


The OpenROAD Project

Andrew B. Kahng
UCSD Depts. of CSE and ECE
abk@ucsd.edu
<https://vlsicad.ucsd.edu>

SoC Labs Workshop
Arm Research Summit

The Crisis of **Design**: It's Too Difficult !

- Modern “Place and Route” tool: 10000+ commands/options
- Hard to design with latest tools in latest technologies
 - Quality and schedule are unpredictable
 - Expert users are required
- **Increased cost and risk block innovation**



U.S. DARPA IDEA: No-Humans, 24-Hours

IDEA will create a no-human-in-the-loop hardware compiler for translating source code to layouts of System-On-Chips, System-In-Packages, and Printed Circuit Boards in less than 24 hours

A. Olofsson, DARPA
ISPD-2018 keynote

- **Part of DARPA Electronics Resurgence Initiative**
- **Traditional focus: ultimate performance, power, area**
- **IDEA focus: ultimate ease of use and runtime**

Restore designer access to silicon

The OpenROAD Project

OpenROAD will create a no-human-in-the-loop digital hardware compiler for translating source code to layouts of System-On-Chips, System-In-Packages, and Printed Circuit Boards in less than 24 hours

A. Olofsson, DARPA
ISPD-2018 keynote

- **Traditional focus: ultimate performance, power, area**
- **OUR focus: ultimate ease of use and runtime**

Initial focus: Digital IC “RTL-to-GDS” flow

theopenroadproject.org

OpenROAD

HomePeopleNews and EventsPublicationsOutreach

DEMOCRATIZING HARDWARE DESIGN

The OpenROAD project attacks the barriers of Cost, Expertise and Uncertainty (i.e., Risk) that block the feasibility of hardware design in advanced technologies.

[READ MORE](#)

Open Source Tools

[SHOW ON GITHUB](#)

User Guide

[GETTING STARTED](#)

Community

[JOIN THE DISCUSSION](#)

OpenROAD

11 Sep

Final submissions of the "ICCAD 2019 LEF/DEF Based Global Routing Contest" are now being evaluated! See the contest description paper by #UCSD and #MentorGraphics here <https://t.co/7mW3mw9uPJ>. Results and open-source bonuses announced at #ICCAD!

#contest #VLSI

Twitter

[Load More...](#)

About OpenROAD

Problem: Hardware design requires too much effort, cost and time.

Challenge: \$\$\$ costs and "expertise gap" block system designers' access to advanced technology.

Foundations and Realization of Open, Accessible Design

Prof. Kahng and the OpenROAD team are aiming to develop open-source tools that achieve autonomous, 24-hour layout implementation.

Latest News and Events

ERI Summit 2019: OpenROAD presentation video posted

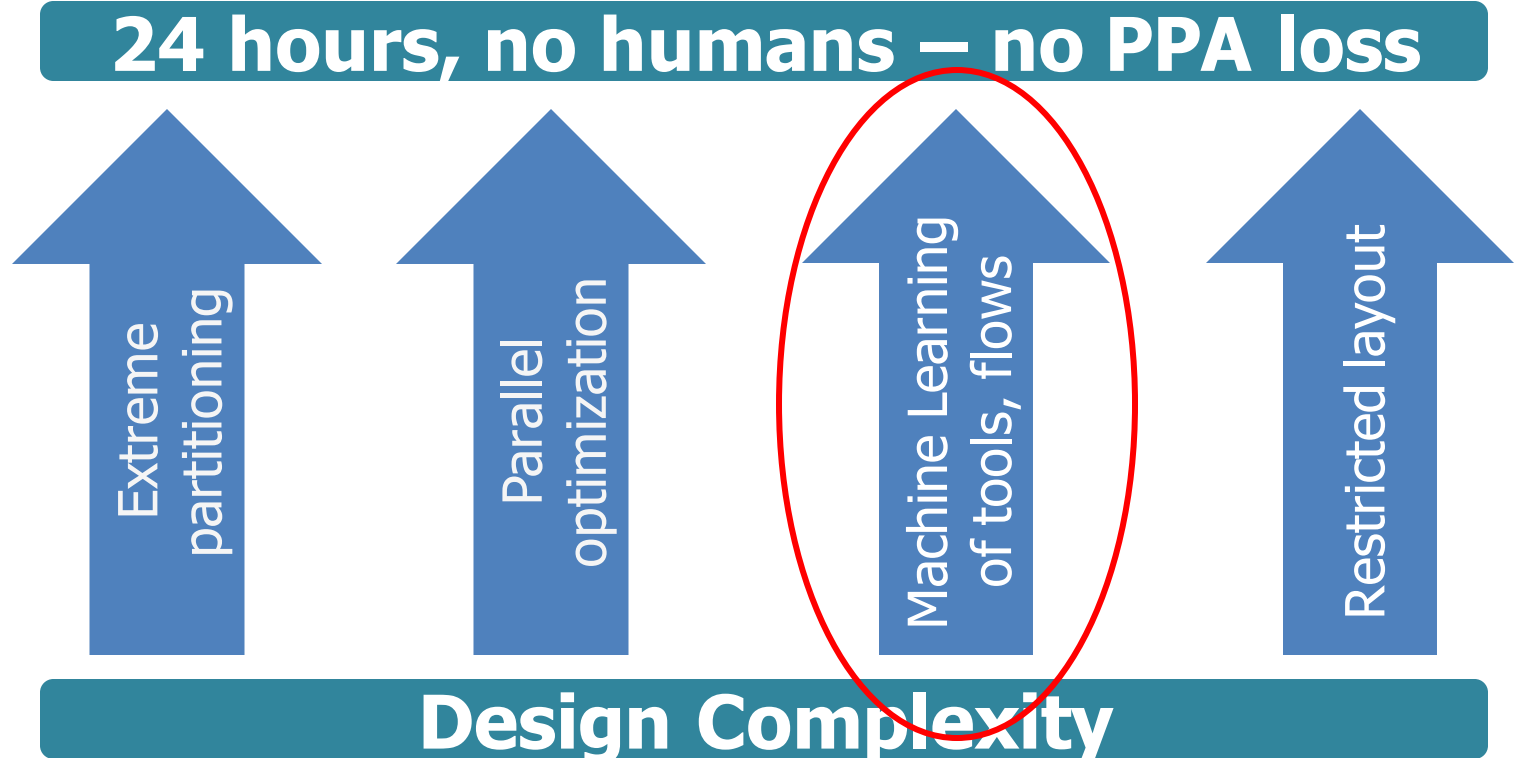
September 11, 2019

OpenROAD Alpha tools in the IEEE CEDA DATC's RDF-2019 flow!

A. B. Kahng, 190917 Research Enabled by Arm

5

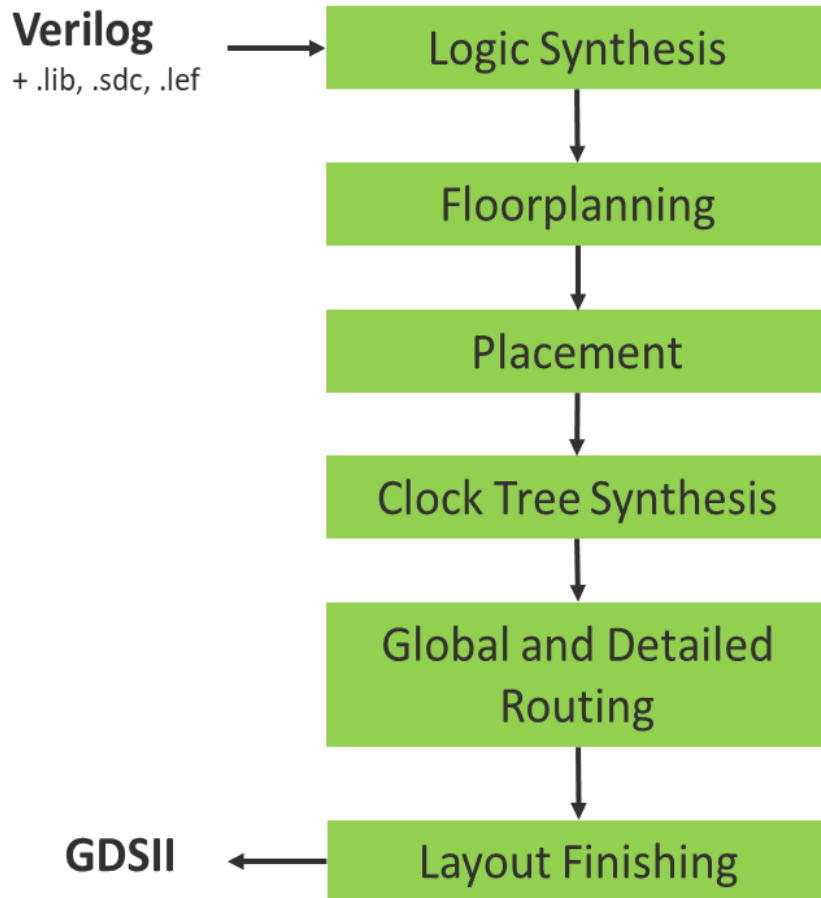
OpenROAD's Foundation Technologies



Mindsets

- Use cloud/parallel to achieve **predictability**, recover solution **quality**
- Relentless focus on **time, effort reduction**





Initial Focus: Digital IC “RTL-to-GDS” Flow



- **Alpha release: July 2019**
- **Alpha milestone: design-rule clean layout in foundry (65nm) node with Arm stdcells, IPs**
 - Essential building block for eventual no-humans layout generation
- **v1.0 release: June 2020**

See: <https://theopenroadproject.org/publications/>
<https://github.com/The-OpenROAD-Project>

July 2019 “alpha” <https://github.com/The-OpenROAD-Project/alpha-release>



July 2019 “alpha”

