

# Challenges with Basilisk and where the Future is going

Integrated Systems Laboratory (ETH Zürich)

|                           |                         |
|---------------------------|-------------------------|
| <b>Philippe Sauter</b>    | phsauter@iis.ee.ethz.ch |
| <b>Thomas Benz</b>        | tbenz@iis.ee.ethz.ch    |
| <b>Paul Scheffler</b>     | paulsc@iis.ee.ethz.ch   |
| <b>Frank K. Gürkaynak</b> | kgf@iis.ee.ethz.ch      |
| <b>Luca Benini</b>        | lbenini@iis.ee.ethz.ch  |

## **PULP Platform**

Open Source Hardware, the way it should be!



@pulp\_platform



pulp-platform.org



youtube.com/pulp\_platform



# The PULP Open-Source Ecosystem



## RISC-V Cores and Vector Units

|                       |                       |        |       |                       |     |
|-----------------------|-----------------------|--------|-------|-----------------------|-----|
| RI5CY<br><i>CV32E</i> | Zero R<br><i>lbex</i> | Snitch | Spatz | Ariane<br><i>CVA6</i> | ARA |
| RV32                  | RV32                  | RV32   | RVV   | RV64                  | RVV |

## Peripherals

|      |      |
|------|------|
| JTAG | SPI  |
| UART | I2S  |
| DMA  | GPIO |

## Interconnects

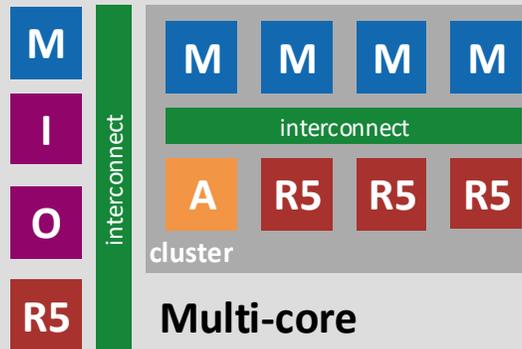
|      |         |
|------|---------|
| LIC  | HCI     |
| APB  | FlooNoC |
| AXI4 |         |

