C. Ma, L. Yu, Y. Gai, T. Xiao, T. He, G. Karypis, J. Li, and Z. Zhang, "Deep Graph Library: A Graph-Centric, Highly-Performant Package for Graph Neural Networks," *arXiv:1909.01315 [cs.ar]*, 2020.

[149] S. Holst, M. E. Imhof, and H.-J. Wunderlich, "High-throughput logic timing simulation on GPGPUs," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 20, no. 3, Jun. 2015.

[150] Y. Zhang, H. Ren, and B. Khailany, "Opportunities for RTL and gate level simulation using GPUs (invited talk)," in *Proc. ICCAD*, 2020, pp. 1–5.

[151] Z. Chen, "Crosstalk superposition of multiple aggressors in electronic package system pre-pd signal integrity simulations," in *Proc. Elec. Perf. of Electron. Packag.*, 2006, pp. 115–118.

[152] J. P. Fishburn and A. E. Dunlop, "TILOS: A posynomial programming approach to transistor sizing," in *Proc. ICCAD*, 1985, pp. 326–328.

[153] J. Hu, A. B. Kahng, S. Kang, M.-C. Kim, and I. L. Markov, "Sensitivity-Guided Metaheuristics for Accurate Discrete Gate Sizing," in *Proc. ICCAD*, 2012, pp. 233–239.

[154] K. Kasamsetty, M. Ketkar, and S. Sapatnekar, "A New Class of Convex Functions for Delay Modeling and its Application to the Transistor Sizing Problem," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 19, no. 7, pp. 779–788, 2000.

[155] A. Sharma, D. Chinnery, T. Reimann, S. Bhardwaj, and C. Chu, "Fast Lagrangian Relaxation-Based Multithreaded Gate Sizing Using Simple Timing Calibrations," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 39, no. 7, pp. 1456–1469, 2020.

[156] C.-P. Chen, C. Chu, and D. Wong, "Fast and Exact Simultaneous Gate and Wire Sizing by Lagrangian Relaxation," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 18, no. 7, pp. 1014–1025, 1999.

[157] A. Mirhoseini, A. Goldie, M. Yazgan, J. Jiang, E. Songhori, W. Shen, Y.-J. Lee, E. Johnson, O. Pathak, A. Nazi, J. Pak, A. Tong, K. Srinivasa, W. Hang, E. Tuncer, Q. Le, J. Laudon, R. Ho, R. Carpenter, and J. Dean, "A Graph Placement Methodology for Fast Chip Design," *Nature*, vol. 594, pp. 207–212, Jun. 2021.

[158] H. Ren, S. Godil, B. Khailany, R. Kirby, H. Liao, S. Nath, J. Raiman, and R. Roy, "Optimizing VLSI implementation with reinforcement learning," in *Proc. ICCAD*, 2021.

[159] H. Wang, K. Wang, J. Yang, L. Shen, N. Sun, H.-S. Lee, and S. Han, "GCN-RL circuit designer: Transferable transistor sizing with graph neural networks and reinforcement learning," in *Proc. DAC*, 2020.

[160] X. Zhou, J. Ye, C.-W. Pui, K. Shao, G. Zhang, B. Wang, J. Hao, G. Chen, and P. A. Heng, "Heterogeneous graph neural network-based imitation learning for gate sizing acceleration," in *Proc. ICCAD*, 2022.

[161] S. Huang, A. Abdolmaleki, G. Vezzani, P. Brakel, D. J. Mankowitz, M. Neunert, S. Bohez, Y. Tassa, N. Heess, M. Riedmiller, and R. Hadsell, "A constrained multi-objective reinforcement learning framework," in *Proc. Conf. on Robot Learning*, 2022, pp. 883–893.

[162] A. Irpan, "Deep reinforcement learning doesn't work yet," 2018, https://www.alexirpan.com/2018/02/14/rl-hard.html.

[163] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing Atari with deep reinforcement learning," in *Adv. NeurIPS*, 2013.

[164] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. van den Berg, I. Titov, and M. Welling, "Modeling Relational Data with Graph Convolutional Networks," in *Proc. Eur. Semantic Web Conf.*, 2018, pp. 593–607.

[165] N. Viswanathan and C. C.-N. Chu, "FastPlace: Efficient analytical placement using cell shifting, iterative local refinement and a hybrid net model," in *Proc. ISPD*, 2004, p. 26–33.

[166] "NanGate 45nm FreePDK and cell library," 2022, https://si2.org/open-cell-library.

[167] Y. Huang, M. Chiang, S. Wang, and J. G. Fossum, "GAAFET versus pragmatic FinFET at the 5nm Si-based CMOS technology node," *IEEE J. Electron Dev. Soc.*, vol. 5, no. 3, pp. 164–169, 2017.

