Leveling Up: A Trajectory of OpenROAD, TILOS and Beyond

Andrew B. Kahng
UC San Diego Departments of CSE and ECE, La Jolla, CA, USA
abk@ucsd.edu

ABSTRACT

Since June 2018, the OpenROAD project has developed an open-source, RTL-to-GDS EDA system within the DARPA IDEA program. The tool achieves no-human-in-loop generation of design-rule clean layout in 24 hours. This enables system innovation and design space exploration, while also democratizing hardware design by lowering barriers of cost, expertise and risk. Since November 2021, The Institute for Learning-enabled Optimization at Scale (TILOS), an NSF AI institute for advances in optimization partially supported by Intel, has begun its work toward a "new nexus" of AI, optimization, and the leading edge of practice for use domains that include IC design. This paper traces a trajectory of "leveling up" in the research enablement for IC physical design automation and EDA in general. This trajectory has OpenROAD and TILOS as waypoints, and advances themes of openness, infrastructure, and culture change.

CCS CONCEPTS

• **Hardware → Electronic design automation**; **Physical design (EDA)**; • **Computing methodologies → Artificial intelligence**; **Machine learning**; • **Applied computing → Computer-aided design**.

KEYWORDS

Physical design; RTL-to-GDS; machine learning; IC design; open source; metrics; design productivity; optimization; system design.

ACM Reference Format:
Andrew B. Kahng. 2022. Leveling Up: A Trajectory of OpenROAD, TILOS and Beyond. In *Proceedings of the 2022 International Symposium on Physical Design (ISPD '22), March 27–30, 2022, Virtual Event, Canada.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3505170.3511479

1 INTRODUCTION

Scaling is the realization of more value while consuming less resources (energy, time, area, cost). Scaling makes the impossible possible, propelling EDA, IC design and the broader semiconductor ecosystem forward into the future. For EDA, scaling means *optimization* that is faster, better and cheaper – because EDA *is* optimization. There are never enough resources for optimization.

With the slowdown of device and process scaling, more burden is placed on "equivalent scaling" from EDA that improves IC product quality, development schedule and cost. Here, *machine learning* (ML) offers important boosters to EDA that enable better chip designs in less time. ML models provide predictions that can be leveraged in design exploration, while also serving as objectives for higher-level optimizations. ML also helps to solve difficult optimizations, while also affording new perspectives on classic optimization formulations via frameworks such as learning to optimize [9] or graphical neural network (GNN)-based embedding [5]. Deploying learning and optimization on modern cloud and GPU compute resources provides new paths to scalable "EDA2.0".

This paper describes a trajectory of "leveling up" in the research enablement for IC physical design automation and EDA in general. At its core, this trajectory is one of openness, infrastructure and culture change. It spans open source, open data, benchmarking and roadmapping – and it ultimately seeks scaling and empowerment of people and a future workforce, as well as culture change at the research interface between academia and industry. One waypoint in this trajectory is The OpenROAD Project, which since June 2018 has developed an open-source, RTL-to-GDS EDA system within the DARPA IDEA program. A hoped-for second waypoint is The Institute for Learning-enabled Optimization at Scale (TILOS), an NSF AI institute for advances in optimization, partially supported by Intel, that began work in November 2021 toward a "new nexus" of AI, optimization, and the leading edge of practice for use domains that include IC design.

ISPD '22, March 27–30, 2022, Virtual Event, Canada

ACM ISBN 978-1-4503-9210-5/22/03.
https://doi.org/10.1145/3505170.3511479

2 THE OPENROAD PROJECT

The OpenROAD project https://theopenroadproject.org/ seeks to deliver an open-source RTL-to-GDS tool that generates manufacturable layout from input RTL – in 24 hours, with no human in the loop – as part of the DARPA IDEA program. To date, the OpenROAD tool has been used in over 180 tapeouts in technologies ranging from 12nm to 130nm. This section will review the project's goals, status, and contributions to a trajectory of "leveling up" for the physical design and EDA field. These contributions include disruptive enablement for semiconductor design innovation and education, and support for research toward a future nexus of AI/ML, optimization, and IC design and EDA.

2.1 OpenROAD Goals

OpenROAD tackles a crisis which has been decades in the making: Hardware design, and system innovation in hardware, are simply too difficult. Design process outcomes are difficult to predict, even with expert tool users. As a result, hardware innovation faces barriers of cost, expertise, and risk. A potential solution is intelligent automation, with "self-driving" design tools and flows. OpenROAD is part of the DARPA IDEA program, which broadly aims for Hardware Compilers 2.0 – automated generation of manufacturable layout in 24 hours, with no human in the loop, and eventually with no loss of quality of results in Power, Performance or Area (PPA). The IDEA program shifts the focus from tools that squeeze out every last picosecond or microwatt from the manufacturing technology, to "self-driving" tools that require neither expertise nor complex, manually-derived tool settings to tape out a working chip.

OpenROAD's scope is digital IC design: the tool takes Verilog hardware description in, and delivers a merged tapeout-ready GDSII

layout database. As described in [1], achieving 24-hour automation requires three foundational base technologies: extreme partitioning to decompose the design problem into bite-sized chunks; intelligent orchestration of distributed and parallel optimization using cloud resources; and machine learning to model and predict what will happen when a given tool is run on a given design input with a given target. *Freedoms from choice* are also expected in a no-humans tool, just as a self-driving car will eventually have no steering wheel.