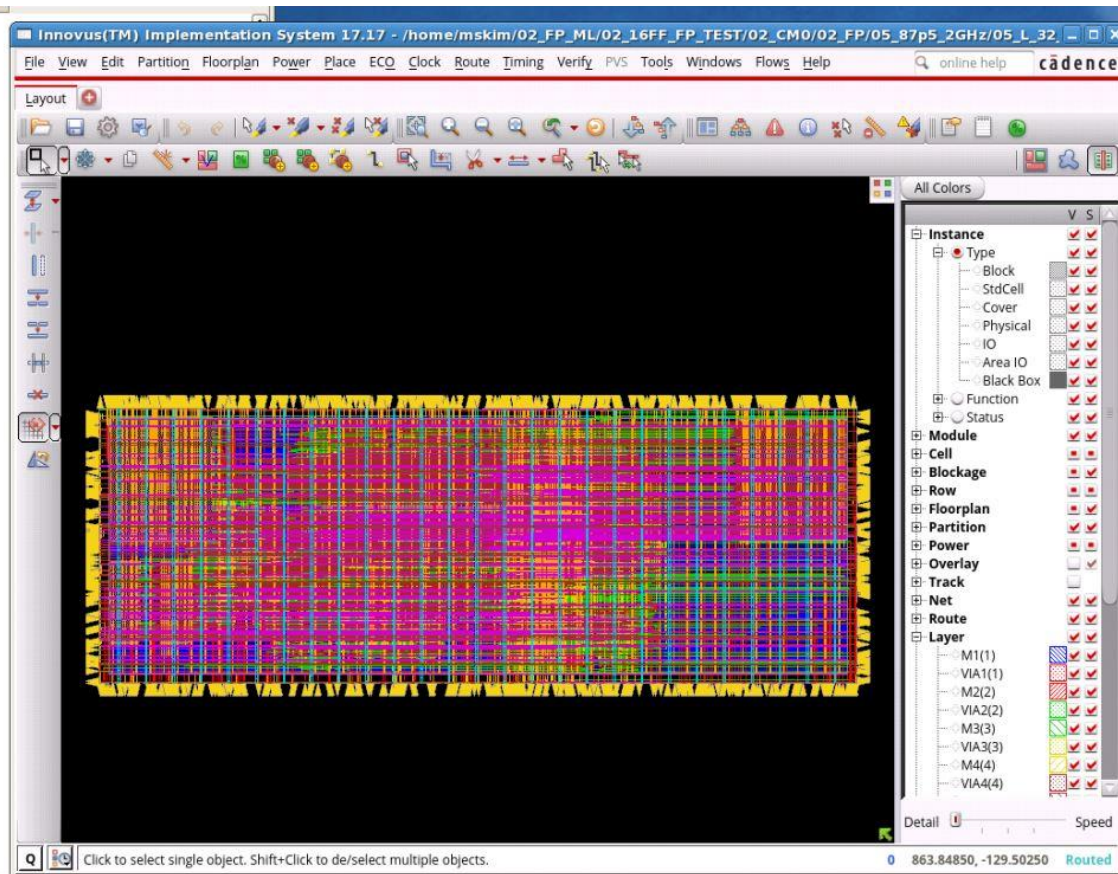
<https://github.com/The-OpenROAD-Project/alpha-release>

```
VERIFY DRC ..... Sub-Area: {254.880 0.000 339.840 80.640} 4 of 65 Thread : 1
VERIFY DRC ..... Sub-Area: {679.680 161.280 764.640 241.920} 35 of 65 Thread : 5
VERIFY DRC ..... Sub-Area: {764.640 0.000 849.600 80.640} 10 of 65 Thread : 6
VERIFY DRC ..... Sub-Area: {84.960 241.920 169.920 322.560} 41 of 65 Thread : 6
VERIFY DRC ..... Sub-Area: {84.960 80.640 169.920 161.280} 15 of 65 Thread : 3
VERIFY DRC ..... Sub-Area: {424.800 0.000 509.760 80.640} 6 of 65 Thread : 2
VERIFY DRC ..... Sub-Area: {594.720 0.000 679.680 80.640} 8 of 65 Thread : 7
VERIFY DRC ..... Sub-Area: {764.640 161.280 849.600 241.920} 36 of 65 Thread : 5
VERIFY DRC ..... Sub-Area: {509.760 241.920 594.720 322.560} 46 of 65 Thread : 4
VERIFY DRC ..... Sub-Area: {169.920 241.920 254.880 322.560} 42 of 65 Thread : 6
VERIFY DRC ..... Sub-Area: {849.600 322.560 934.560 400.000} 63 of 65 Thread : 4
VERIFY DRC ..... Sub-Area: {339.840 241.920 424.800 322.560} 44 of 65 Thread : 0
VERIFY DRC ..... Sub-Area: {934.560 322.560 1019.520 400.000} 64 of 65 Thread : 4
VERIFY DRC ..... Sub-Area: {254.880 80.640 339.840 161.280} 17 of 65 Thread : 1
VERIFY DRC ..... Sub-Area: {1019.520 322.560 1100.000 400.000} 65 of 65 Thread : 1
VERIFY DRC ..... Sub-Area: {84.960 161.280 169.920 241.920} 28 of 65 Thread : 6
VERIFY DRC ..... Sub-Area: {764.640 80.640 849.600 161.280} 23 of 65 Thread : 5
VERIFY DRC ..... Sub-Area: {679.680 322.560 764.640 400.000} 61 of 65 Thread : 3
VERIFY DRC ..... Sub-Area: {424.800 80.640 509.760 161.280} 19 of 65 Thread : 2
VERIFY DRC ..... Sub-Area: {254.880 241.920 339.840 322.560} 43 of 65 Thread : 0
VERIFY DRC ..... Sub-Area: {594.720 80.640 679.680 161.280} 21 of 65 Thread : 7
VERIFY DRC ..... Sub-Area: {764.640 241.920 849.600 322.560} 49 of 65 Thread : 5
VERIFY DRC ..... Sub-Area: {764.640 322.560 849.600 400.000} 62 of 65 Thread : 5
VERIFY DRC ..... Sub-Area: {424.800 241.920 509.760 322.560} 45 of 65 Thread : 4
VERIFY DRC ..... Sub-Area: {169.920 80.640 254.880 161.280} 16 of 65 Thread : 6
VERIFY DRC ..... Sub-Area: {254.880 322.560 339.840 400.000} 56 of 65 Thread : 0
VERIFY DRC ..... Sub-Area: {339.840 80.640 424.800 161.280} 18 of 65 Thread : 2
VERIFY DRC ..... Sub-Area: {509.760 80.640 594.720 161.280} 20 of 65 Thread : 7
VERIFY DRC ..... Sub-Area: {594.720 241.920 679.680 322.560} 47 of 65 Thread : 3
VERIFY DRC ..... Sub-Area: {679.680 80.640 764.640 161.280} 22 of 65 Thread : 1
VERIFY DRC ..... Sub-Area: {424.800 322.560 509.760 400.000} 58 of 65 Thread : 4
VERIFY DRC ..... Sub-Area: {254.880 161.280 339.840 241.920} 30 of 65 Thread : 2
VERIFY DRC ..... Sub-Area: {424.800 161.280 509.760 241.920} 32 of 65 Thread : 7
VERIFY DRC ..... Sub-Area: {679.680 241.920 764.640 322.560} 48 of 65 Thread : 3
VERIFY DRC ..... Thread : 2 finished.
VERIFY DRC ..... Thread : 6 finished.
VERIFY DRC ..... Sub-Area: {594.720 322.560 679.680 400.000} 60 of 65 Thread : 7
VERIFY DRC ..... Thread : 7 finished.
VERIFY DRC ..... Sub-Area: {594.720 161.280 679.680 241.920} 34 of 65 Thread : 3
VERIFY DRC ..... Thread : 3 finished.

Verification Complete : 0 Viols.


*** End Verify DRC (CPU: 0:00:04.5 ELAPSED TIME: 1.00 MEM: 2.0M) ***


innovus 1> win
innovus 2> 
```



<https://github.com/The-OpenROAD-Project/>





← → ↻ GitHub, Inc. [US] | <https://github.com/The-OpenROAD-Project> 🔍 ☆ 📄 🗨️ 🔄 📁 📄

 Search or jump to... / Pull requests Issues Marketplace Explore 🔔 + 📁




The OpenROAD Project


OpenROAD seeks to develop and foster an autonomous, 24-hour, open-source layout generation flow (RTL-to-GDS).
<https://theopenroadproject.org/>

 **Repositories** 38  Packages  People 28  Teams 5 ⚙️ Settings

Pinned repositories

[Customize pinned repositories](#)

 **alpha-release** ≡
Builds, flow and designs for the alpha release
● Verilog 🍴 1

 **docs** ≡
OpenROAD Documentation
● Python ★ 1

Type: All ▾ Language: All ▾ New

alpha-release

Builds, flow and designs for the alpha release

● Verilog 🍴 1 ★ 0 ⓘ 0 🔄 0 Updated 1 hour ago

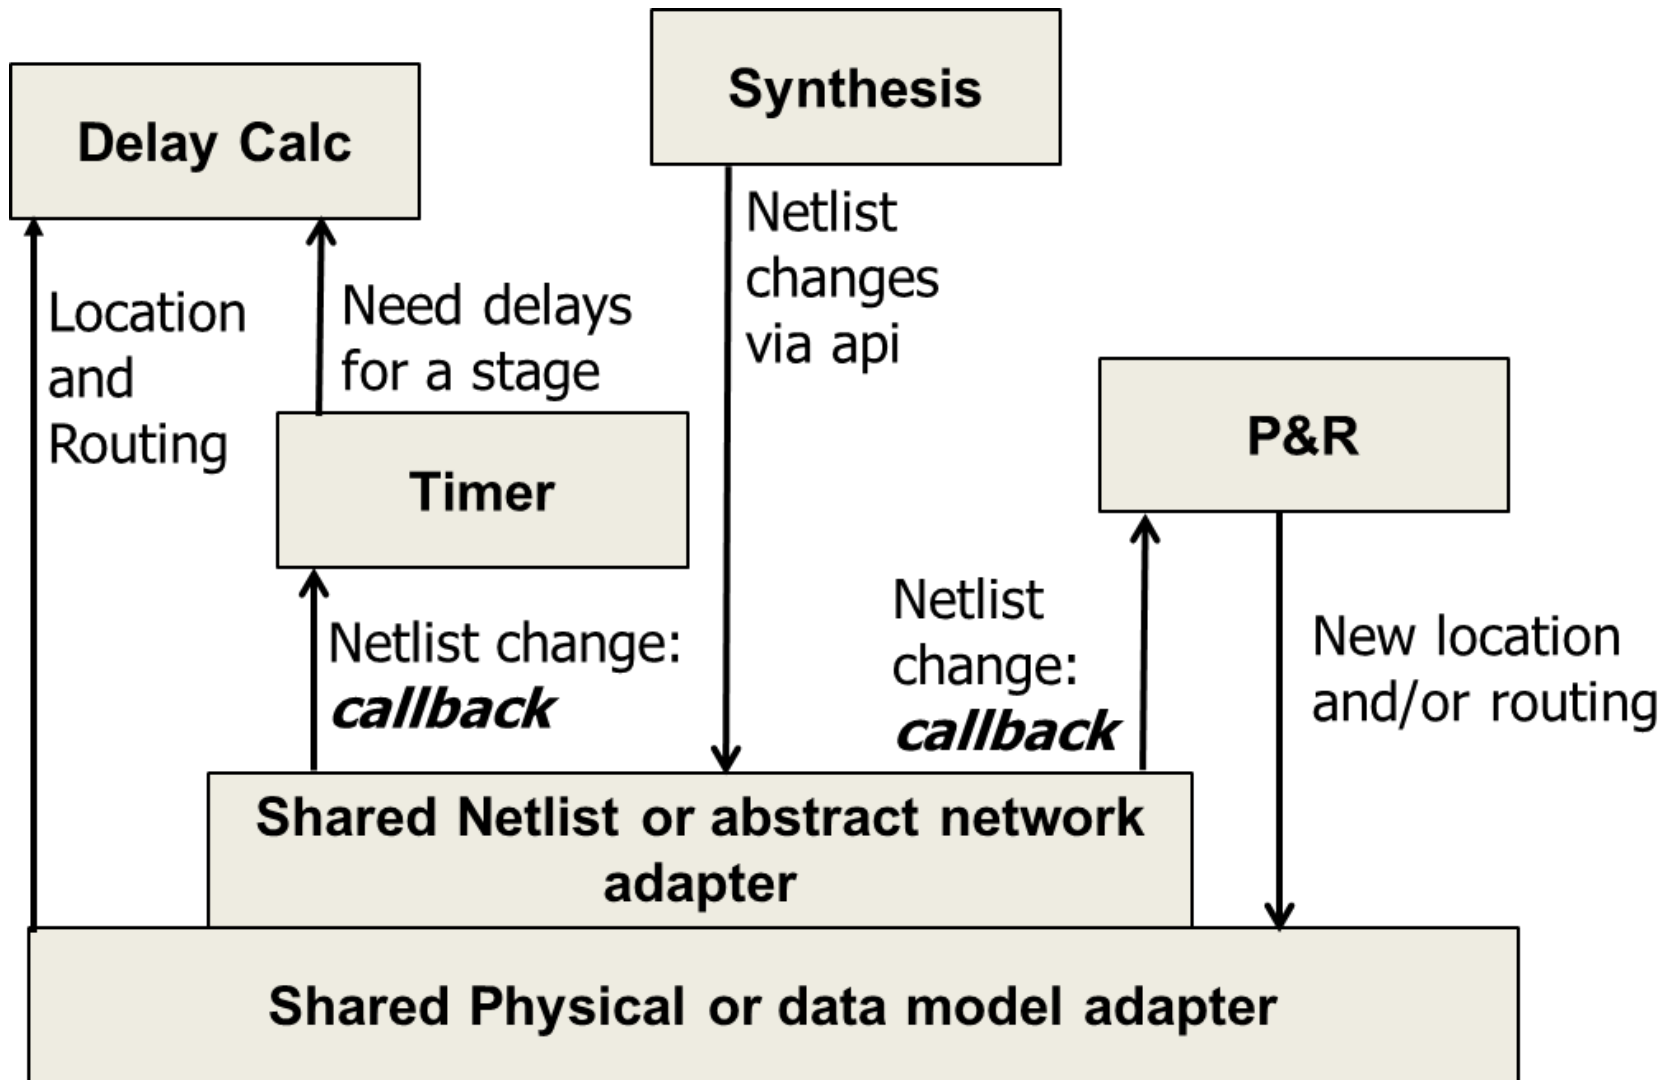
Top languages

● C++ ● Shell ● C ● Verilog
● Python

Looking Toward OpenROAD v1.0 (July 2020)