## Platforms



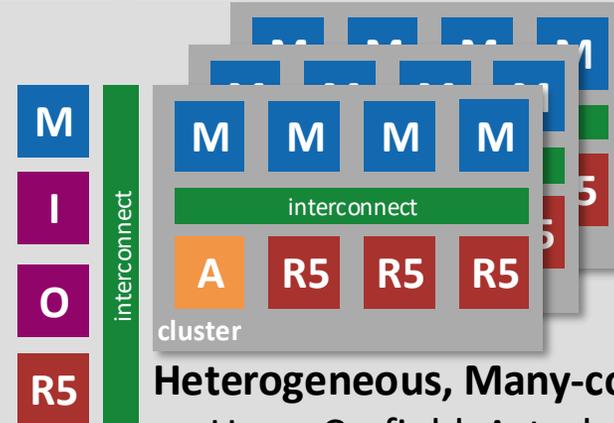
### Single core

- PULPino, PULPissimo
- **Cheshire**



### Multi-core

- OpenPULP
- ControlPULP



### Heterogeneous, Many-core

- Hero, Carfield, Astral
- Occamy, Mempoool

IOT

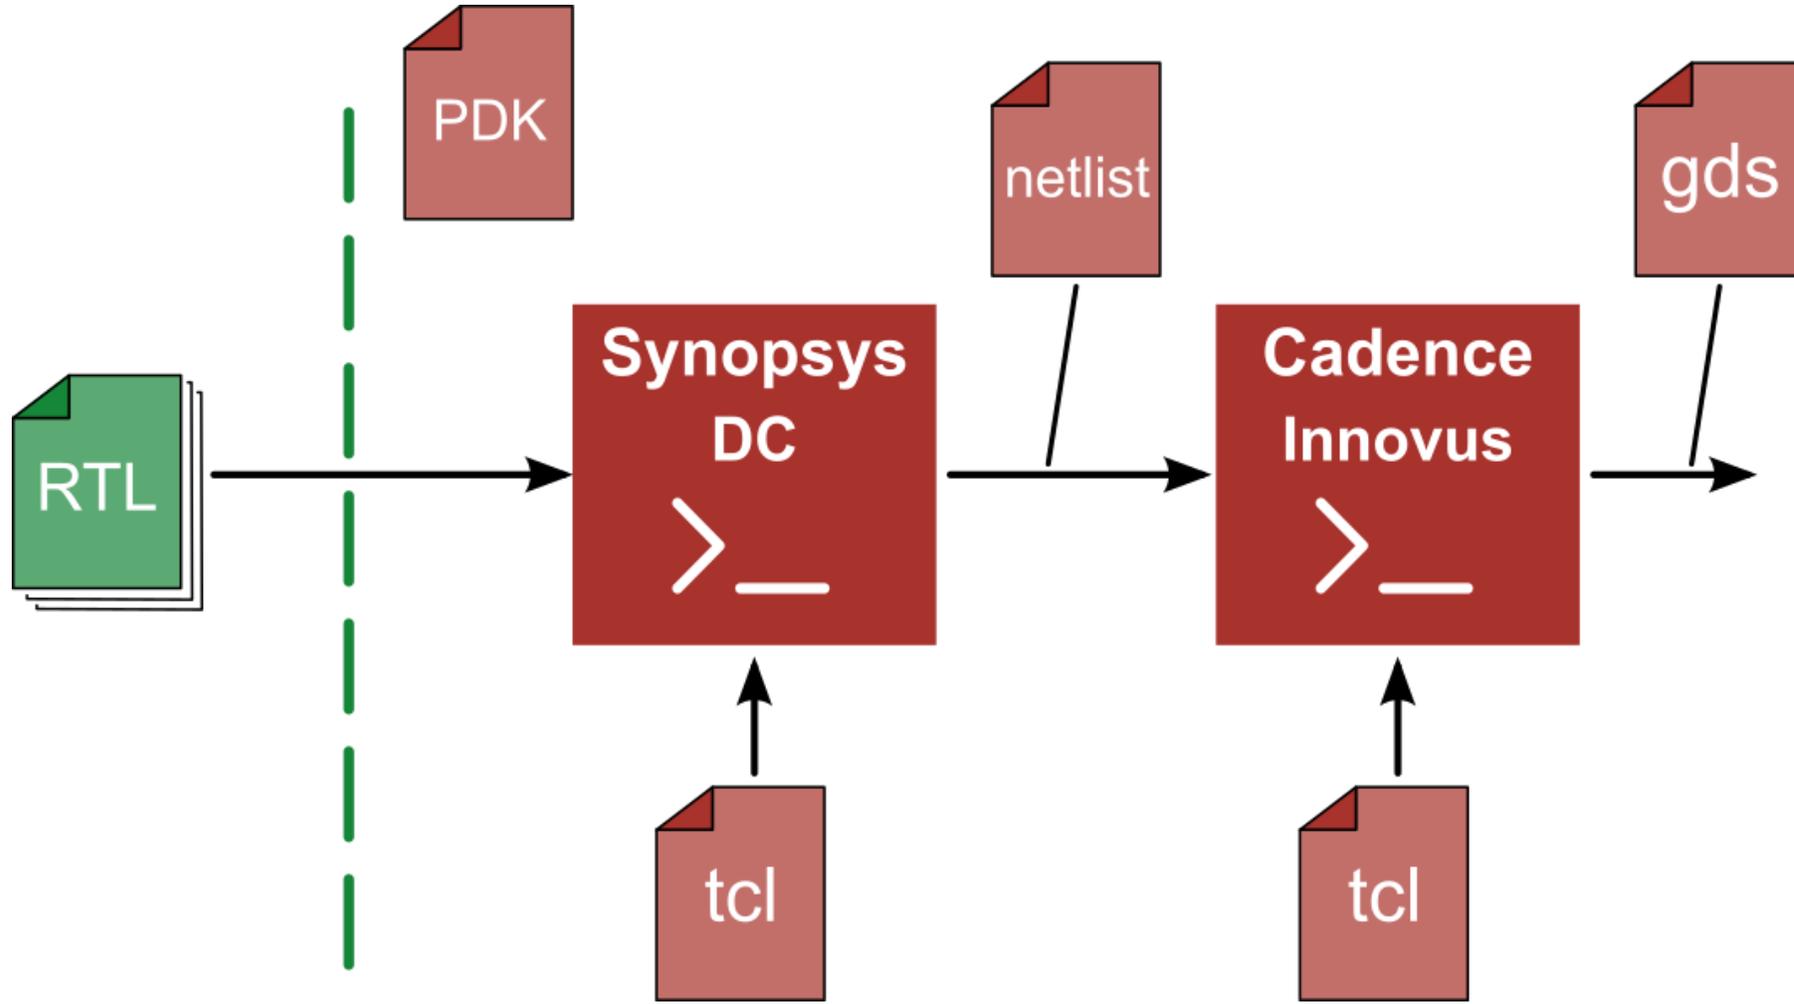
HPC

## Accelerators and ISA extensions

|                      |                       |                        |              |                        |
|----------------------|-----------------------|------------------------|--------------|------------------------|
| XpulpNN,<br>XpulpTNN | ITA<br>(Transformers) | RBE, NEUREKA<br>(QNNs) | FFT<br>(DSP) | REDMULE<br>(FP-Tensor) |
|----------------------|-----------------------|------------------------|--------------|------------------------|



# The Year is 2023, PULP is Open-Source



# Why are we Exploring more Openness?



- **Research**

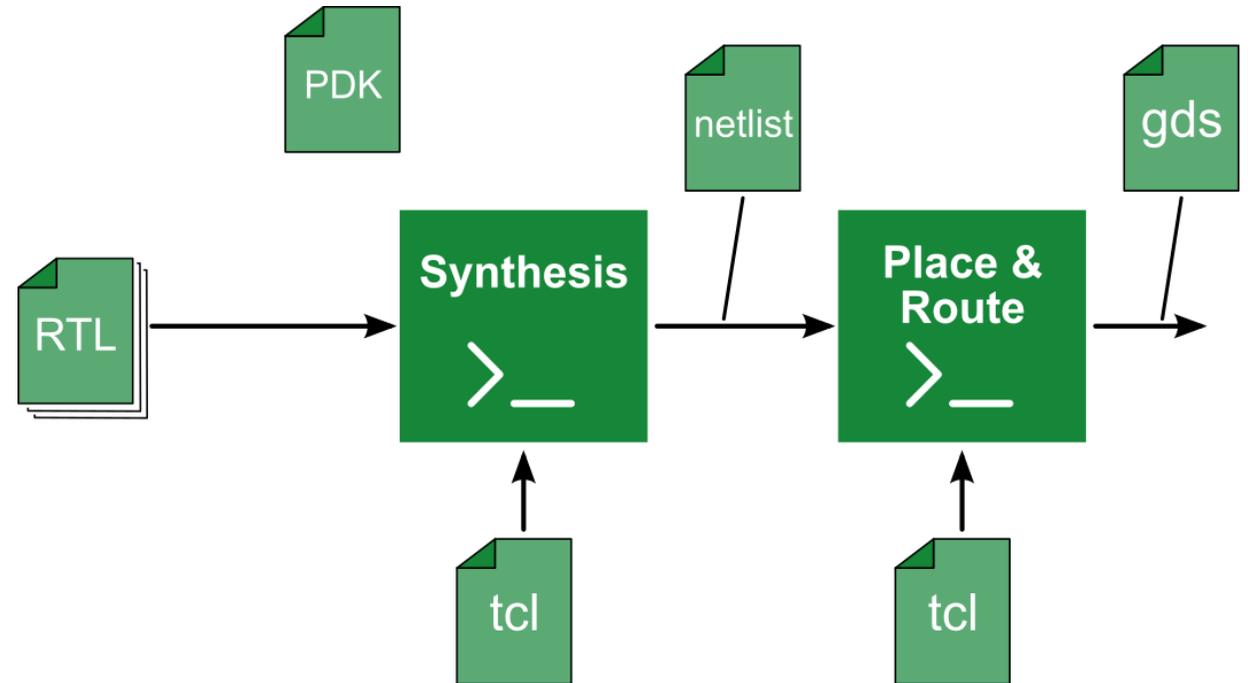
- Easier collaboration
- Reproducing results
- New research (using tools or data)

- **Education**

- Increased access
- Experiment with flows and tools
- A look "under the hood"