2.2 OpenROAD Status

OpenROAD v1.0 in Summer 2020 achieved a modern, integrated EDA tool architecture – industry-strength database and timing analysis, Tcl and Python scripting interfaces for users – plus design-rule clean layout generation in the GF12LP process. OpenROAD v2.0 in Summer 2021 brought new features such as early chip-package planning, and parasitic extraction to feed timing signoff. Users have brought up OpenROAD on new commercial technologies, and given feedback that has resulted in numerous improvements to usability. v2.0 PPA improved over v1.0 by roughly the equivalent of a full technology node of scaling. A recent no-human-in-the-loop *auto-tuning* capability finds better tool configurations than expert human users.

The use of OpenROAD for tapeouts of user designs includes over 180 tapeouts in SkyWater 130nm and an SoC tapeout in GF12LP; Army Research Labs has a project in flight in Intel 22FFL. OpenROAD is also the heart of the ASIC design platform from Efabless called OpenLane, which has been used in the design work for tapeouts on four Google-SkyWater SKY130 shuttle runs. The OpenLane repo became part of The OpenROAD Project's organization GitHub in July 2021. OpenROAD is also the default EDA solution in Efabless's commercial offering called ChipIgnite. The project is growing a community of users and developers. Users range from novices to experts, with applications that include Trust, 3D-IC, and AI and machine learning. Figure 1 shows daily statistics from a two-week period at the end of July 2021. The upper plot shows nearly a thousand downloads a day from nearly four thousand unique GitHub users.

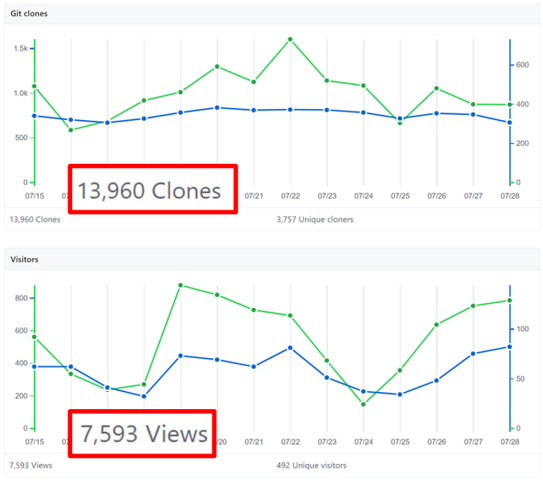

Figure 1: OpenROAD GitHub metrics in July 2021.

Key evolutions and facets of the OpenROAD ecosystem span infrastructure, engines, user support, and partnerships with industry and professional societies. Following are 12 examples.

- **The DATC Robust Design Flow.** OpenROAD works closely with IEEE CEDA's Design Automation Technical Committee, whose main task for several years has been to develop and support the RDF academic reference design flow [7] [8]. RDF was launched in 2016, based on winning tools and benchmark suites from academic contests in the physical design field. Since 2019, RDF has included the OpenROAD tool chain and hence a complete RTL-to-GDS flow. This backplane for research stimulates flow-scale research as well as new cross-stage optimizations.
- **OpenDB as *RosettaStone*.** The IEEE CEDA RDF includes contest benchmarks and contest-winning tools that in some cases date from over 15 years ago. These benchmarks and tools exist in a "parallel universe", whereby old tools could not be run on modern designs, and old benchmarks could not be fed to modern tools. A root cause is that past academic contests never established a consistent, underlying *data model*. On the other hand, OpenROAD's OpenDB database realizes an industry-standard (LEF/DEF 5.8) physical design data model. Along with OpenDB APIs, this enables *RosettaStone* [21] to connect not just the past, present and future of academic physical design research, but research and industry practice as well. With RosettaStone, all Bookshelf-based contest benchmarks since 2005 can be run through modern P&R tools, and modern designs in current enablements (e.g., a RISC-V core in foundry 12nm) can be run through old academic tools that are available only as executable binaries. Figure 2 contrasts the file-based tool chain paradigm with integration using an in-memory database and common data model.
- **AutoTuner.** Autotuning encompasses black-box optimization or adaptive sampling that seeks successful combinations of tool hyperparameters within a given budget of tool runs. OpenROAD's AutoTuner provides true no-human-in-the-loop operation and obtains better QoR than can be found by internal expert users. As detailed in [16], AutoTuner uses the Python Ray and Tune APIs [27], which afford easy-to-use parallelization and in-built search algorithms. AutoTuner is now being deployed on Google Cloud in collaboration with Google engineers.
- **METRICS2.1.** The open-source nature of OpenROAD lets it enable research on machine learning for EDA and IC design, in a way that is currently impossible with commercial EDA tools. Again in collaboration with IEEE CEDA, the project has published large-scale metrics datasets harvested from thousands of runs of the OpenROAD design flow, to stimulate data collection and publication that we hope will enable machine learning research. The recent METRICS2.1 syntax and OpenROAD-based realization, plus public repos of metrics data and example Jupyter notebooks for visualization and ML applications, are detailed in [16]. OpenROAD's logging and metrics infrastructure also feed into the score function evaluations used by the AutoTuner.
- **RTL-MP.** Another recent advance in the project is RTL-MP, a new RTL macro placer [25]. RTL-MP tries to "mimic'' the behavior of human experts. It captures the dataflow defined by RTL designers by exploiting logical hierarchy, and by extracting logical modules-based connection signatures. Further, RTL-MP understands relevant constraints such as macro placement blockages and preferred locations for macros.