- EDA industry learning curve: architecture of synthesis, place & route
 - 1980s: file-based integration **tool chain** (“Alpha”, July 2019)
 - 2000s: tightly coupled algorithms on a shared incremental substrate
tight integration (v1.0, July 2020) → implies first-ever shared DB layer in permissive open source!
- Critical foundations: architecture, devops, database, AE

Incremental EDA Architecture: Shared Netlist



Looking Toward OpenROAD v1.0 (July 2020)




- EDA industry learning curve: architecture of synthesis, place & route
 - 1980s: file-based integration **tool chain** (“Alpha”, July 2019)
 - 2000s: tightly coupled algorithms on a shared incremental substrate **tight integration** (v1.0, July 2020) → implies first-ever shared DB layer in permissive open source!
- Critical foundations: architecture, devops, database, AE
 - Tom Spyrou: Well known EDA system architect, OpenAccess database, PrimeTime
 - *Full-time in San Diego; chief architect / technical project manager*
 - Lukas van Ginneken: Magma co-founder, synthesis and optimization
 - James Cherry: Pearl and Parallax static timing engines



OpenSTA Slack, WNS, TNS 28nm

<https://github.com/The-OpenROAD-Project/OpenSTA>


aes_cipher_top (28nm, 12T, clkp=1000ps)



aes_cipher_top	WNS (ps)	TNS (ps)	#viol.
Signoff STA	-61	-289	7
OpenSTA (arnoldi)	-57	-314	9

OpenSTA Slack, WNS, TNS 16nm

<https://github.com/The-OpenROAD-Project/OpenSTA>





Coyote (16nm, 9T, clkp=2000ps)

	Signoff STA	OpenSTA
WNS (ns)	-0.660	-0.603
TNS (ns)	-1758.004	-1219.239
#viol.	8096	6926


“Research Enabled by Arm” in OpenROAD



- **Proposal brainstorming**
 - Qualcomm + Arm = compelling for SOC/IP/methodology
- **Key technical contributions**
 - P/G mesh layout in Alpha release
 - 65LP bringup: ibex, swerve, blackparrot, ... (CA53)
- **Enabling 14nm NHIL, DRC-clean**
 - Scripted “veneers” for on-grid pins
 - Correct by construction P/G mesh
 - How to bring NHIL within grasp of tools
- **+ SOC scope, methodology**
- **+ Machine learning of flow, tools**



Area overhead ~1.5%



Some Broader Views

Open Source Enables EDA Research

- **Clarity**
 - *Today:* EDA license terms include “no benchmarking”
 - *Open Source:* Leading edge becomes visible, well-defined
- **Better science**
 - *Today:* Research in EDA is constrained to be irreproducible !
 - Cannot publish the commercial EDA Tcl scripts used in experiments
 - *Open Source:* Verifiable advances, no more “irreproducible results”
- **Avoid reinventing wheels**
 - *Today:* Students waste months trying to reconstruct papers. reimplementing basic algorithms and engines
 - *Open Source:* Field advances more rapidly

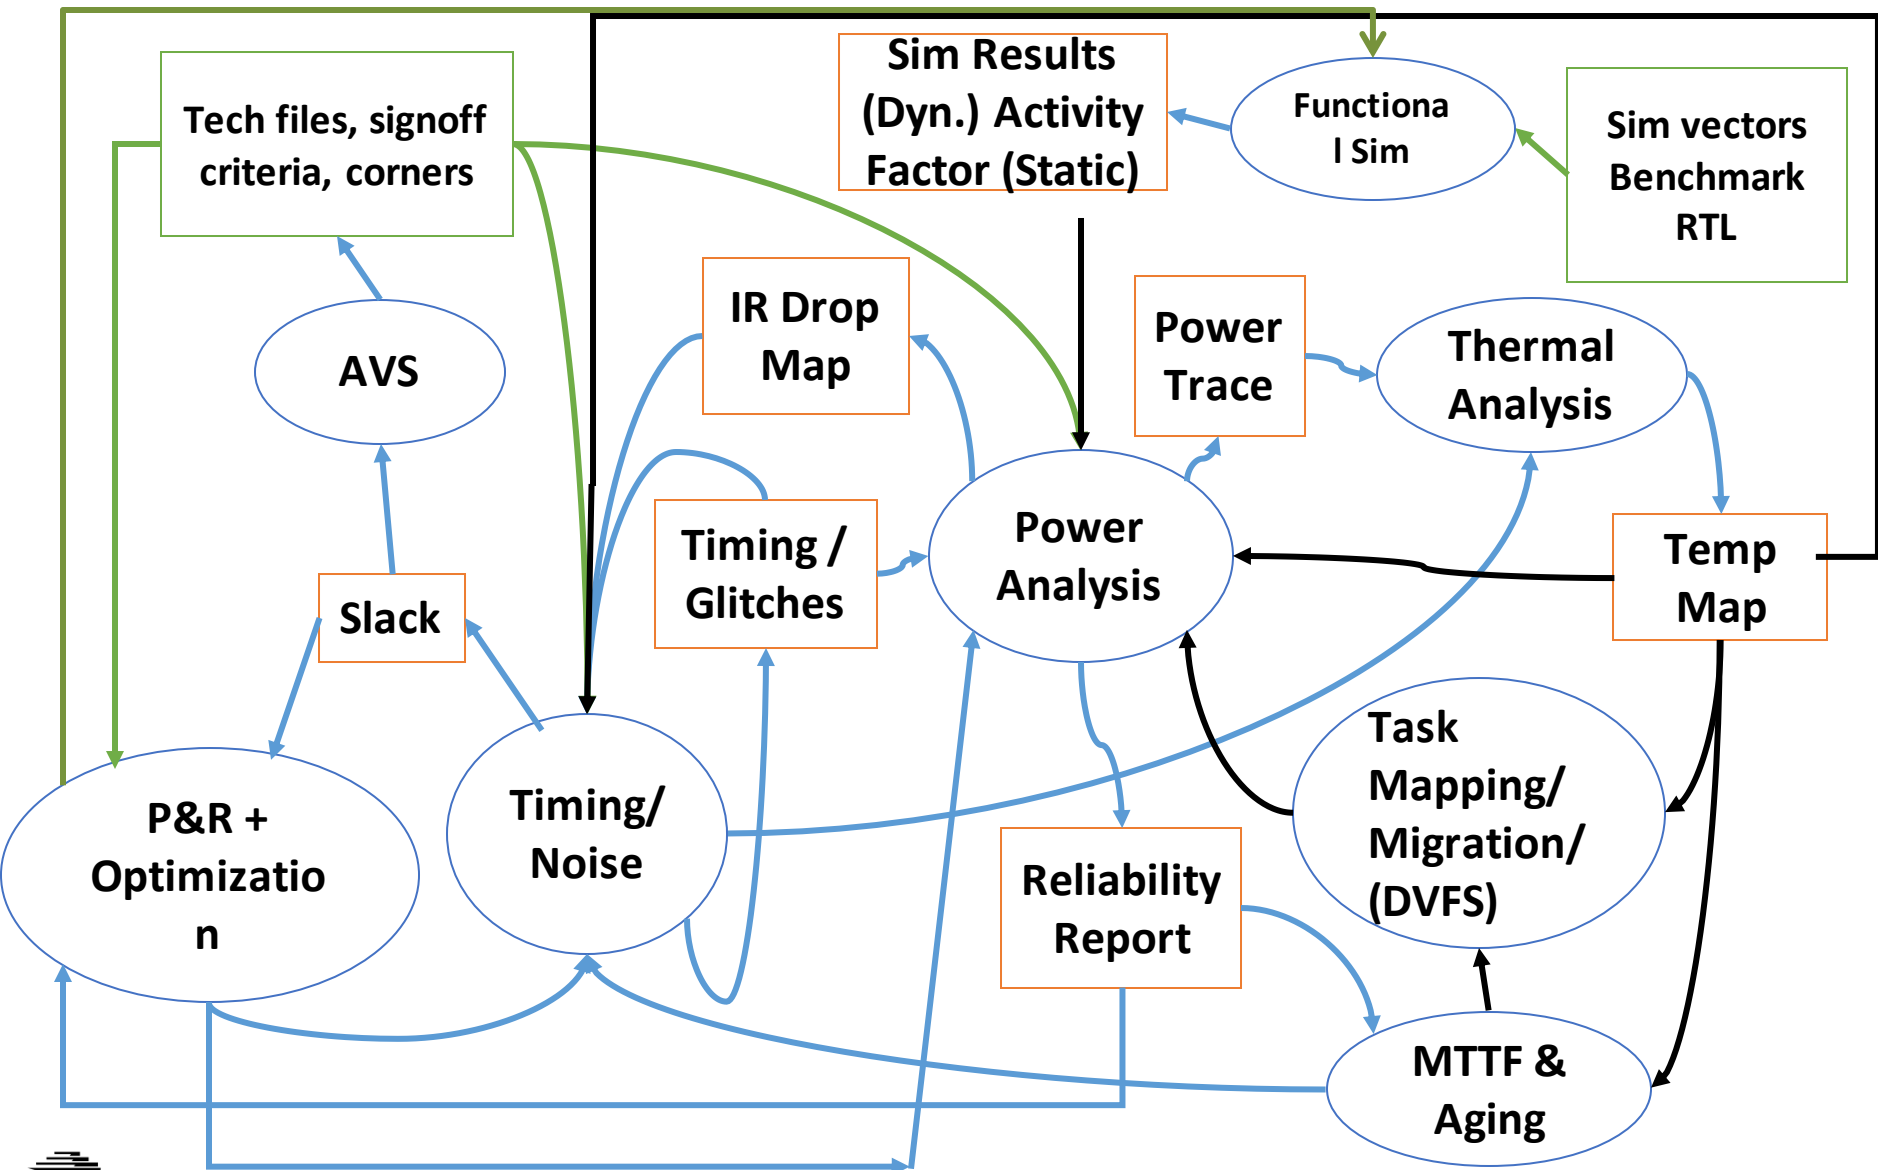
→ A new era for EDA research and research impact