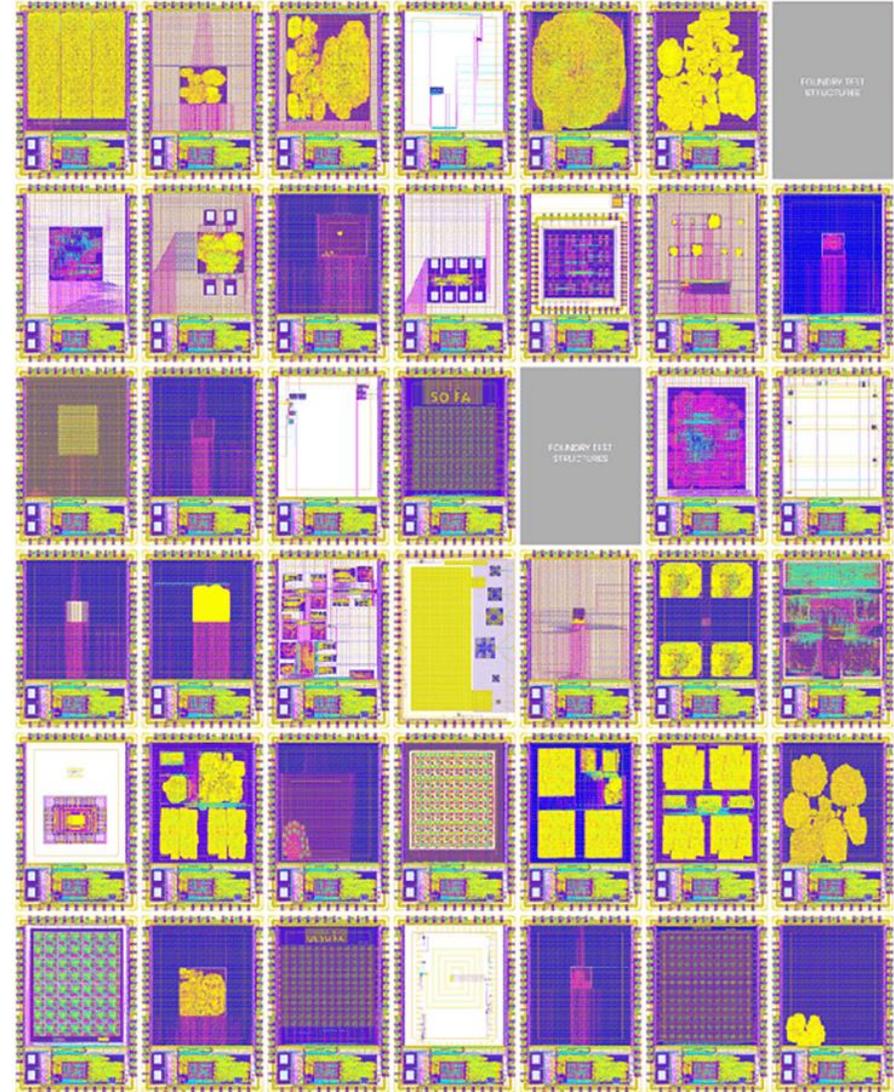
- **Industry**

- Transparent Chain-of-Trust
- Lower Cost



# What Have Others Done?

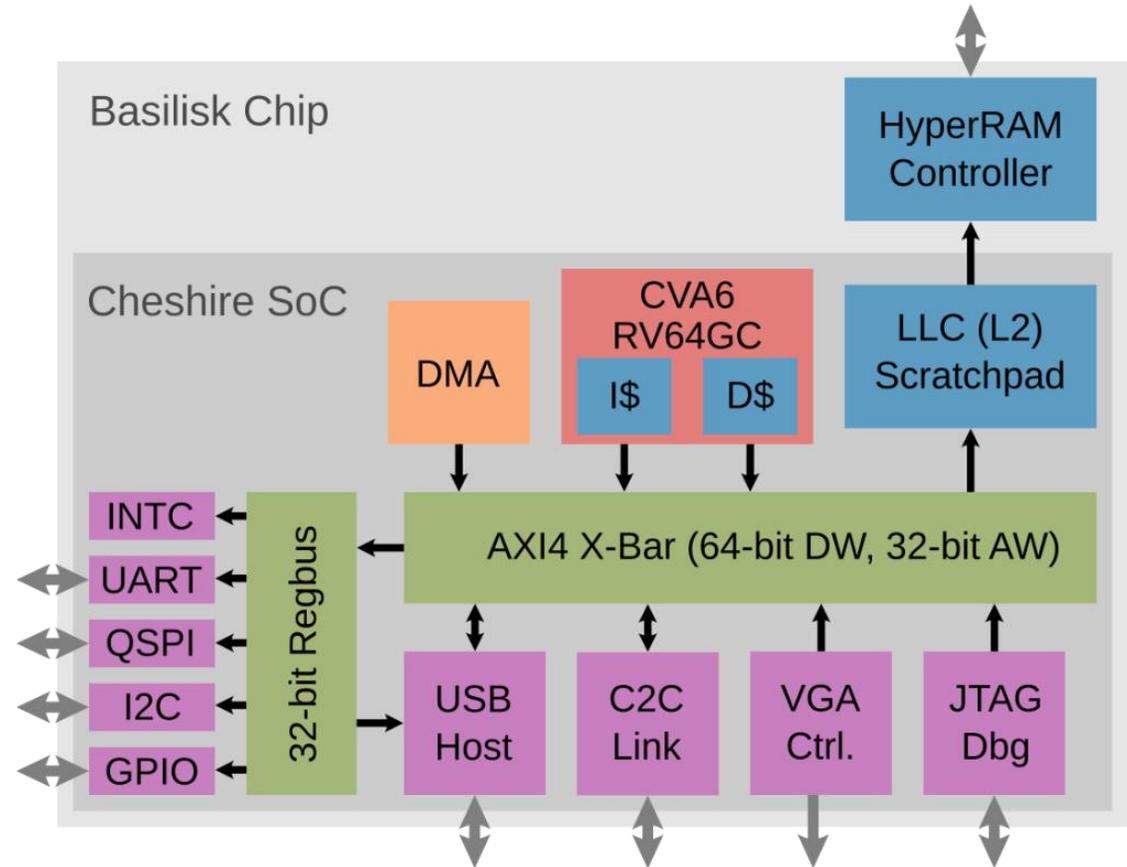
- **Over 600 tapeouts**
  - Caravel Multi Project Wafers
  - Most <50kGE, <10% utilization
  - Big ones 200kGE
- **10x smaller than Cheshire**
- **Claim: Works “out of the box”**
  - True for existing designs but...
- **We want to go further!**
  - Build on what is already there



# Find the Limits with our Cheshire SoC Platform



- **Multi-million gate design**
- **64-bit RISC-V Core**
  - Complete Linux-capable SoC
  - Simple “Raspberry Pi”
- **Rich Peripherals**
  - Including an open USB 1.1 host
- **Open-source DRAM interface**
  - Double data rate (DDR)
- **Silicon-proven**
  - Multiple successful tapeouts with closed-source commercial EDA