- **External Contributions.** There have been many contributions from outside the formal project team. [37] currently lists over 50 contributors to the main OpenROAD application. An early external contribution by students in Brazil spanned global routing, clock tree synthesis and pin placement [13]. A more recent example is Professor Andrew Kennings' contribution of the *dpo* detailed placement optimizer, a toolkit that includes MIS matching, global moves and swaps, vertical moves and swaps, optimal reordering, and greedy random improvement.
- **Software Methodology.** A key strength of OpenROAD is its software development methodology and testing within an open-source, "fork and pull" framework. Project members [38] include both student researchers and EDA industry veterans, enabling a balance between prototyping for innovation and robust feature development. A continuous integration framework using Jenkins and GitHub hooks ensures that no code that breaks tests can be integrated. Secure confidential data is also used to test the system, ensuring tool stability for partners with confidential data [24].
- **ASAP7/ASAP5.** Since 2020, OpenROAD has been a home to the ASAP PDK and library development [10]. Open-sourcing of the ASAP7 PDK and libraries in December 2020 was a milestone for the research community. An ASAP5 PDK and libraries is nearing release; ASAP5 uses horizontal nanowire transistors, with ground rules that track advanced-node scaling boosters such as single diffusion break and contact over active gate.
- **Calibrations.** OpenROAD and the IEEE CEDA DATC also collaborate to support academic research on signoff electrical analyses. The *Calibrations* effort [34] leverages open PDKs to produce open data against which academic researchers can measure and report progress. Reference analysis results for STA, RC extraction, and IR drop analysis are obtained based on running OpenROAD with open designs and enablements (currently four designs and three PDKs); data include post-route .def and .spef, .v, .sdc, and worst paths and endpoint slacks in an open *.json format. Available scripts introduce noise for data obfuscation as contributors desire.
- **GUI.** The quest to achieve critical mass, critical quality, maximum velocity and maximum openness – while supporting clean foundry tape-out automation and the needs of early-adopter users – is why a project that aims for no-human-in-loop automation has a GUI. The GUI supports both developers and users, as shown in the project's 2021 DARPA ERI Summit presentation [35]. Figure 3 shows a placement density heatmap and the option (requested by a Ph.D. student in Brazil) to dump this data to a file for machine learning.
- **Efabless and OpenLane.** Beyond the use of Efabless's OpenLane flow in numerous SKY130 tapeouts, OpenLane is widely used in IC design education around the world (e.g., [40]), and the two teams are increasingly focused on OpenLane development.
- **Outreach.** Outreach to stakeholders has also been a key project goal. Active engagements beyond IEEE CEDA include CHIPS Alliance and Si2. The project has organized and sponsored prizes for academic contests (e.g., ICCAD-2019, TAU-2020), organized new workshops and meetings (e.g., the Workshop on OpenSource EDA Technology (WOSET)), and presented tutorials and invited talks at numerous conferences. This outreach complements OpenROAD's organic growth as an open-source software project.

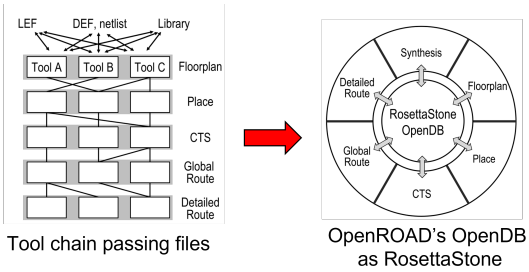

Figure 2: *RosettaStone* based on OpenDB.

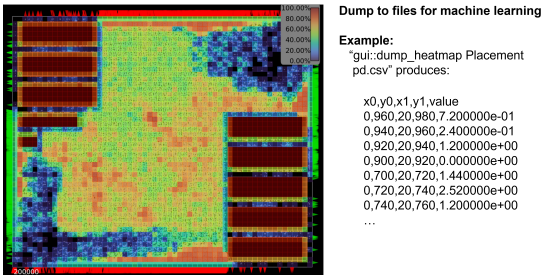

Figure 3: Heatmap of placement density and writing to CSV.

2.3 Trajectory Impacts of OpenROAD

OpenROAD has provided "leveling up" in several key dimensions that amplify researchers' efforts and reduce wasted energy. This helps scale the efficiency of discovery and innovation in physical design automation, as well as the education and training of a future workforce.