(1) “Looking Into the Mirror of Open Source”, ICCAD-2019, <https://vlsicad.ucsd.edu/Publications/Conferences/374/c374.pdf>
(2) DAC-2019 panel: <https://vlsicad.ucsd.edu/NEWS19/ABK-PositionStatement-Sess75-v2-posted.pptx>

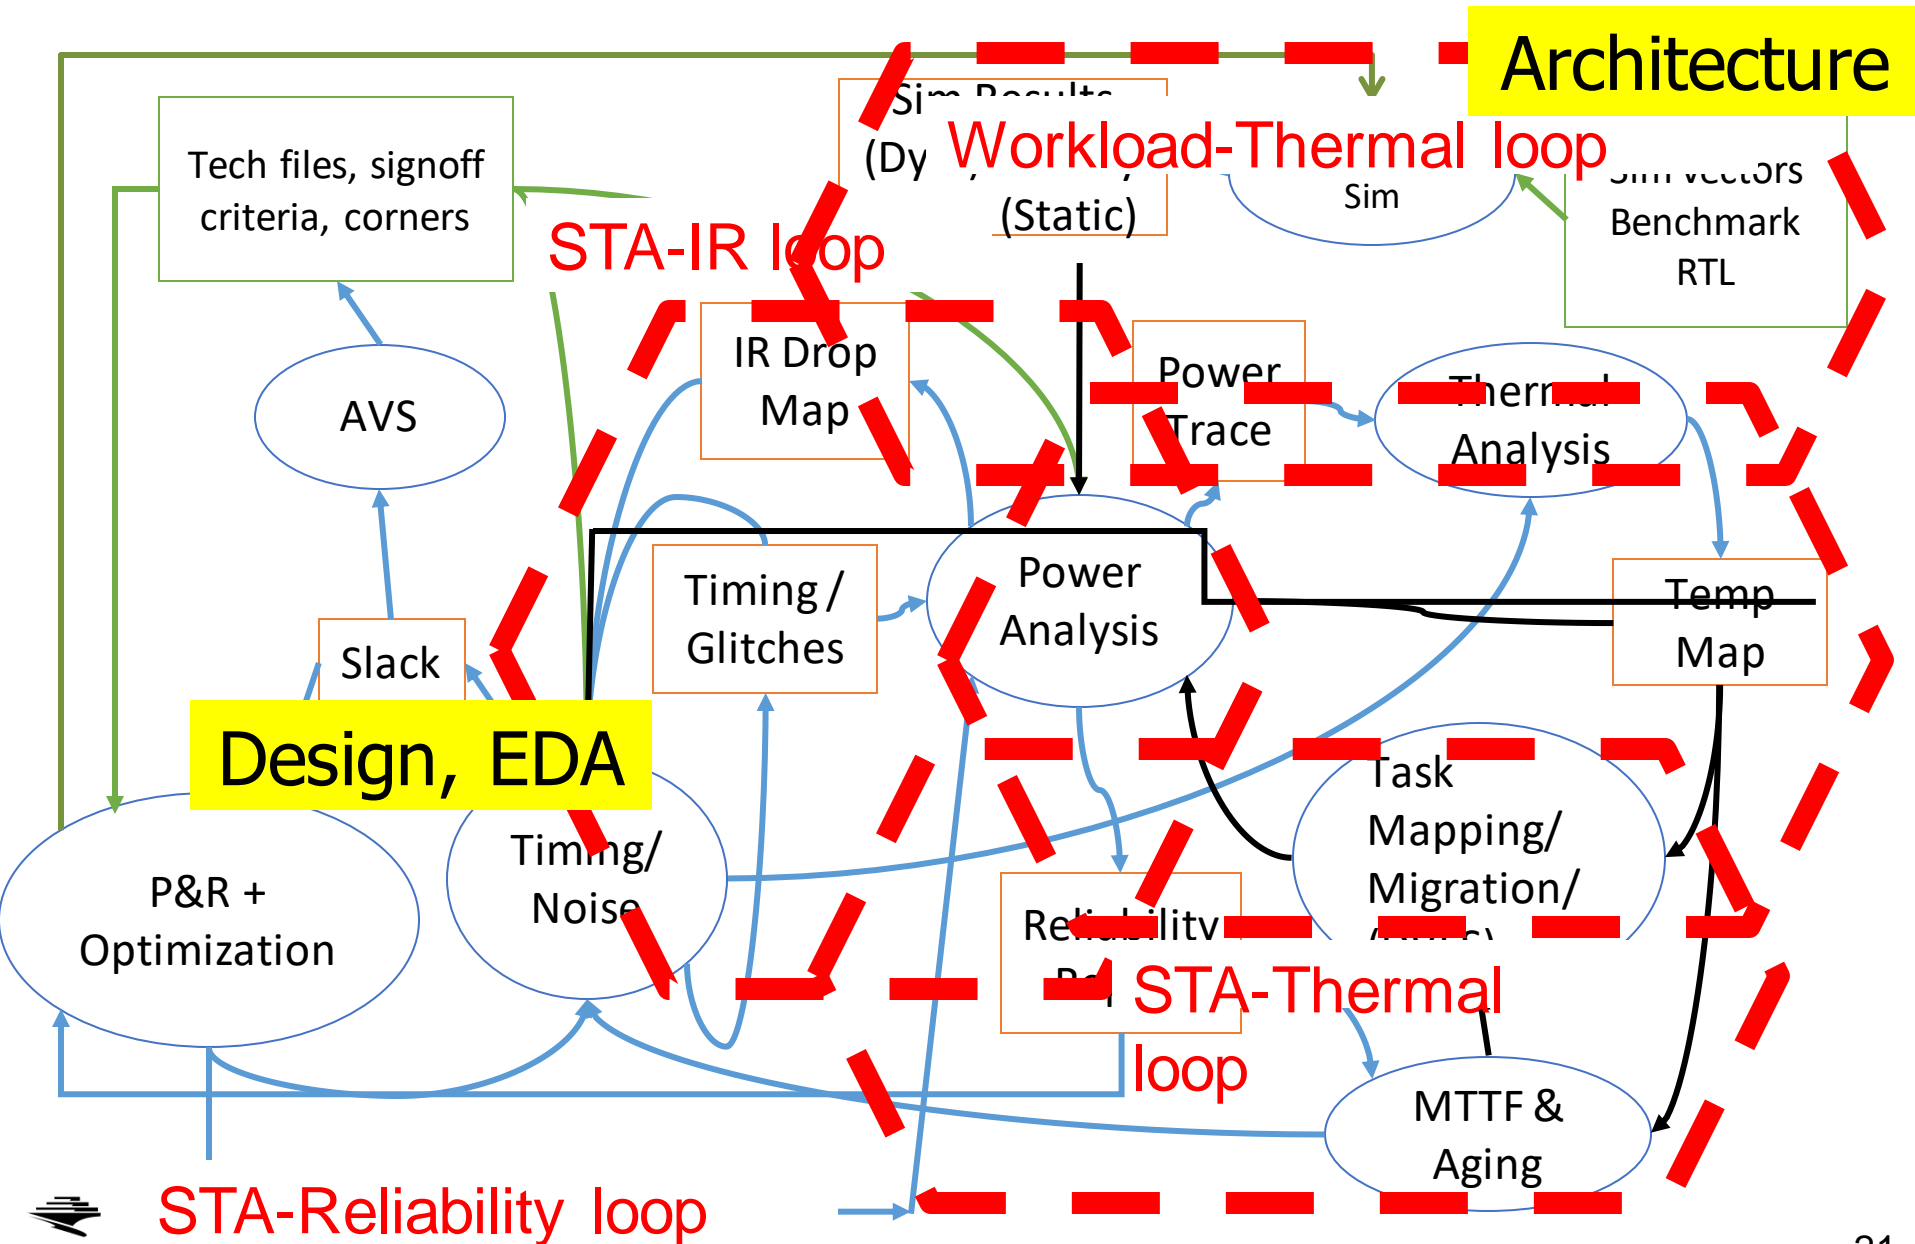
Open-Source EDA Enables Systems Research

- **IC design today: struggle for 1% power reduction**
 - → For EDA, 0.5% clock power reduction is a home run in a competitive benchmark today
- **Systems (ML, DSLs, IRs, uArch, ...) research today: 40%, 25x, etc. improvements frequently achieved**
 - My own recent draft on multi-die system integration: 50% iso-cost performance improvement, 30% iso-performance cost improvement
- **Begs question: What happened after cycle-accurate modeling, FPGA implementation – on the way to integrated silicon?**
 - Open-source EDA scales and is transparent, instrumentable
 - Can shine a light

How to Estimate... (power, reliability, ...) ?



(across analysis loops, scales, layers, ...)?



Open-Source EDA Enables ...

- ... **ML-powered, self-driving EDA tools and flows?**
 - ML needs data
 - Data comes from {designs} + {tools} + {technology}
 - → **collaboration opportunities for EDA, systems, ML**
- ... **Better modeling of what can be achieved in hardware for given design + tools + technology**
 - “40%” or “25x” can be more rigorously supported
 - Or, maybe 50% or 35x could have been found !
 - ML in and around (open-source?) tools and flows could help
 - → **paths to improve design space exploration, “pathfinding”, system synthesis ...**

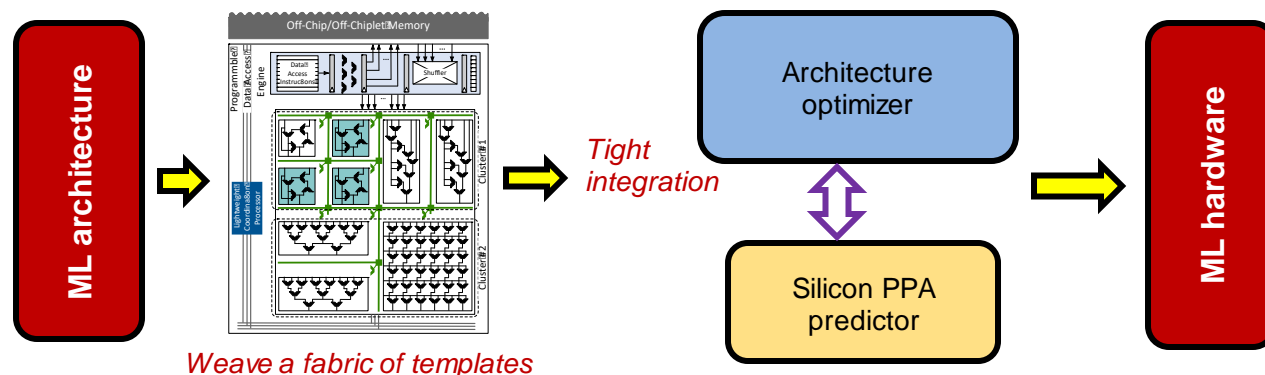
VeriGOOD-ML: Open-source Verilog Generator for RTML

DARPA Real-Time Machine Learning (RTML) program

Team: S. Sapatnekar (U. Minnesota), A. B. Kahng, H. Esmaeilzadeh (UCSD), J. Gu (Northwestern)

• Overview:

- No-human-in-loop Verilog generator for RTML ASICs to meet aggressive performance goals
- Real-time data input through fast I/O interfaces (e.g., AIB) in a chiplet-based ecosystem
- Wide variety of ML architectures targeted: SVM, CNN, RNN, Logistic regression, Decision tree, ...
- **Input:** ML spec (e.g., ONNX), performance specs (e.g., TOPS/W)
- **Output:** Verilog implementation that can be synthesized through a standard design flow
 - Will be benchmarked through hardware tapeouts, measured hardware performance
- **Approach:** **Template-based design/pathfinding** with *predictable performance, Pareto tradeoffs*



Broader View of “Research Enabled by Arm”


- **Opens doors, gives instant relevance and credibility**
 - Early relationship with Common Platform foundries: **2005**
 - M* cores and 7nm library: **2013-2014**
 - Studies at design-manufacturing interface
 - Opened door to IMEC 7nm, pre-release versions of Cadence Innovus
 - CA53, NEON, FinFET enablement: **early 2017**
 - Studies of “ultimate PPAC”
 - CA53 = “gold standard” for PPAC evaluation → credibility with large semis
- **Culture, empathy, impedance match with academia**
 - Mutual “pick up the phone” help, unblocking research
 - Conference organization, panels, invited papers/talks ...
- **Price of admission**
 - Machines in locked cages, technology control plans, long LUL process, “bundles” in DesignStart, ... and, stay within bounds
= What one would expect from an IP company

UCSD Machine Integrated Computing & Security (MICS)

<http://mics.ucsd.edu>

Focus: Data analytics, security, machine learning, and privacy

- Paradigm shift in performance and trust through automated co-optimization of HW, SW, data and algorithms
- 18 faculty in UCSD ECE, CSE depts and Supercomputing
- With industry partners, developing foundations for the next generation machine learning and trust for emerging applications



Example recent projects and results

- ARM2GC^[1]: ARM to Garbled Circuit translation for privacy-preserving compute
- XONN^[2]: State-of-the-art in provably privacy preserving machine learning and the first to scale to very deep learning (DL) networks (e.g., 22 layers)
- DeepAttest^[3]: Automated HW-based attestation of deep learning/IP protection
- Coda^[4]: Decompiler (reverse engineering) of binary executables for DL

[1] Songhori, E.M. et al., Succinct Garbled Processor for Secure Computation. DAC'19

[2] Riazi, S. et al. XNOR-based Oblivious Deep Neural Network Inference. USENIX Sec'19

[3] Chen, H. et al. DeepAttest: An End-to-End Attestation Framework for Deep Neural Networks. ISCA'19

[4] Fu, C. et al., Coda: An End-to-End Neural Program Decompiler. NeurIPS'2019

Conclusions

- **OpenROAD aims straight at RTL-to-GDS heart attacks**
 - **Reboots EDA research** and academic-industry bridges
 - **Enables system design space exploration** to reach beneath RTL and FPGAs (e.g., in latest RTML project)
- **Open-source EDA = vehicle for “last scaling levers” (cost, schedule, quality) and machine learning for IC design**
- **Many branching points in path forward: guidance needed!**
 - Balance of needs across academic researchers, EDA developers, EDA tool users, IC design ecosystem
- **Please watch this space!**

<https://theopenroadproject.org/>

THANK YOU !

Research at UCSD is supported by U.S. National Science Foundation, U.S. DARPA, Samsung, Qualcomm, NXP Semiconductors, Mentor Graphics, and the C-DEN Center.

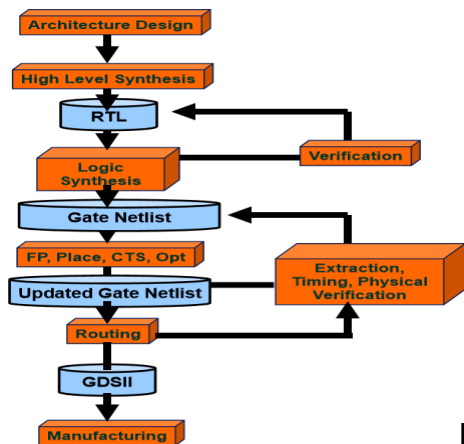
A few links:

- Presentations linked under News at <https://vlsicad.ucsd.edu/>
- <https://theopenroadproject.org/> and <https://theopenroadproject.org/outreach/>
- <https://github.com/The-OpenROAD-Project/>
- Machine learning “in and around IC design tools”:
 - <https://vlsicad.ucsd.edu/Publications/Conferences/356/c356.pdf>
 - <https://vlsicad.ucsd.edu/Publications/Conferences/360/c360.pdf>
 - Etc.

Backup

Design-Tech Co-Optimization (DTCO)


Design Technology



Design-Aware
Manufacturing


Manufacturing-Aware
Design

Manufacturing Technology

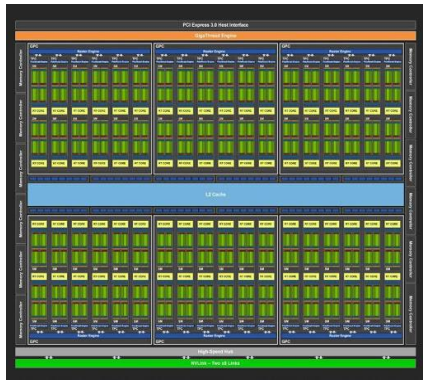


Design for Manufacturability (DFM)

Key Design Types

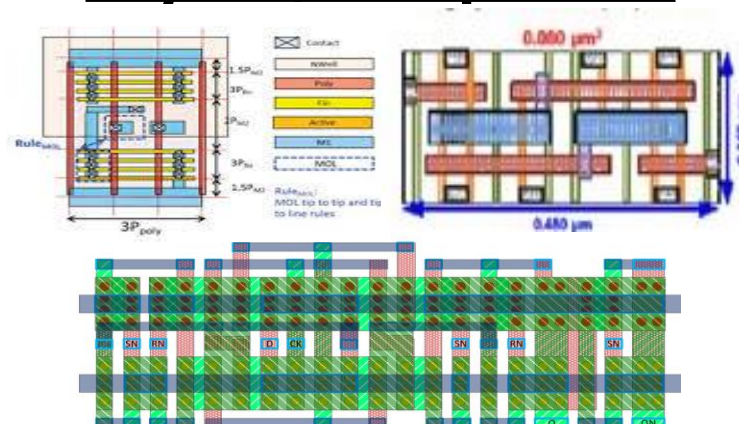


RV12 RISC-V



GPU Core


Key IPs/ Components



NAND2, 6T SRAM, D Flip-Flop

DTCO → Pathfinder

Applications, Markets

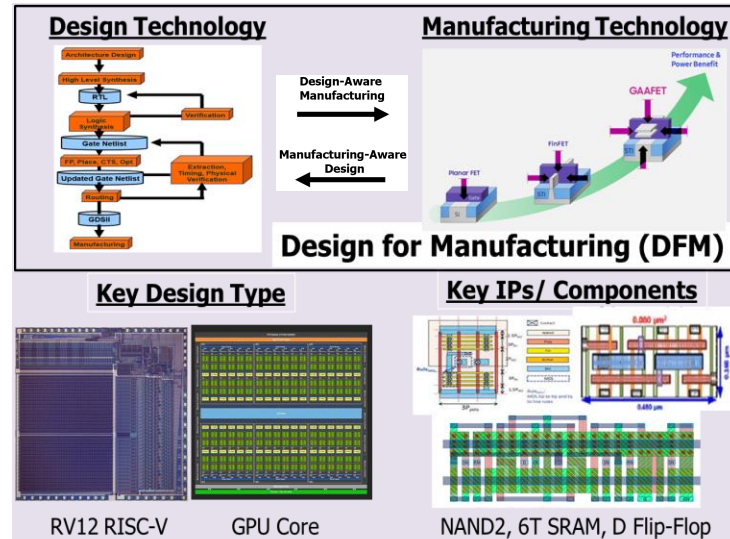


Systems

Quantum Processor
 Package On Package D2W
 D2D Monolithic
 Neuromorphic computing




Loop back is missing!!!

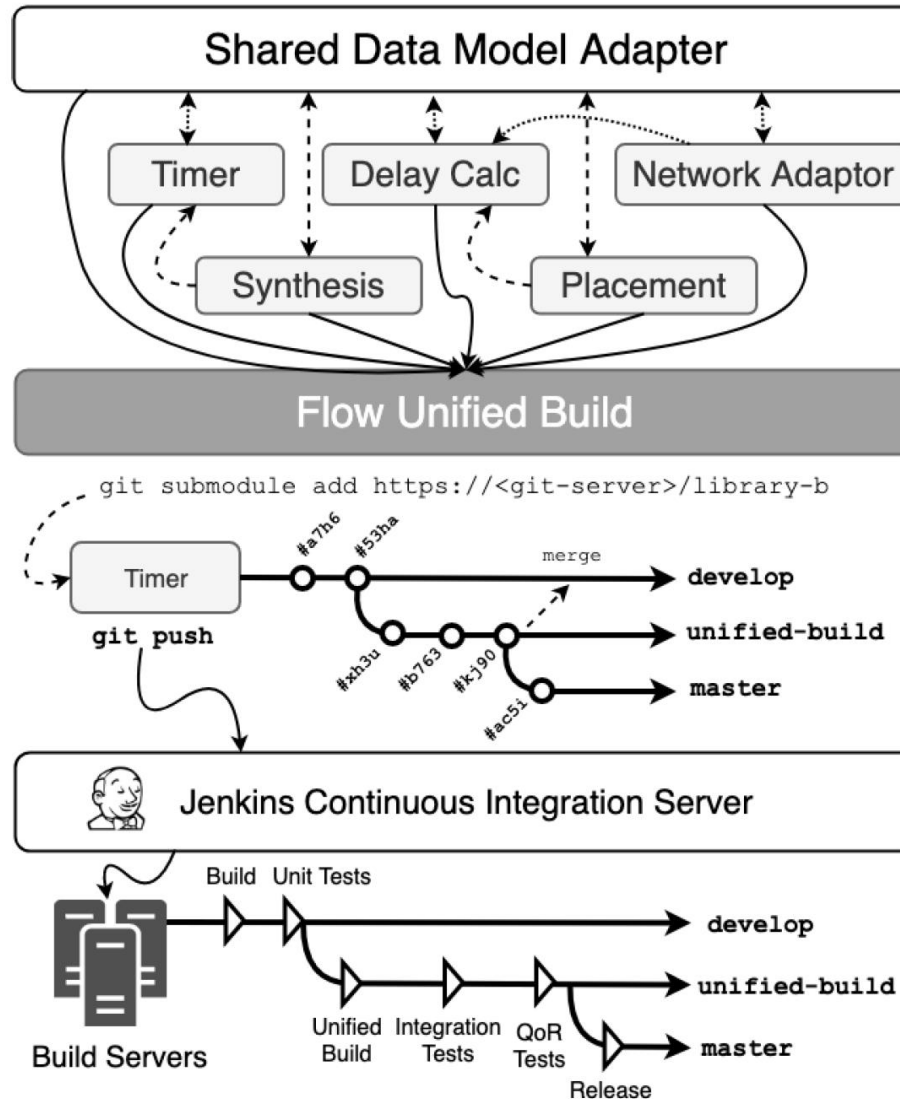


“PROBE”-based QOR Comparison

- PROBE: Start with a placement, iteratively perturb to tangle it, and record “hockey stick” of #DRCs in final routing
- Stronger placements will survive more tangling (iso-router) before the bend in the “hockey stick”
- Stronger routers will survive more tangling (iso-placer)
- OpenROAD tools currently weaker than commercial tools



Build from Sources, Continuous Integration



Warning: 40 Years of Industry Learning Curve

- Commercial EDA has gone up a learning curve regarding how to architect the complex interactions between tools in the Synthesis, Place and Route tool chain
- 1980's - focus was on point tools connected by files
- 1990's – pervasive timing driven steps required tight connection to timing
- 2000's – tightly coupled algorithms on a shared incremental substrate
- 2010's – Advanced node effects, parallel processing, **hyper-optimization**


OpenROAD Outreach: Workshops, Contests...