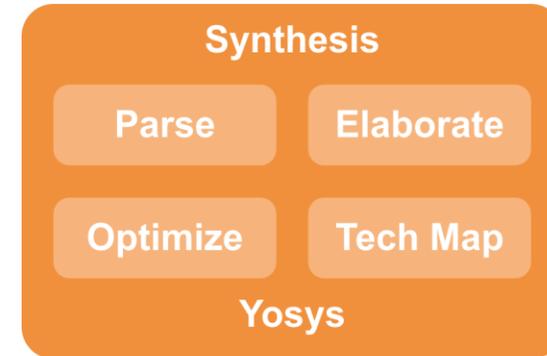


**Ideal benchmark design to stress OS EDA**



# The Basic Battle Plan

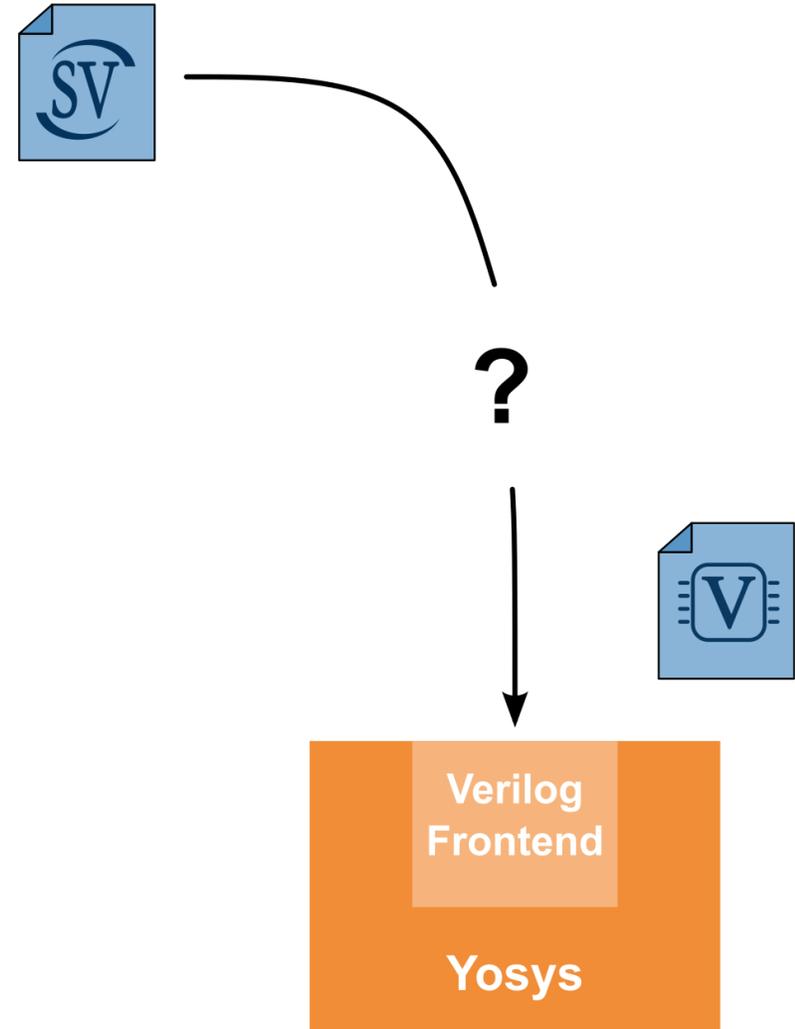
- Get **PULP IPs into the synthesis tool**
  - If needed, simplify SystemVerilog
- **Synthesize the design**
  - Yosys takes it from RTL to generic-cells
  - Calls ABC for logic optimization and mapping to the technology (IHP 130nm)
- **Implement the designs backend**
  - Collection of research tools into OpenRoad
  - Places the cells and routes wires between them
- *Lets look at a few challenges we encountered*



# First Challenge: How do we Read the Code?



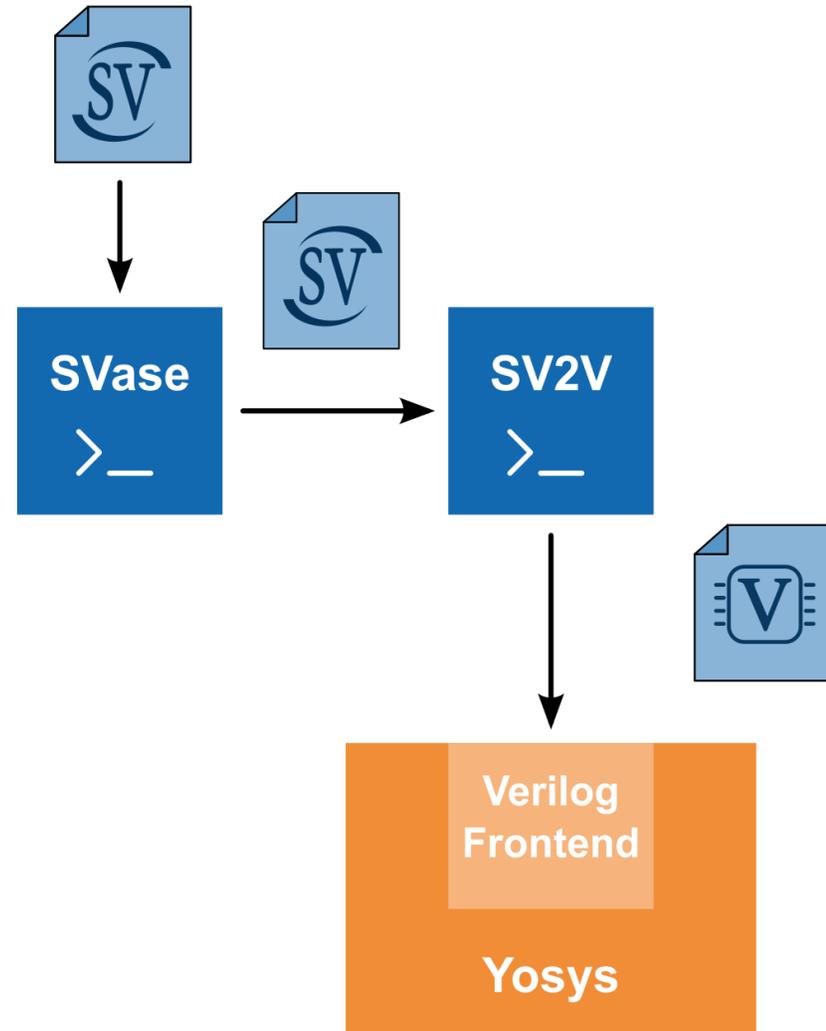
- **SystemVerilog is widespread**
  - **PULP-Platform: Cheshire**
  - lowRISC: OpenTitan and its IPs
  - OpenHW Group: CVA6, CV32E40P
  - BlackParrot RISC-V core
  - Most **industrial IPs** are implemented in SV



# First Challenge: How do we Read the Code?



- **SystemVerilog is widespread**
  - **PULP-Platform: Cheshire**
  - lowRISC: OpenTitan and its IPs
  - OpenHW Group: CVA6, CV32E40P
  - BlackParrot RISC-V core
  - Most **industrial IPs** are implemented in SV
- **For Basilisk:**
  - Pre-process the RTL
  - SVase simplifies SystemVerilog
  - SV2V converts to Verilog



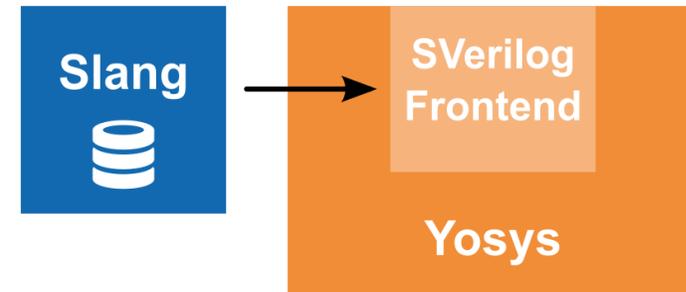
# First Challenge: How do we Read the Code?