- **Openness can succeed, and infrastructure matters.** OpenROAD has developed and open-sourced core EDA infrastructure, including an industry-strength database and static timing engine, within a modern, integrated EDA tool architecture. The tool is built to last, with a robust software engineering methodology. It is now a viable platform for future research and innovation. From a training and workforce development standpoint, OpenROAD uniquely documents in open source how robust EDA software is architected and developed. For physical design research, the project provides an integration backplane that reduces the effort barrier to starting research, accelerates exploration and evaluation of new ideas, improves transparency and reproducibility, and otherwise "lifts all boats".
- **There is a virtuous cycle of democratization and innovation.** OpenROAD has enabled high-schoolers to tape out working chips and served as the tool platform for thousands of students in online VLSI design courses. The project has the potential to be a disruptive enabler for semiconductor design innovation, education and workforce development. Seeing the virtuous cycle of openness, democratization and impact is rewarding for project contributors, and sparks broader culture change as well.[1]

[1] Ongoing culture change is seen in the growing body of open-sourced research software and research data in the literature. The research community is also experiencing the zeitgeist of replication and reproducibility that has swept over fields such as AI [14] and funding agencies such as the U.S. NIH [36].

- **Reaching critical mass and critical quality is essential.** OpenROAD has from its outset been "Not Research As Usual". The project has afforded further understanding of the concept of open-source EDA [17] [18], adding to the learnings from past efforts such as the Bookshelf [4], OA Gear [32] and METRICS [12] [23]. In retrospect, the qualitative aiming point of "critical mass and critical quality" has been invaluable to help keep in mind the bar for new, lasting research foundations. Aiming points such as this also reveal the skill sets, attitudes and collaborations needed to develop new foundations.[2]
- **Sustainability and transition are first-class concerns.** Developing a tool in open source naturally raises the issue of sustainability and transitions beyond the traditional research funding model. Comprehending and managing this issue with adequate lead time is another important aspect of "leveling up". As a project, OpenROAD has had to sharpen its thinking about many aspects of its transition to sustainability, such as project governance, stable funding and/or revenue model, organizational structure and non-technical skillsets, technology roadmap, and stakeholders.

After four years, original aims of OpenROAD such as accelerating research, enabling reproducibility of leading-edge research, and providing an open platform for data collection and machine learning have become important features of the project. The project team plans to continue pursuing "critical mass and critical quality" on multiple axes, and to continue as a platform for basic research on EDA, including research toward a future nexus of AI/ML, optimization, and IC design and EDA. Looking forward, OpenROAD's technology roadmap will advance both functionality (improved support for power intent, hierarchical and incremental flows, low-power and other PPA optimizations, and exploitation of modern compute fabrics) and adoptability (robustness, usability, and the traditional criteria of QOR, capacity and turnaround time). Transition and sustainability demand continued attention to building both the developer and user communities.[3]

3 THE TILOS AI INSTITUTE

TILOS – The Institute for Learning-enabled Optimization at Scale – is an NSF National AI Research Institute for advances in optimization, partially supported by Intel Corporation. The institute began operations in November 2021 with a mission to "make impossible optimizations possible, at scale and in practice".

There are six universities in TILOS: UCSD, MIT, National University, Penn, UT-Austin, and Yale. The institute seeks a *new* nexus of AI and machine learning, optimization, and use in practice. Figure 4 shows four virtuous cycles envisioned for the institute: (i) mutual advances of AI and optimization provide the foundations; (ii) challenges of scale, along with breakthroughs from scaling, bind together foundations and the use domains of chip design, networks and robotics; (iii) the cycle of translation and impact brings research and the leading edge of practice closer together; and (iv) the cycle of research, education, and broadening participation grows the field and its workforce.[4]

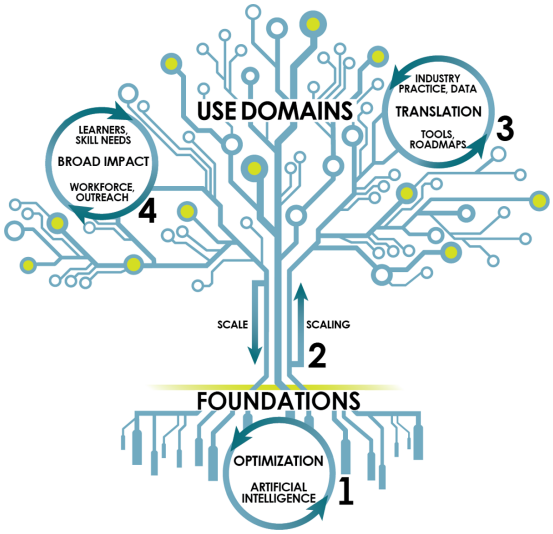

Figure 4: Four virtuous cycles in TILOS.

3.1 Learning and Optimization Research

Practical optimization brings well-known challenges: (i) instances have enormous scale; (ii) representations and abstractions are crucial to success; (iii) objectives are hazy, particularly with multi-stage optimizations and dynamic settings; (iv) optimization tools must provide reliability and generalization; and (v) scaling of productivity increasingly demands new ways to learn and optimize using modern compute fabrics.[5] Addressing these challenges requires new foundations and an interplay of finding the right representations, developing the right machine learning methods, and advancing optimization methods.

At a high level, TILOS research at the AI-optimization interface will pursue five thrusts [39] [20]: (i) bridges between continuous and discrete optimization; (ii) parallel, distributed and federated optimization; (iii) optimization on manifolds; (iv) dynamic decision making under uncertain environments; and (v) nonconvex optimization in deep learning.

3.2 The Chip Design Use Domain

Chip design brings challenges that include hierarchical-system context; extreme cost of training data; "multi-everything" (physics, objectives) constrained optimization; and pervasive security aspects.