<https://theopenroadproject.org/outreach/>

OpenROAD

Subscribe

Home People News and Events Publications Outreach



Home / Outreach



Open Source Community
Contribution Awards
(Nominations Open)



DAC'19 Birds of a Feather



WOSET 2019



ICCAD-2019 Contest C co-
sponsorship



DAC'18 Birds of a Feather

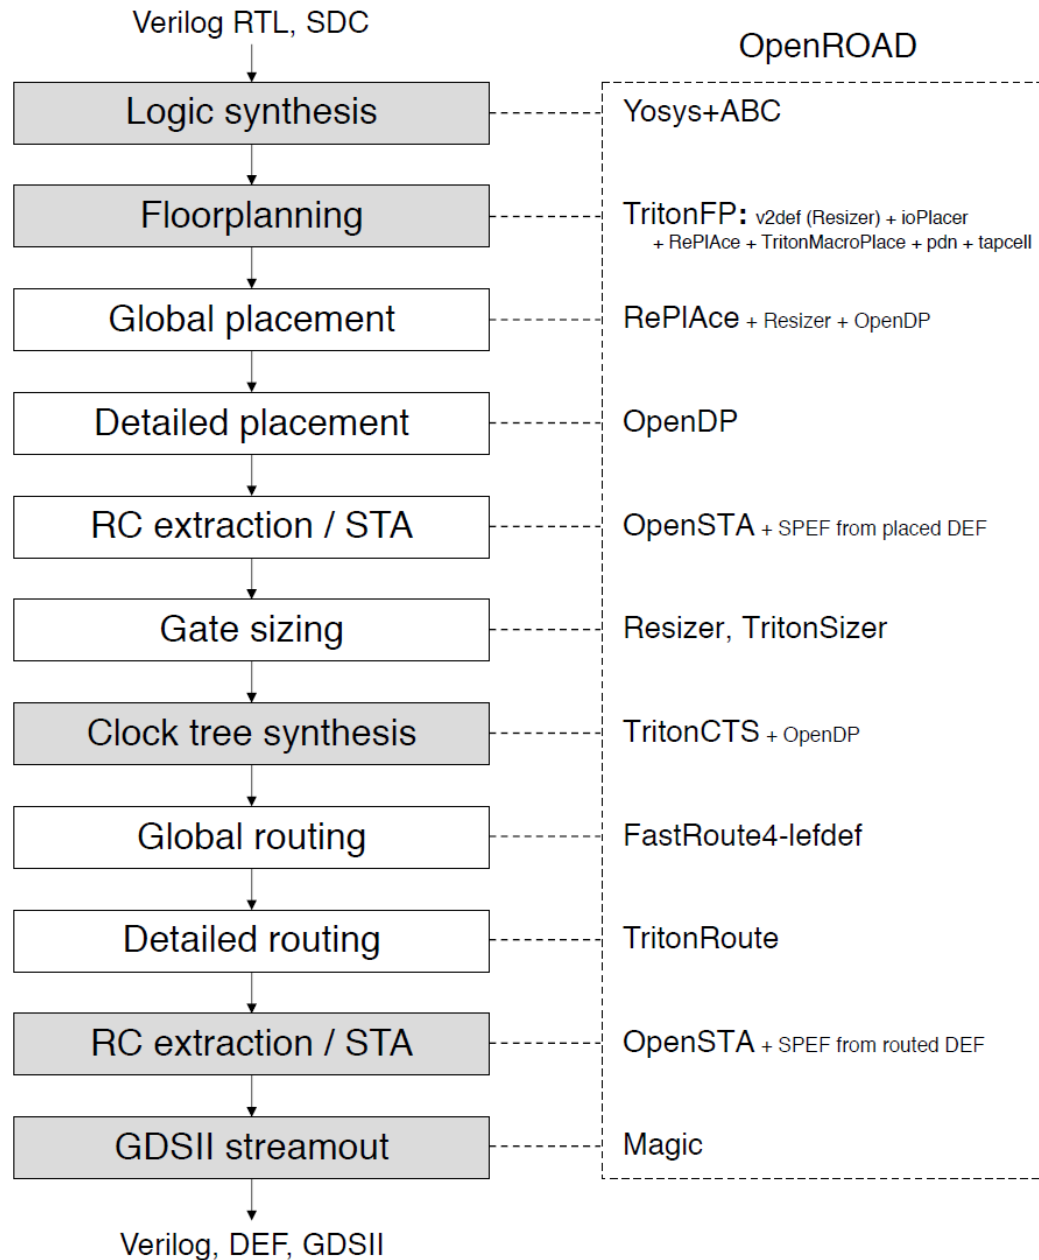


WOSET 2018



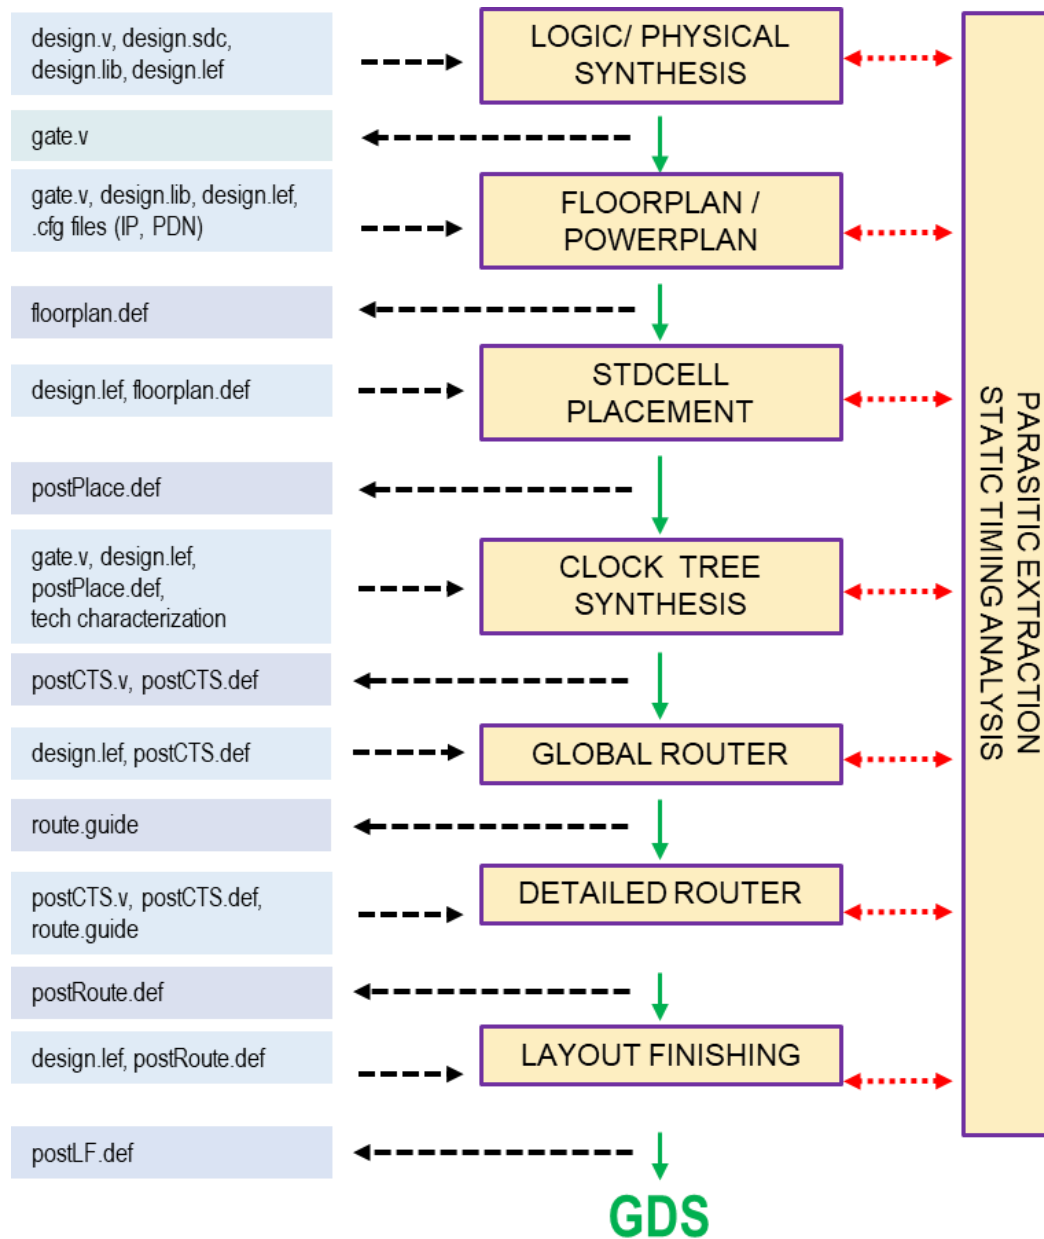
EDA Futures Workshop

RDF-2019: Merging of “Parallel Universes”



J. Chen, I. H.-R. Jiang, J. Jung et al., “DATC RDF: Towards a Complete Reference Flow”, Proc. ICCAD, 2019, to appear.

(DAC-2019 Paper)



Not an Architect, Just a Fan...

- **ORION 2.0/3.0:**
 - **Network on Chip Power and Area Model**
ORION 2.0 ([download](#))
ORION 3.0 ([website](#))
- **CACTI-IO:**
 - **Power, Area and Timing Models for The IO and PHY of Off-Chip Memory Interfaces** ([report](#))
- **CACTI 7:**
 - **New Tools for Interconnect Exploration in Innovative Off-Chip Memories** ([website](#))
- **ITRS System Drivers, System Integration roadmaps**
- **Collaborations with architects**
- ...