- **SystemVerilog is widespread**

- **PULP-Platform: Cheshire**
- lowRISC: OpenTitan and its IPs
- OpenHW Group: CVA6, CV32E40P
- BlackParrot RISC-V core

For Reading SystemVerilog with  
one command using  
**yosys-slang**



# Next Challenge: We need better Synthesis Results

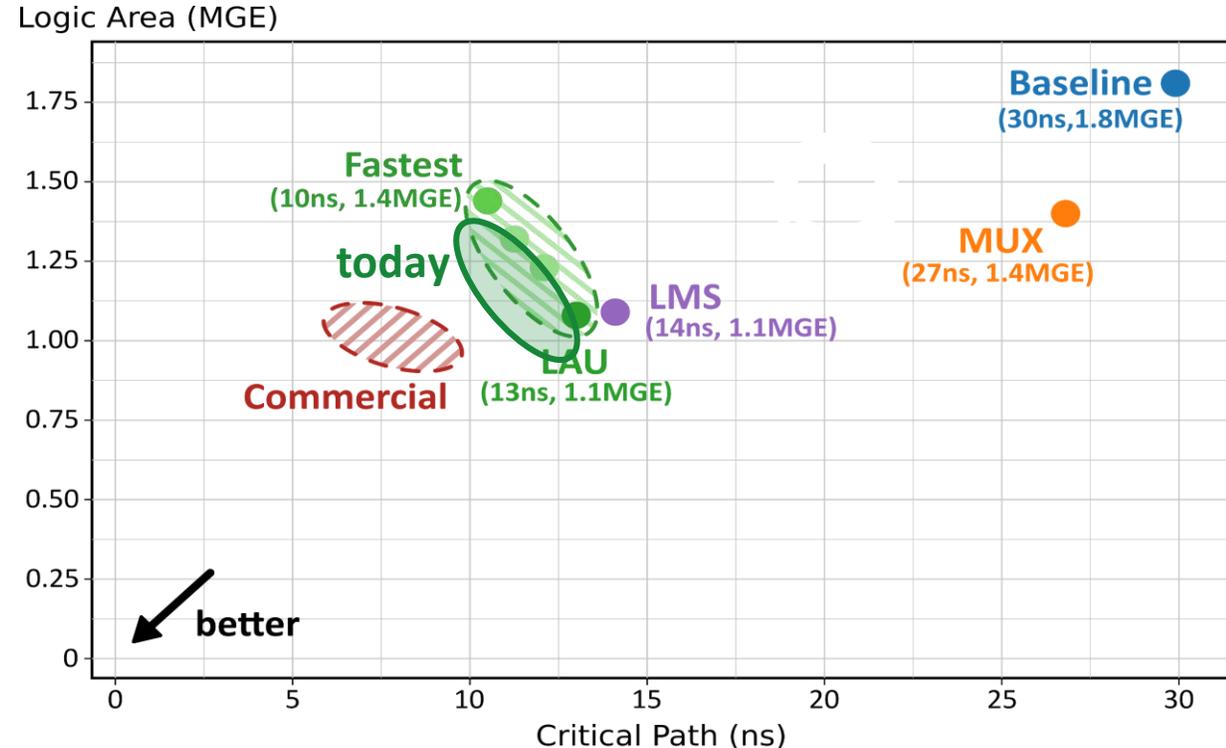


## • Improvements until June 2024

- SV-to-Verilog chain @ <2min runtime
- Yosys synthesis:
  - 1.1 MGE (1.6x) @ 77 MHz (2.3x)
  - 1.4x less runtime, 2.4x less peak RAM
- OpenROAD P&R: tuning
  - -12% die area, +10% core utilization

## • Improvements June-October

- Yosys-slang replaces SV2V
  - 1.6x less runtime, 10x less peak RAM
- -10% logic area

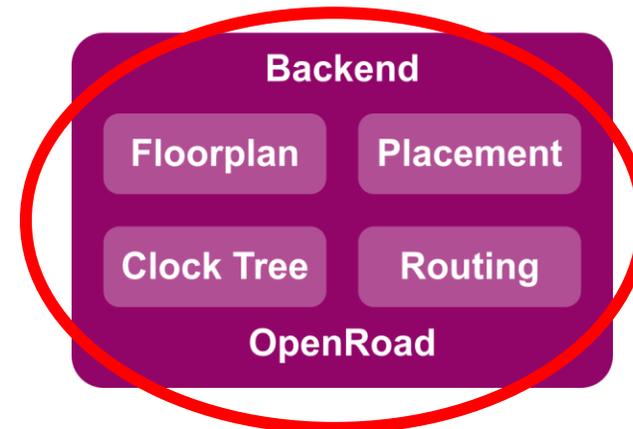
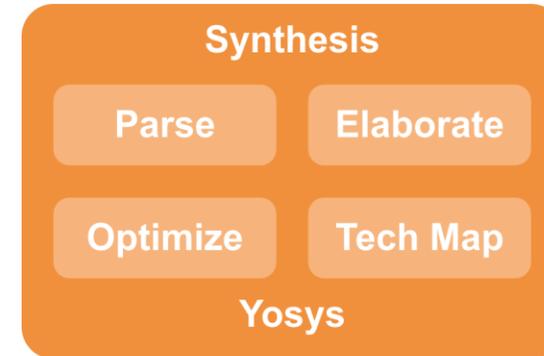


[github.com/pulp-platform/cheshire-ihp130-o](https://github.com/pulp-platform/cheshire-ihp130-o)