[168] S. Ramey, Y. Lu, I. Meric, S. Mudanai, S. Novak, C. P. Prasad, and J. Hicks, "Aging model challenges in deeply scaled tri-gate technologies," *Proc. IIRW*, pp. 56–62, 2015.

[169] S. E. Liu, J. S. Wang, Y. R. Lu, D. S. Huang, C. F. Huang, W. H. Hsieh, J. H. Lee, Y. S. Tsai, J. R. Shih, Y. H. Lee, and K. Wu, "Self-heating effect in FinFETs and its impact on devices reliability characterization," in *Proc. IRPS*, 2014, pp. 4A.4.1–4A.4.4.

[170] D. Rossi, V. Tenentes, S. Yang, S. Khursheed, and B. M. Al-Hashimi, "Aging benefits in nanometer CMOS designs," *IEEE Trans. Circuits Syst. II*, vol. 64, no. 3, pp. 324–328, 2017.

[171] S. Mishra, H. Y. Wong, R. Tiwari, A. Chaudhary, R. Rao, V. Moroz, and S. Mahapatra, "TCAD-based predictive NBTI framework for sub-20-nm node device design considerations," *IEEE Trans. Electron Devices*, vol. 63, no. 12, pp. 4624–4631, 2016.

[172] Z. Yu, J. Zhang, R. Wang, S. Guo, C. Liu, and R. Huang, "New insights into the hot carrier degradation (HCD) in FinFET: New observations, unified compact model, and impacts on circuit reliability," in *Proc. IEDM*, Dec 2017, pp. 7.2.1–7.2.4.

[173] I. Messaris, N. Fasarakis, T. A. Karatsori, A. Tsormpatzoglou, G. Ghibaudo, and C. A. Dimitriadis, "Hot carrier degradation modeling of short-channel n-FinFETs," in *Proc. DRC*, 2015, pp. 183–184.

[174] K.-D. Lee, "Electromigration critical length effect and early failures in Cu/oxide and Cu/low k interconnects," Ph.D. dissertation, Univ. Texas Austin, Austin, TX, 2003.

[175] R. Wang, R. Huang, D. Kim, Y. He, Z. Wang, G. Jia, D. Park, and Y. Wang, "New observations on the hot carrier and NBTI reliability of silicon nanowire transistors," in *Proc. IEDM*, 2007, pp. 821–824.

[176] M. Si, S. Shin, N. J. Conrad, J. Gu, J. Zhang, M. A. Alam, and P. D. Ye, "Characterization and reliability of III-V gate-all-around MOSFETs," in *Proc. IRPS*, 2015, pp. 4A.1.1–4A.1.6.

[177] S. V. Kumar, C. H. Kim, and S. S. Sapatnekar, "NBTI-aware synthesis of digital circuits," in *Proc. DAC*, 2007, pp. 370–375.

[178] H. Mertens, R. Ritzenthaler, A. Hikavyy, M. S. Kim, Z. Tao, K. Wostyn, S. A. Chew, A. D. Keersgieter, G. Mannaert, E. Rosseel, T. Schram, K. Devriendt, D. Tsvetanova, H. Dekkers, S. Demuynck, A. Chasin, E. V. Besien, A. Dangol, S. Godny, B. Douhard, N. Bosman, O. Richard, J. Geypen, H. Bender, K. Barla, D. Mocuta, N. Horiguchi, and A. V. Y. Thean, "Gate-all-around MOSFETs based on vertically stacked horizontal Si nanowires in a replacement metal gate process on bulk Si substrates," in *Proc. VLSIT*, 2016, pp. 1–2.

[179] M. Rabus, A. Fiory, N. Ravindra, P. Frisella, A. Agarwal, T. Sorsch, J. Miner, E. Ferry, F. Klemens, R. Cirelli, and W. Mansfield, "Rapid thermal processing of silicon wafers with emissivity patterns," *J. Electron. Mater.*, vol. 35, pp. 877–891, 12 2006.

[180] E. H. Granneman, H. Terhorst, A. Falepin, E. Rosseel, K. Verheyden, K. Vanormelingen, H. Bourdon, A. Halimaoui, and K. Funk, "3D pattern effects in RTA radiative vs. conductive heating," *ECS Transactions*, vol. 3, no. 2, p. 85, oct 2006.

[181] Y. Wei, J. Hu, F. Liu, and S. S. Sapatnekar, "Physical design techniques for optimizing RTA-induced variations," in *Proc. ASP-DAC*, 2010, pp. 745–750.

[182] R. Gunawan, M. Jung, E. Seebauer, and R. Braatz, "Optimal control of rapid thermal annealing in a semiconductor process," *J. Process Control*, vol. 14, no. 4, pp. 423–430, 2004.

[183] COMSOL, "Rapid Thermal Annealing," https://www.comsol.com/model/rapid-thermal-annealing-504.