[2] Maintaining a long view is also helpful. Time scales, moving parts and overheads in the project (e.g., recruiting personnel, tool development with robust software methodology, meeting bars for adoption, cycles of user interaction and support, emergence of collaborations, establishing dialogues with stakeholders, etc.) can be off the charts in a field where a new idea can become a Ph.D. student's accepted conference paper submission in a matter of weeks. However, the long view ultimately pays off in the form of failure-proofing and lasting impact.

[3] OpenROAD's community of contributors began with members of the project, then expanded to others with interest in developing the core technology. It now includes users interested in completing designs and adding new methodologies to the flow.

[4] Impacting the leading edge of optimization in practice requires new collaborations between foundations, use-domains and industry practitioners. This "tripartite matching" challenge, along with an ethos of being "a whole that is greater than the sum of its parts", brings hurdles of "team science" [11]. However, only with such an aiming point and ethos can fewer than 30 faculty, with a similar number of Ph.D. students and postdoctoral scholars, move the needle at the the leading edge.

[5] In TILOS, the three use domains of chip design, networks and robotics bring diverse optimization challenges but inspire shared solutions, due to commonalities such as physical embeddedness, hierarchical-system context, underlying graphical models, safety and robustness as first-class concerns, and the bridging of human-guided and autonomous systems.

The scale of configuration spaces motivates research on (i) discovery of exploitable structure in cost landscapes; (ii) distributed, data-driven sampling and search methods; and (iii) metaheuristic "templates" to match discovered instance structure with optimization strategies. Given its decades-long history as a driver of applied optimization and automation, chip design also highlights *augmenting rather than rediscovering* domain expertise, by encoding expert knowledge and intuition to serve optimization and decision-making agents.

Five TILOS faculty comprise the Chips team: Sicun (Sean) Gao, Andrew Kahng, Farinaz Koushanfar, David Pan and Tajana Rosing. At a high level, Chips research is organized into four thrusts.

- **Layout** has a long-range goal of creating manufacturable chip layouts directly from a circuit netlist description. This thrust includes topics such as optimal embedding of hypergraphs; development of a modern partitioning toolkit; and studying the nexus of sampling, sequential decision-making, L2O (learning to optimize) and cloud compute fabrics.
- **Verification** broadly seeks scaling of verification and validation methods. This thrust includes the topic of interior search methods to scale SMT solving and verification methods.
- **Quantifying the cost of "X"** (in particular, X = Security) studies the intrinsic cost of robustness in optimization and learning, with respect to aspects such as data anonymity, data integrity, and privacy in federated and distributed settings.
- **Data, Benchmarking and Roadmapping** supports (i) relevance and translation of research into real-world contexts, (ii) availability of abundant and relevant data for machine learning, and (iii) development of consensus benchmarks and roadmaps to support research on practical optimization. Aspects of this thrust range from generation of "artificial but realistic" benchmarks to principles and frameworks for ethical and fair benchmarking.

3.3 Challenges of Translation

Recall from Figure 4 that *translation* at the interface between industry and academia is the third virtuous cycle of TILOS. Figure 5 shows an idealized life cycle of translation, in which real-world practitioners supply problems and data, researchers bridge foundations and use domains to discover new methods, and these results go back into the real world.

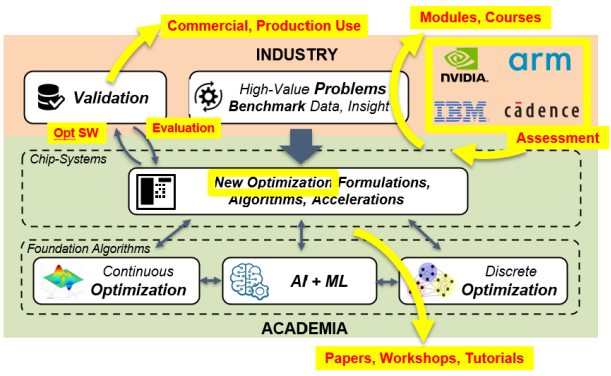

Figure 5: Life cycle of research and translation.

Unfortunately, today this picture is spoiled by fundamental technical and cultural obstacles. For example: (i) relevant datasets are unavailable to researchers; (ii) access to research data is a world of haves (e.g., TSMC7 + Arm Cortex) vs. have-nots (e.g., FreePDK45 + ISCAS89); and (iii) a culture of reproducibility and benchmarking is still work in progress. Obstacles such as these cause inefficiencies that waste or, worse yet, drive away precious human resources. Thus, TILOS aims to scale *people* in addition to optimization in practice, by pioneering new democratizations, new cultural and scientific or technical norms, and principled bases for looking forward and investing resources. This is only fitting for an AI Institute for advances in optimization that has chip design as a key use domain. Examples of associated research challenges include the following.