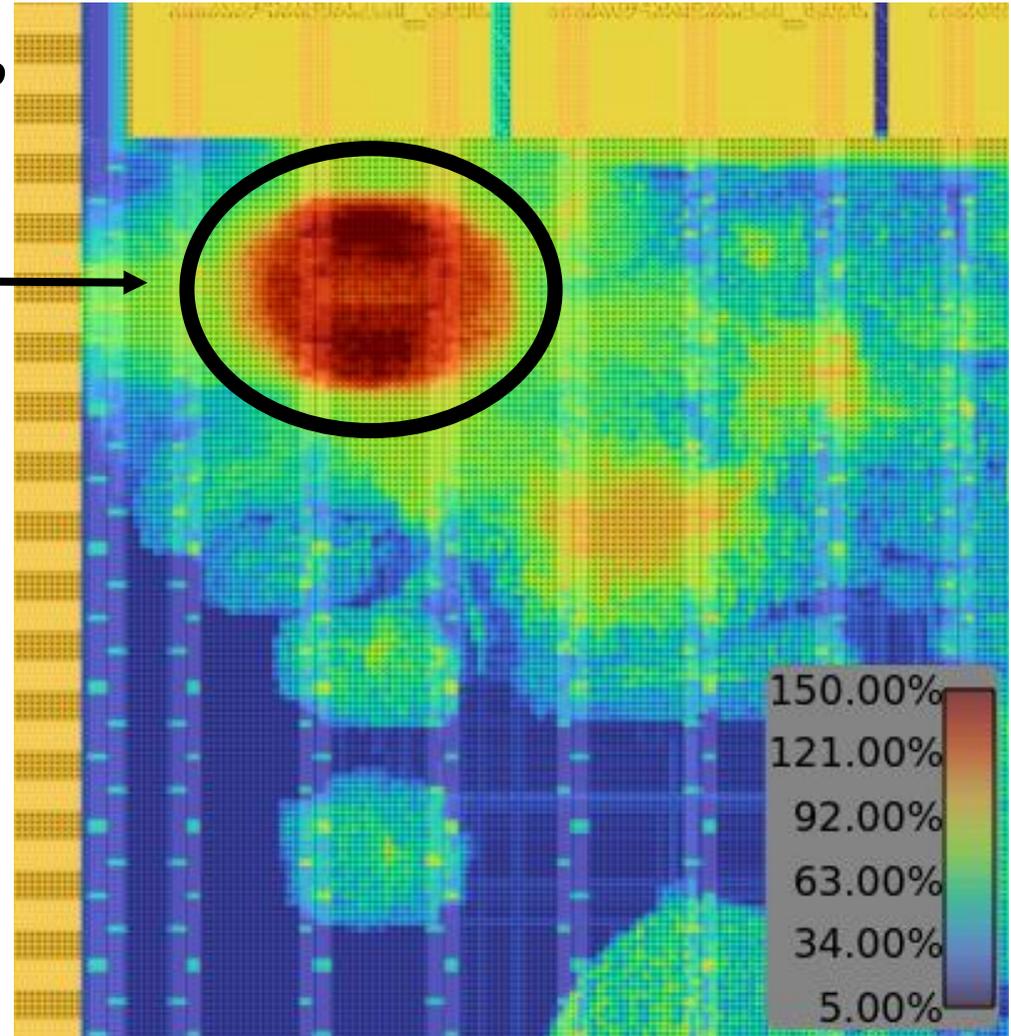
# The Basic Battle Plan

- ~~Get PULP IPs into the synthesis tool~~
  - ~~If needed, simplify SystemVerilog~~
- ~~Synthesize the design~~
  - ~~Yosys takes it from RTL to generic cells~~
  - ~~Calls ABC for logic optimization and mapping to the technology (HP 130nm)~~
- **Implement the designs backend**
  - Collection of research tools into OpenRoad
  - Places the cells and routes wires between them
- *Lets look at a few challenges we encountered*