[184] T. Gebel, L. Rebohle, R. Fendler, W. Hentsch, W. Skorupa, M. Voelskow, W. Anwand, and R. A. Yankov, "Millisecond annealing with flashlamps: Tool and process challenges," in *Proc. RTP*, 2006, pp. 47–55.

[185] P. Timans, J. Gelpey, S. McCoy, W. Lerch, and S. Paul, "Millisecond annealing: Past, present and future," *MRS Proceedings*, vol. 912, pp. 0912–C01–01, 2006.

[186] B. Walsh, H. Utomo, E. Leobandung, A. Mahorowala, D. Mocuta, K. Miyamoto, M. Kumar, S. Huang, M. Gribelyuk, A. Gabor, G. Freeman, B. Dirahoui, S. Deshpande, A. Azuma, A. Chan, E. Maciejewski, J. Herman, G. Berg, J. Zimmerman, H. Kimura, E. Nowak, R. Logan, O. Glushchenkov, N. Zamdmer, and I. Ahsan, "RTA-driven intra-die variations in stage delay, and parametric sensitivities for 65nm technology," in *Proc. VLSIT*, 2006, pp. 170–171.

[187] B. H. Calhoun and A. Chandrakasan, "Static noise margin variation for subthreshold SRAM in 65-nm CMOS," *IEEE J. Solid-St. Circ.*, vol. 41, no. 7, pp. 1673–1679, 2006.

[188] S. V. Kumar, K. H. Kim, and S. S. Sapatnekar, "Impact of NBTI on SRAM read stability and design for reliability," in *Proc. ISQED*, 2006, pp. 6–218.

[189] A. B. Kahng, "A mixed open-source and proprietary EDA commons for education and prototyping," in *Proc. ICCAD*, 2022.

[190] Parallax Software, Inc., "OpenSTA," 2022, https://github.com/The-OpenROAD-Project/OpenSTA.

[191] G. Guennebaud, B. Jacob *et al.*, "Eigen v3," 2010, http://eigen.tuxfamily.org.

[192] The OpenROAD Project, "OpenROAD-flow-scripts regressions," https://jenkins.openroad.tools/job/OpenROAD-flow-scripts-Nightly-Public/lastSuccessfulBuild/Report/.

# Appendix A

# PDNSim TCL-based Usage

In this appendix, we describe the various TCL commands and arguments needed to generate the different outputs of PDNSim in the OpenROAD application. PDNSim can be downloaded and installed the same way as OpenROAD [125]. It uses a TCL interface to interact with the user.

Listing A.1: PDNSim TCL commands in the OpenROAD application.

```
1 set_pdnsim_net_voltage -net <net_name> -voltage <voltage_value>
2 check_power_grid -net <net_name>
3 analyze_power_grid -net <net_name>
                     [-vsrc <voltage_source_location_file>]
                     [-outfile <filename>]
                     [-enable_em]
                     [-em_outfile <filename>]
                     [-dx]
                     [-dy]
                     [-em_outfile <filename>]
                     [-node_density <node_pitch>]
                     [-node_density_factor <factor>]
4 write_pg_spice  -net <net_name> `
                  -outfile <netlist.sp>
                  [-vsrc <voltage_source_location_file>]
```

OpenROAD has four TCL commands related to PDNSim that are presented in listing A.1. The first TCL command listed on line 1 sets the supply on the power or

ground net specified. It takes two arguments a net name and a voltage value. If this command is not invoked, PDNSim uses a default supply voltage from the liberty value that is available in OpenDB. The second command on line 2 performs the connectivity check on the net specified as an argument. The third command performs power grid analysis and has several arguments as detailed below:

- `net:` (mandatory) is the name of the net to analyze, power or ground net name.
- `vsrc:` (optional) file to set the location of the power C4 bumps/IO pins.
- `dx,dy`: (optional) these arguments set the bump pitch to decide the voltage source location in the absence of a vsrc file. Default bump pitch of 140um used in absence of these arguments and vsrc.
- `enable_em:` (optional) is the flag to report current per power grid segment. outfile: (optional) filename specified per-instance voltage written into file.
- `em_outfile:` (optional) filename to write out the per segment current values into a file, can be specified only if `enable_em` is flag exists.
- `node_density:` (optional) This value can be specified by the user in um to determine the node density on the std. cell rails. Cannot be used together with `node_density_factor`.
- `node_density_factor:` (optional) Integer value factor which is multiplied by standard cell height to determine the node density on the std. cell rails. Cannot be used together with `node_density`. Default value is 5.

The fourth command on line 4 of Listing A.1 creates a SPICE netlist representation of the power grid ($GV = J$ system of equations). The arguments include the net name whose netlist must be extracted, the name of the output netlist file (`outfile`), and the optional `vsrc` file with the locations of the C4 bumps or I/O power pins. The description of all these commands can also be found in PDNSim readme as a part of the OpenROAD GitHub repository [39].