ORION3.0:
A Power-Performance Simulator for Interconnection Networks

[Overview](#) | [Description](#)

CACTI-IO Technical Report

Norman P. Jouppi, Andrew B. Kahng, Naveen Muralimanohar, Vaishnav Srinivas

HP Laboratories
HPL-2013-79

Keyword(s):
IO; Interconnect; SERDES; PHY; Memory bus; DDR; LPDDR

Abstract:
We describe CACTI-IO and PHY of the enables quick and describe the memory capacity for the input port for a custom configuration within 0-15% error and low-latency interconnection growing importance, NoC aid architects and designers in estimators are required. ORION er and area models, which have p networks. These models enable tem levels. The current version, state-of-the-art non-parametric

CACTI 7: New Tools for Interconnect Exploration in Innovative Off-Chip Memories

[Overview](#) | [Description](#)

[Download](#)
[People](#)
[Publications](#)
[Contact Us](#)

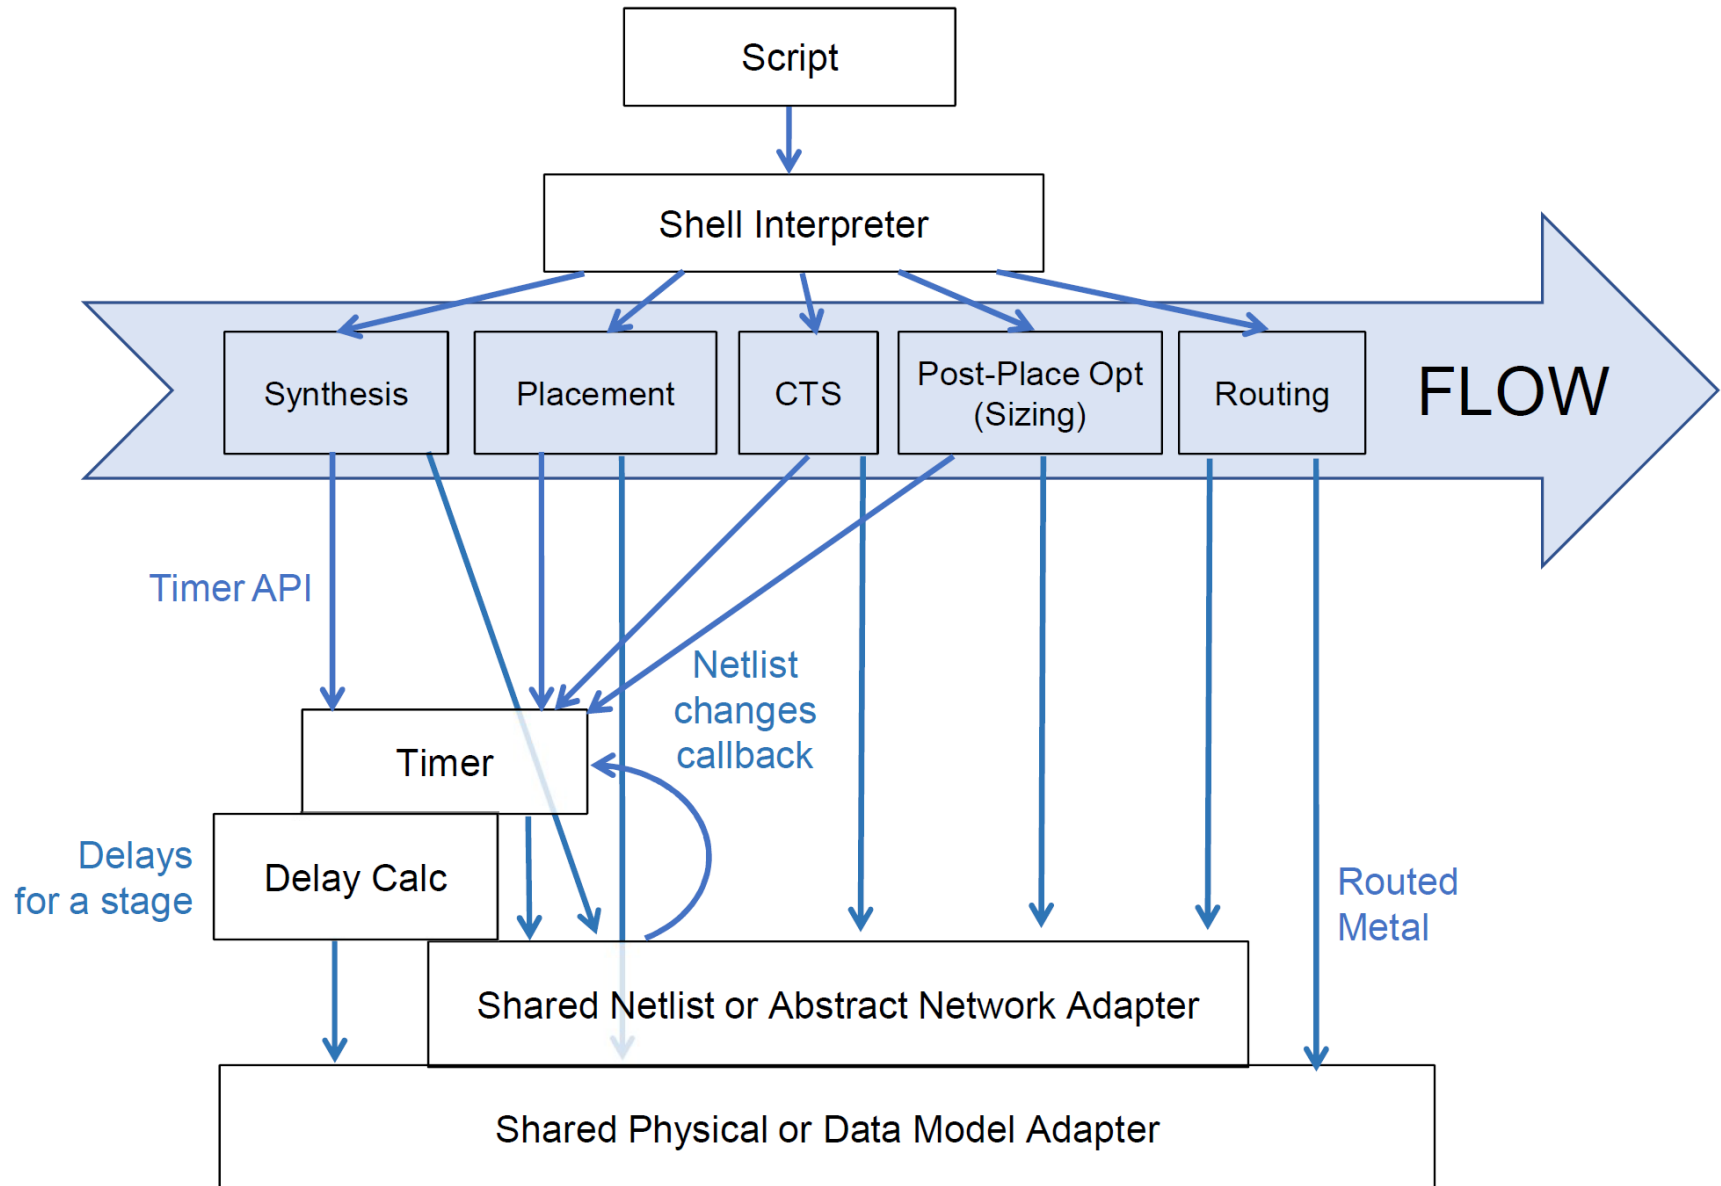
CACTI7 is an extended version of CACTI that includes power, area and timing models for I/O [2]. CACTI7 allows for a quick search of optimal IO configuration parameters that help optimize power and performance of the IO along with the DRAM and cache subsystem. CACTI has analytical models for all the basic building blocks of a memory: decoder, sense-amplifier, crossbar, on-chip wires, DRAM/DRAM cell, and latch. We extend it to include the OFF-chip models. This requires modifying CACTI's global on-chip interconnect to include buffers at the PHY and drivers at the bank edge to connect to the IO circuit. Since all calculations are based on the ITRS technology parameters, the energy and delay values calculated by CACTI are guaranteed to be mutually consistent. When a user inputs memory parameters and energy/delay constraints into CACTI, the tool performs an exhaustive design space exploration involving different array sizes, degrees of multiplexing, and interconnect choices to identify an optimal configuration. CACTI7 is capable of performing an additional search for OFF-chip parameters, including optimal number of ranks, memory data width (x4, x8, x16, or x32 DRAMs), OFF-chip bus frequency, and bus width. This allows for optimal tradeoffs between OFF-chip power, area, and timing. For further details on the off-chip I/O models used in CACTI7, please refer to [CACTI-IO tech report](#).

Leveraging Thermally-Aware Chiplet Organization in 2.5D Systems to Reclaim Dark Silicon
Furkan Eris¹, Ajay Joshi¹, Andrew B. Kahng^{2,3}, Yenai Ma¹, Saiful Mojumder¹ and Tiansheng Zhang¹
¹ECE Department, Boston University, Boston, MA, USA; ²ECE and ³CSE Departments, UC San Diego, La Jolla, CA, USA
{fe, joshi, yenai, msam, tszhang}@bu.edu, abk@cs.ucsd.edu

A Cross-Layer Methodology for Design and Optimization of Networks in 2.5D Systems
Ayse K. Coskun¹, Furkan Eris¹, Ajay Joshi¹, Andrew B. Kahng^{2,3}, Yenai Ma¹, and Vaishnav Srinivas²
¹ECE Department, Boston University, Boston, MA, USA; ²ECE and ³CSE Departments, UC San Diego, La Jolla, CA, USA
acoskun@bu.edu, fe@bu.edu, joshi@bu.edu, yenai@bu.edu, abk@eng.ucsd.edu, vaishnav@ucsd.edu



Cross-layer Floorplan Optimization For S... In Many-core Systems
Ayse K. Coskun¹, Anjun Gu¹, Warren Jin⁴, Ajay Joshi³, Andrew B. Kahng^{1,2}, Jonathan Klamkin⁴, Yenai Ma³, John Recchio¹, Vaishnav Srinivas¹ and Tiansheng Zhang³
UCSD ¹ECE and ²CSE Departments, La Jolla, CA; ³Boston University ECE Department, Boston, MA; ⁴UCSB ECE Department, Santa Barbara, CA

Incremental Shared Netlist Architecture



A Root of Evil: Unpredictability

- Intractable optimizations → heuristics piled on heuristics
- **“Noise” or “Chaos” when EDA tools “try hard”**
- **Unpredictability → added margin and schedule**
14nm PULPino: $\Delta\text{area} = 6\%$ from $\Delta\text{freq} = 10\text{MHz}$!



CRISES of Schedule, Quality, Cost – and RISK


ML in IC Design: Not Like Chess or Cat Photos

- **Getting to self-driving IC design: not so obvious**
 - Do recent ML successes transfer well?
 - 3-week SP&R&Opt run is NOT like playing chess!
- **Design lives in a {servers, licenses, schedule} box**
- **Distributions** of outcomes matter cloud, parallel
- A “stack of models” is mandatory: Predictions of downstream outcomes are also optimization objectives
- **Still uncharted road to self-driving tools and flows**
 - How do we overcome “small, expensive data” challenges?
 - **Standards:** Learning comes from {design + tool + technology}, all of which are highly proprietary
 - **Need mechanisms for IP-preserving sharing of data and models**

4 Stages of ML to Recover Time, Effort

Four Stages of Machine Learning

1. Mechanization and Automation
2. Orchestration of Search and Optimization
3. Pruning via Predictors and Models
4. From Reinforcement Learning through Intelligence



Huge space of tool, command, option trajectories through design flow

ML in IC Design Requires Infrastructure !