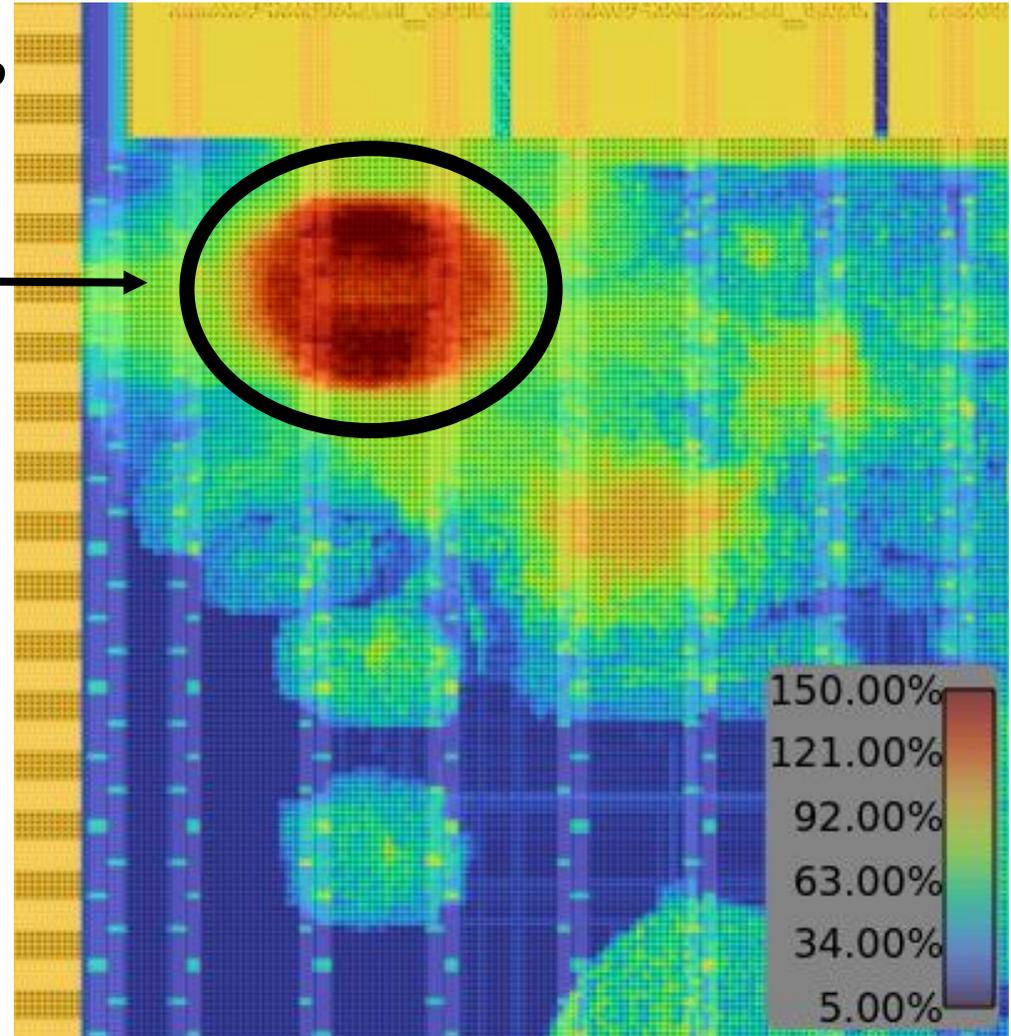
# Challenge: Cell Placement

- *Remember the Routing Congestion Task?*
- **High Congestion in Bootrom** →



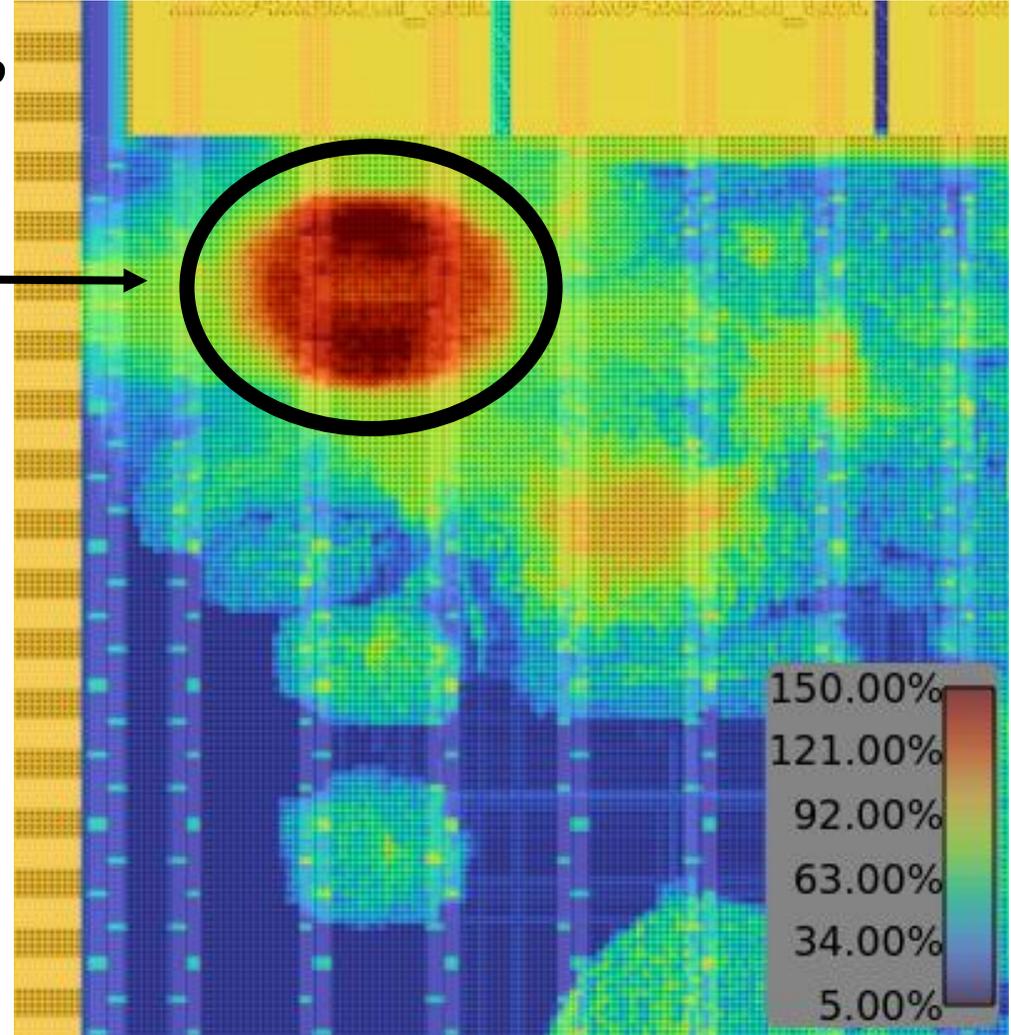
# Challenge: Cell Placement

- *Remember the Routing Congestion Task?*
- **High Congestion in Bootrom** →
- **Decrease local density**
  - More routing resources per cell
  - Decreases congestion



# Challenge: Cell Placement

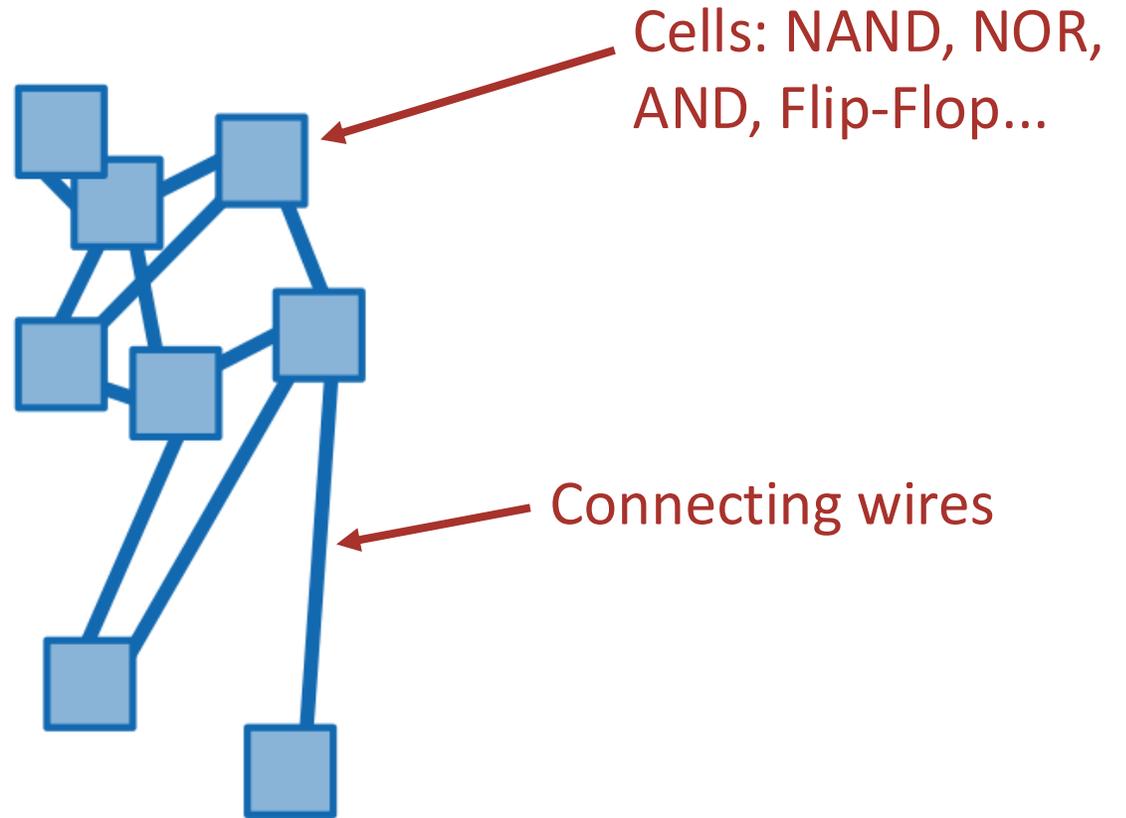
- *Remember the Routing Congestion Task?*
- **High Congestion in Bootrom** →
- **Decrease local density**
  - More routing resources per cell
  - Decreases congestion
- **OpenROAD**
  - Doesn't lower density enough
- *Can we help OpenRoad?*



# Attracting Wires, Repelling Cells

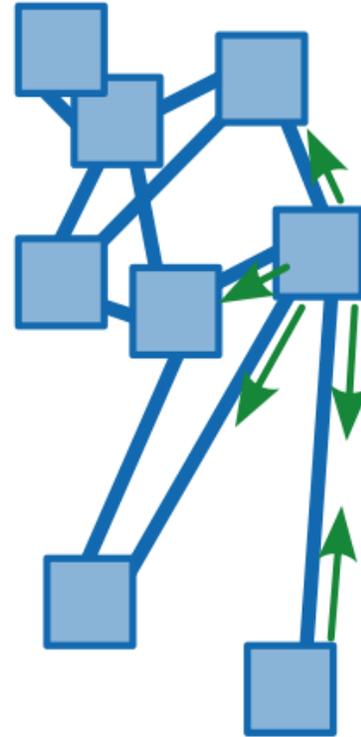


- **Goal:**
  - Minimize wire-length
  - Achieve target density



# Attracting Wires, Repelling Cells

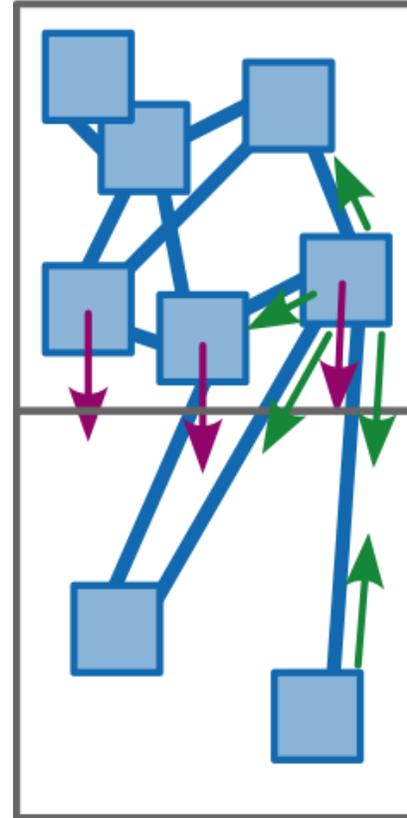
- **Goal:**
  - Minimize wire-length
  - Achieve target density
- **Wire-force attracts** connected cells