- Real data is rare and proprietary, and is shared only with elite researchers. An example research challenge is to democratize: Can we develop a science of "data virtual reality", enabling generation of artificial circuit designs that are indistinguishable from real circuit designs from the perspective of optimizers? Can artificial-but-realistic instances help quantify suboptimality gaps and distributions? A related challenge is to learn with less real data (e.g., via data augmentation and transfer learning).
- Not all researchers have the same access to the strongest optimizers. An example research challenge is to model outcomes of a strong optimizer based on instance attributes and outcomes of weak optimizers.
- Data sources, designs and tools often must be kept anonymous. An example research challenge is to develop trusted methods for privacy-preserving anonymization and obfuscation of design and related data. A complementary challenge is to develop trusted tests for identity leakage.[6]
- Companies forbid benchmarking in large part because benchmarks can be misused. An example challenge for both research and culture is to develop the principles and mechanisms that provide a foundation for fair benchmarking. More broadly, consensus on reporting and comparison methodology is needed at the nexus of applied ML and optimization [3] [31].
- Researchers hope to work on relevant, high-impact problems. However, it is difficult to discern the most crucial research targets for learning and optimization. An example challenge is to foster a practice of roadmapping needs and potential solutions for learning and optimization technology, according to projections of what drives (i.e., demands) such technology.

3.4 Trajectory Impacts of TILOS

It will be some time before outcomes and impacts from TILOS are visible. However, *if* TILOS achieves its goals, there will be "leveling up" in additional dimensions such as the following.

- **Cracking the code of translation.** TILOS aims to broadly address the challenge of *translation* via methodologies and tools, infrastructures and culture change. Efforts to unblock data and benchmarking will promote transparency and reproducibility, as well as more efficient use of researchers' time. Similarly, efforts toward roadmapping of core optimization problems will support

[6] Sometimes, private data cannot be exposed but research tools can be brought in and applied to the data. OpenROAD's secure CI infrastructure provides a turnkey framework for this.

the steering of time and effort to problems for which progress matters the most.

- **Improved understanding of suboptimality.** Quantifying the suboptimality in practice of EDA optimizations gives insight into where further QOR or TAT improvements might be found. Research on federated learning and optimization, and on sequential decision-making and sampling, will help escape today's preoccupation with the single-server, overnight optimization context [19] [20]. Theory and methods will be needed to respond to questions such as, "What QOR benefit is expected from running *N* more copies of the optimizer for *M* more days?"
- **A new nexus of learning, optimization and practice.** Aspects of a "new nexus" include directions noted earlier (e.g., data efficiency, representation and embedding) along with (i) learning to learn (e.g., appropriate loss functions / optimization objectives within a multi-stage flow); (ii) learning to optimize [26] and hybridize learning with combinatorial optimization [2] [5]; (iii) application of available ML techniques such as transformers or deep reinforcement learning [29] that previously required prohibitive computational resources; and more [22].

4 AND BEYOND ...

The past four years of OpenROAD and the coming five years of TILOS are just one segment in a CAD/EDA trajectory that is already well into its seventh decade.

- OpenROAD has contributed the following aspects of "leveling up": (i) culture changes toward openness, and open-sourcing of core EDA infrastructure; (ii) democratization of access and innovation, and enablement of education and workforce development; (iii) improved transparency, replication and reproducibility at the leading edge of research; (iv) showing how the qualitative goal of "critical mass and critical quality" informs many facets of project development as well as a long view; and (v) improved understanding of open-source EDA's unique challenges and potential paths to sustainability and transition.
- TILOS aspires to additional aspects of "leveling up": (i) cracking the code of translation, via research and infrastructure that unblocks data as well as the benchmarking and roadmapping of core EDA optimizations; (ii) refocusing on quantification of suboptimality in practical optimization, as a complement to roadmapping in guiding research and investment; and (iii) understanding how to form and nurture integrative, multi-disciplinary collaborations that bridge learning, EDA optimization and industry practice to achieve impact in practice.
- Both OpenROAD and TILOS aim to scale and amplify the efforts of *people* – students, researchers, a future workforce – through democratizations and lowered barriers to access. Both lie along a path to "machine learning-enabled DA" or "EDA2.0", as shown in Figure 6. In addition, TILOS research in the Chips use domain is able to leverage OpenROAD and other open-sourced components from both UCSD and the University of Texas [28] [6].

What next? Figure 6 suggests that some elements of a future trajectory, such as autotuning, will mature earlier than other elements, such as an ecosystem-wide consensus on fair EDA benchmarking. Hopefully, the bulk of Figure 6 will become well-understood over the next 5-10 years.

What other elements in a trajectory of "leveling up" are important to solve in the coming decade? (i) *How will a new nexus of AI, optimization and chip design draw in more young people and be more exciting?* Training a future workforce depends on attracting it in the first place. It is critical to further improve accessibility and openness. (ii) *How will we fully leverage the centuries of past advances in optimization?* It is critical to mind the suboptimality gap and to scale the reach of optimizers, e.g., by splitting up problems, or making them smaller or sparser, without losing solution quality. (iii) *What are the most critical optimizations to target with development of future MLDA or EDA2.0?* One candidate is the partition-cluster-shape-pack-plan co-optimization that is inevitably at the core of system physical implementation. High-accuracy early design space exploration and pathfinding depend on this as we enter a beyond-everything (Moore, CMOS, von Neumann, ...) era. (iv) *What advances in representation are needed, to complement advances in learning and optimization?* Current ML for EDA largely ignores expert designer knowledge, as well as the stack of models and symbolic representations that underlie chip design. Thus, a candidate challenge is to unite data-centric and knowledge-centric approaches along lines of "third-wave AI".