- **Support for ML in IC design**
 - Standards for model encapsulation, model application, and IP preservation when models are shared
- **Standard ML platform for EDA modeling**
 - Design metrics collection, (design-specific) modeling, prediction of tool/flow outcomes
 - This recalls “METRICS” <http://vlsicad.ucsd.edu/GSRC/metrics>
- **Datasets to support ML**
 - Real designs, Artificial designs and “Eycharts”
 - Shared training data – e.g., analysis correlation, post-route DRV prediction, sizer move trajectories and outcomes, ...
 - Challenges and incentives: “Kaggle for ML in IC design”



The METRICS Initiative

Recent Updates

- [Survey \(1st draft\)](#) for design quality and productivity ([the multiple-choice version](#))
- [Reduced survey \(2nd draft\)](#) for design quality and productivity that is distributed at June 2001 GSRC workshop
- [Updated survey \(3rd draft\)](#) for design quality and productivity that reflects the discussion at June 2001 GSRC workshop
- [Workshop notes](#) for METRICS discussion at June 2001 GSRC workshop
- [List of prediction/estimator models](#) enabled by METRICS System 新
- [DAC02 Birds-of-a-Feather meeting summary](#) (June 12, 2002) 新

- METRICS (1999; DAC00, ISQED01):

“Measure to Improve”

- Goal #1: Predict outcome
- Goal #2: Find sweet spot (field of use) of tool, flow
- Goal #3: Dial in design-specific tool, flow knobs

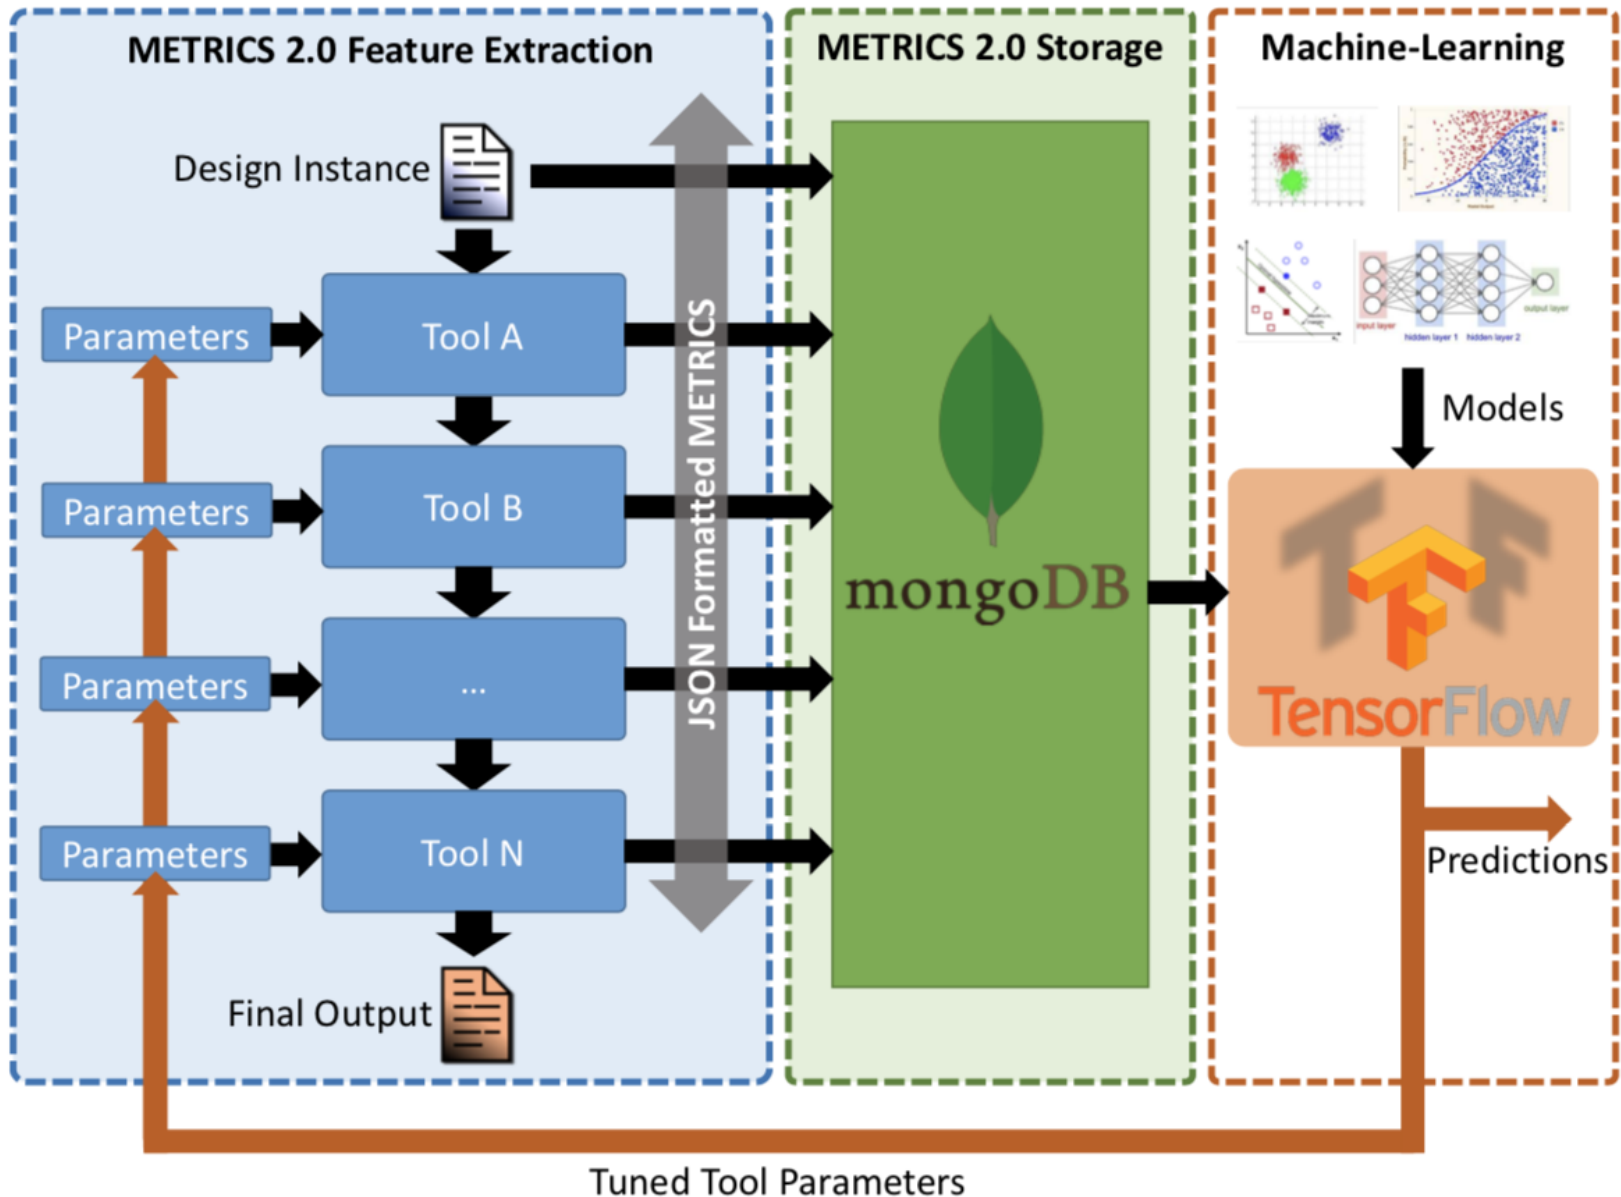
<http://vlsicad.ucsd.edu/GSRC/metrics>

METRICS 2.0 Evolution Path

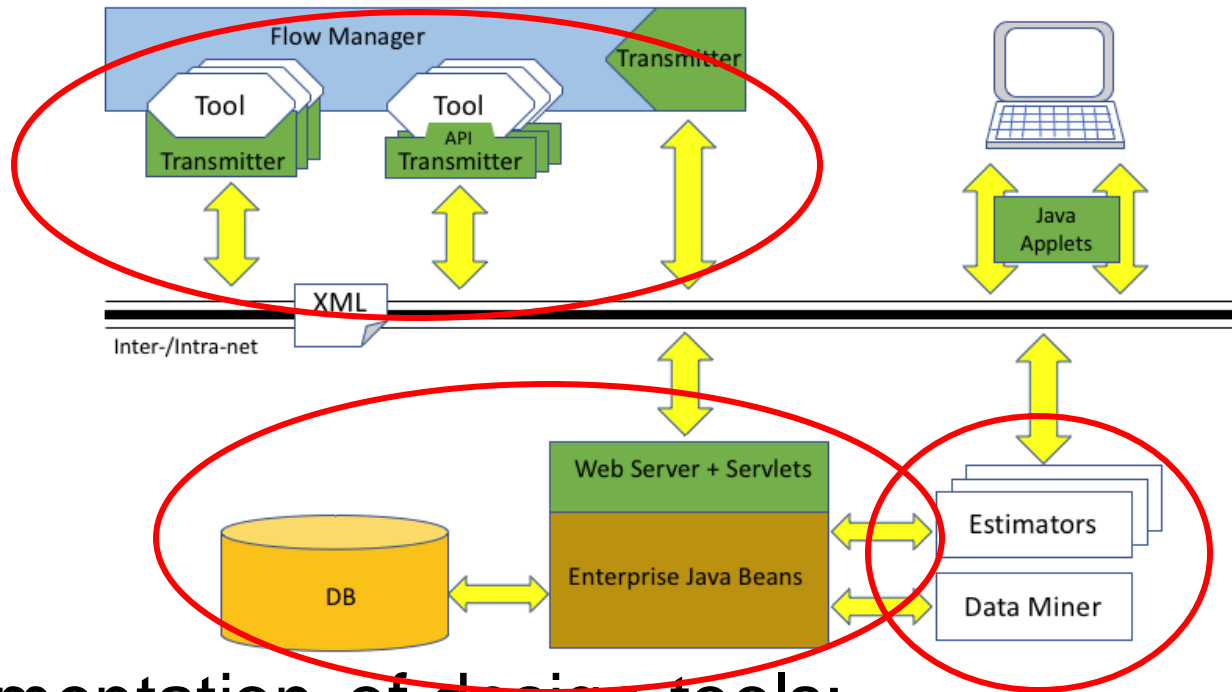
METRICS2.0 metric	Definition	Flow Stage	Collectable from tool reports?	Derivable from tool run data?
SYNTHESIS::INST::COMB	Total #Comb Instances	Synth	Y	Y
SYNTHESIS::INST::REG	Total #Registers	Synth	Y	Y
PLACEMENT::INSTANCES::TOTAL	Total #Instances	Place	Y	Y
PLACEMENT::PIN_DENSITY	Ratio of Total #Pins / Total Area	Place	N	Y
PLACEMENT_OPT::POWER::TOTAL	Total power after preCTS Opt	Place	Y	Y
CLOCK::INST	Total #Clock Insts	CTS	Y	Y
CLOCK::TIMING::SETUP::TNS	Total TNS of timing paths	CTS	Y	Y
CLOCK::TIMING::SKEW::AVG	Average clock skew	CTS	Y	Y
ROUTING::DRC::TOTAL	Total DRC Count	Routing	Y	Y
ROUTING::VIA::TOTAL	Total #Vias	Routing	Y	Y
ROUTING:WIRELENGTH::TOTAL	Total Wirelength	Routing	Y	Y

- METRICS 2.0 → learning from recent runs → adaptive flows → reinforcement learning
- METRICS 2.0 + Grid Computing = shared burden of “big data”
- METRICS 2.0 + Federated ML = privacy-preserving models

Proposed METRICS 2.0 Architecture



Original METRICS Architecture



- Instrumentation of design tools:
 - Wrapper scripts to extract data from outputs and logfiles,
 - Callable API codes that allow direct interaction from within the design tools
- METRICS server: central data collection (Oracle8i)
- Data mining process: analyzes existing data to improve existing design flow (CUBIST, etc.)

Why Open Source has Strategic Value Now

- **Extreme consolidation:** 2 fabs, 2 EDA companies, at most 3 standing in GPU, FPGA, Mobile SOC, ...
- **Design technology is higher-value**
- **Private-label EDA is inevitable**
 - **One of few levers for differentiation**
- **Foundry:** protect IP leakage, roadmap better than current DTCO mechanisms
- **Fabless:** protect key SOC methodology, design innovation that otherwise leaks to competitors
- **EDA:** more useful research, in more usable forms
- **Academia:** see previous slides (and, engage ... or not)