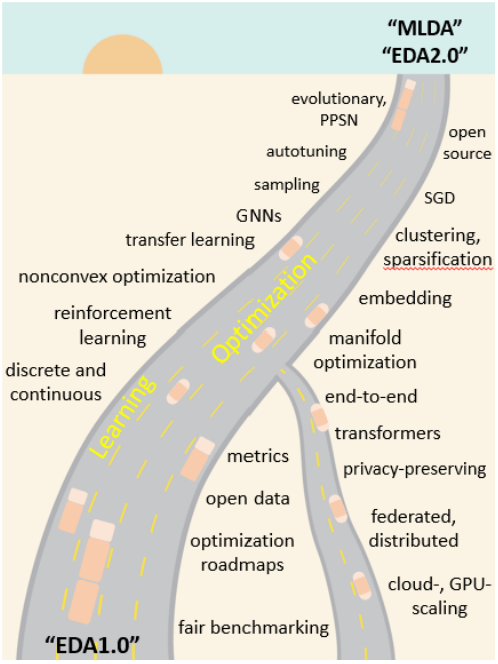

Figure 6: OpenROAD and TILOS provide several initial steps along this trajectory of machine learning and EDA (figure adapted from [19]).

ACKNOWLEDGMENTS

Many thanks are due to colleagues in the OpenROAD project and the TILOS AI Institute, as well as numerous students and collaborators. Research at UCSD is partially supported by NSF CCF-2112665 and CCF-1564302, DARPA HR0011-18-2-0032 and FA8650-20-2-7009, Samsung, Qualcomm, NXP Semiconductors and the C-DEN center.

REFERENCES

[1] 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", *Proc. DAC*, 2019, pp. 76:1-76:4.

[2] Y. Bengio, A. Lodi and A. Prouvost, "Machine Learning for Combinatorial Optimization: a Methodological Tour d'Horizon", *Eur. J. Oper. Res.*, 290(2) (2021), pp. 405-421.

[3] A. E. Caldwell, A. B. Kahng, A. A. Kennings and I. L. Markov, "Hypergraph Partitioning for VLSI CAD: Methodology for Heuristic Development, Experimentation and Reporting", *Proc. DAC*, 1999, pp. 349-354.

[4] A. E. Caldwell, A. B. Kahng and I. L. Markov, "Toward CAD-IP Reuse: The MARCO GSRC Bookshelf of Fundamental CAD Algorithms", *IEEE Design and Test of Computers* 19(3) (2002), pp. 70-79.

[5] Q. Cappart, D. Chételat, E. Khalil, A. Lodi, C. Morris and P. Veličković, "Combinatorial Optimization and Reasoning with Graph Neural Networks", *arXiv*, 2102.09544, 2021. https://arxiv.org/abs/2102.09544

[6] H. Chen, M. Liu, B. Xu, K. Zhu, X. Tang, S. Li, Y. Lin, N. Sun and D. Z. Pan, "MAGICAL: An Open-Source Fully Automated Analog IC Layout System from Netlist to GDSII" *IEEE Design & Test*, 38(2) (2020), pp. 19-26.

[7] J. Chen, I. H.-R. Jiang, J. Jung, A. B. Kahng, V. N. Kravets, Y.-L. Li, S.-T. Lin and M. Woo, "DATC RDF-2020: Strengthening the Foundation for Academic Research in IC Physical Design", *Proc. ICCAD*, 2020.

[8] J. Chen, I. H.-R. Jiang, J. Jung, A. B. Kahng, S. Kim, V. N. Kravets, Y.-L. Li, R. Varadarajan and M. Woo, "DATC RDF-2021: Design Flow and Beyond", *Proc. ICCAD*, 2021, pp. 1-6.

[9] T. Chen, X. Chen, W. Chen, H. Heaton, J. Liu, Z. Wang and W. Yin, "Learning to Optimize: A Primer and A Benchmark", *arXiv preprint* 2103.12828, 2021. https://arxiv.org/abs/2103.12828

[10] L. 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 J.* (53) (2016), pp. 105-115. https://github.com/The-OpenROAD-Project/asap7

[11] N. J. Cooke and M. L. Hilton, Eds., *Enhancing the Effectiveness of Team Science*, The National Academies Press, 2015. https://www.nap.edu/download/19007

[12] S. Fenstermaker, D. George, A. B. Kahng, S. Mantik and B. Thielges, "METRICS: A System Architecture for Design Process Optimization", *Proc. DAC*, 2000, pp. 705-710.

[13] M. Fogaça, E. Monteiro, M. Danigno, I. Oliveira, P. F. Butzen and R. Reis, "Contributions to OpenROAD from Abroad: Experiences and Learnings", *Proc. ICCAD*, 2020, pp. 1-8.

[14] W. D. Heaven, "AI is Wrestling with a Replication Crisis", *MIT Technology Review*, 2020. https://www.technologyreview.com/2020/11/12/1011944/artificial-intelligence-replication-crisis-science-big-tech-google-deepmind-facebook-openai/

[15] G. Huang, J. Hu, Y. He, J. Liu, M. Ma, Z. Shen, J. Wu, Y. Xu, H. Zhang, K. Zhong, X. Ning, Y. Ma, H. Yang, B. Yu, H. Yang and Y. Wang, "Machine Learning for Electronic Design Automation: A Survey", *ACM TODAES* 26(5) (2021), pp. 40:1-40:46.

[16] J. Jung, A. B. Kahng, S. Kim and R. Varadarajan, "METRICS2.1 and Flow Tuning in the IEEE CEDA Robust Design Flow and OpenROAD", *Proc. ICCAD*, 2021, pp. 1-9. https://github.com/ieee-ceda-datc/datc-rdf-Metrics4ML

[17] A. B. Kahng, "Looking Into the Mirror of Open Source", *Proc. ICCAD*, 2019, pp. 1-8.

[18] A. B. Kahng, "Open-Source EDA: If We Build It, Who Will Come?", *Proc. VLSI-SoC*, 2020, pp. 1-6.

[19] A. B. Kahng, "Machine Learning for CAD/EDA: The Road Ahead", *IEEE Design and Test (Special Issue on Machine Learning for CAD/EDA)* (2022), to appear.

[20] A. B. Kahng, "AI/ML, Optimization and EDA in TILOS, an NSF National AI Research Institute", Synopsys APUP Special Session, January 2022. https://vlsicad.ucsd.edu/NEWS22/Synopsys-APUP-v3-ACTUAL-NOBACKUP.pptx

[21] A. B. Kahng, M. Kim, S. Kim and M. Woo, "RosettaStone: Connecting the Past, Present and Future of Physical Design Research", *manuscript in submission*, 2021. https://github.com/ABKGroup/RosettaStone

[22] A. B. Kahng, J. Lienig, I. L. Markov and J. Hu, *VLSI Physical Design: From Graph Partitioning to Timing Closure*, 2nd Ed., Springer, 2022.

[23] A. B. Kahng and S. Mantik, "A System for Automatic Recording and Prediction of Design Quality Metrics", *Proc. ISQED*, 2001, pp. 81-86. https://vlsicad.ucsd.edu/GSRC/metrics/

[24] A. B. Kahng and T. Spyrou, "The OpenROAD Project: Unleashing Hardware Innovation", *Proc. GOMACTech*, 2021.

[25] A. B. Kahng, R. Varadarajan and Z. Wang, "RTL-MP: Toward Practical, Human-Quality Chip Planning and Macro Placement", *Proc. ISPD*, 2022, to appear.

[26] K. Li and J. Malik, "Learning to Optimize", *arXiv:1606.01885*, 2017. https://arxiv.org/abs/1606.01885

[27] R. Liaw, E. Liang, R. Nishihara, P. Moritz, J. E. Gonzalez and I. Stoica, "Tune: A Research Platform for Distributed Model Selection and Training", *arXiv:1807.05118*, 2018. https://arxiv.org/abs/1807.05118 https://docs.ray.io/en/latest/tune/index.html

[28] Y. Lin, Z. Jiang, J. Gu, W. Li, S. Dhar, H. Ren, B. Khailany and D. Z. Pan, "DREAMPlace: Deep Learning Toolkit-Enabled GPU Acceleration for Modern VLSI Placement", *IEEE Trans. on CAD* 40(4) (2021), pp. 748-761.

[29] A. Mirhoseini, A. Goldie, M. Yazgan, J. W. Jiang, E. Songhori, S. Wang et al., "A Graph Placement Methodology for Fast Chip Design", *Nature* 594 (2021), pp. 207-212. https://github.com/google-research/circuit_training

[30] R. Sutton, "The Bitter Lesson", March 2019. http://www.incompleteideas.net/IncIdeas/BitterLesson.html

[31] T. Weise, X. Wang, Q. Qi, B. Li, and K. Tang, "Automatically Discovering Clusters of Algorithm and Problem Instance Behaviors as well as Their Causes from Experimental Data, Algorithm Setups, and Instance Features"' *Applied Soft Computing Journal*, 73 (2018), pp. 366–382.

[32] Z. Xiu, D. A. Papa, P. Chong, C. Albrecht, A. Kuehlmann, R. A. Rutenbar and I. L. Markov, "Early Research Experience with OpenAccess Gear: an Open Source Development Environment for Physical Design", *Proc. ISPD*, 2005, pp. 94-100

[33] (MARCO GSRC) VLSI CAD Bookshelf. http://vlsicad.eecs.umich.edu/BK/

[34] DATC RDF Calibrations. https://github.com/ieee-ceda-datc/datc-rdf-calibrations

[35] Electronics Resurgence Initiative (ERI) Summit & MTO Symposium, "Curation-Driven Open-Source Transition of the OpenROAD Chip Design Platform", October 2021. https://youtu.be/gG8Lc9IRbQ4

[36] NIH Data Sharing Policy and Implementation Guidance. https://grants.nih.gov/grants/policy/data_sharing/data_sharing_guidance.htm

[37] The OpenROAD Project. https:github.com/The-OpenROAD-Project [OpenROAD application: https:github.com/The-OpenROAD-Project/OpenROAD]

[38] The OpenROAD team. https://theopenroadproject.org/our-team/

[39] The TILOS Foundations team. https://tilos.ai/foundations/

[40] VLSI Systems Design. https://www.vlsisystemdesign.com/