# Attracting Wires, Repelling Cells

- **Goal:**
  - Minimize wire-length
  - Achieve target density
- **Wire-force attracts** connected cells
- **Density-force push** overflowing cells away

target: 4  
actual: 6 **over: 2**



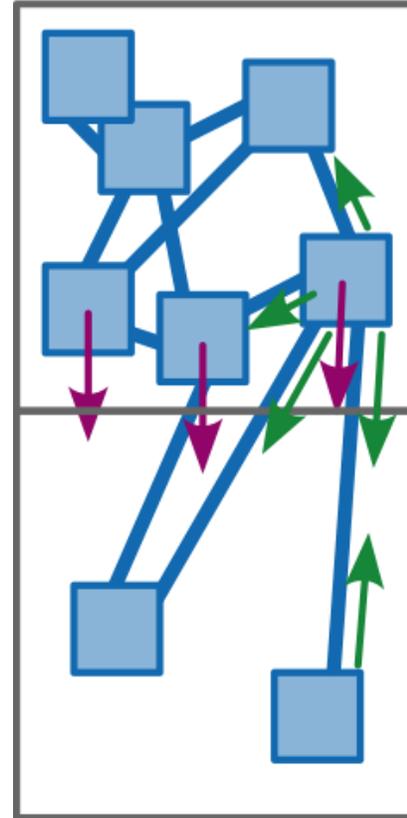
target: 4  
actual: 2 **over: 0**

# Attracting Wires, Repelling Cells

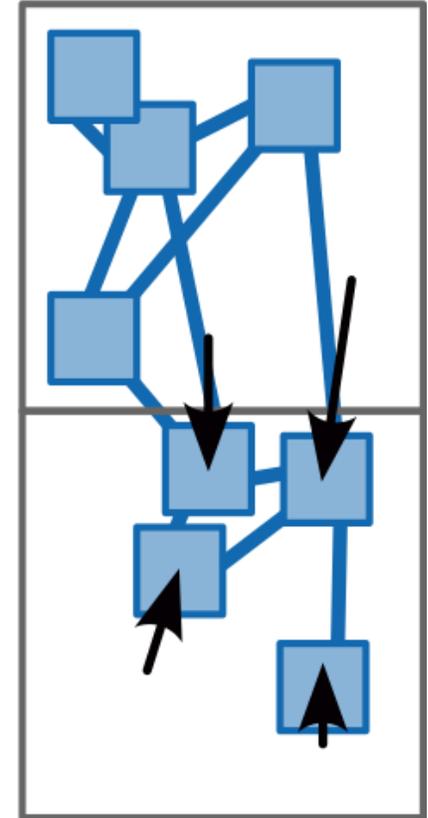


- **Goal:**
  - Minimize wire-length
  - Achieve target density
- **Wire-force** attracts connected cells
- **Density-force** push overflowing cells away
- **Total force per cell**
  - Cell starts moving in this direction
  - Incrementally adds velocity
- **Placement Tool**
  - Find useful force-balance
  - Weight critical nets/cells

target: 4  
actual: 6 **over: 2**



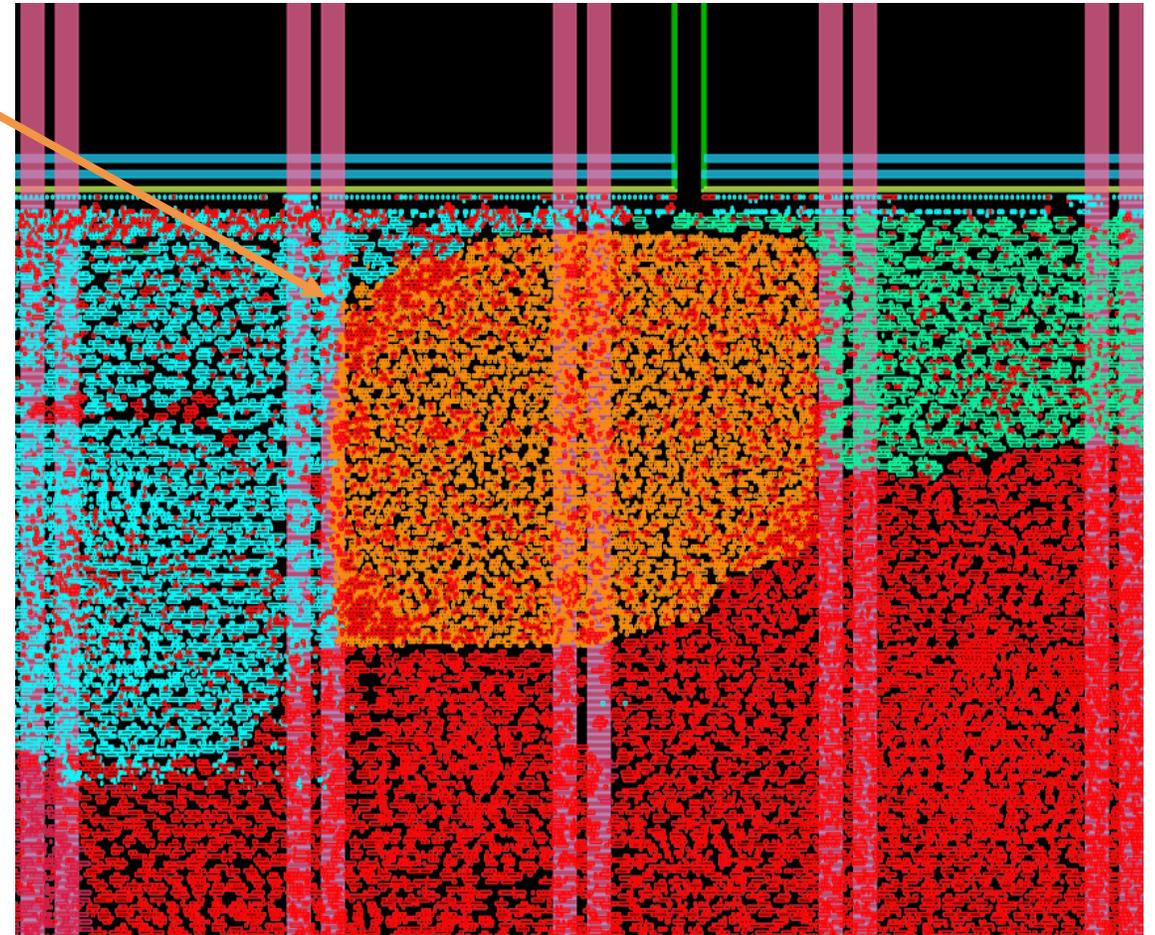
target: 4  
actual: 2 **over: 0**



# Power Stripes are Fences

- **Initial power stripes**
  - Wide VSS/VDD in pairs

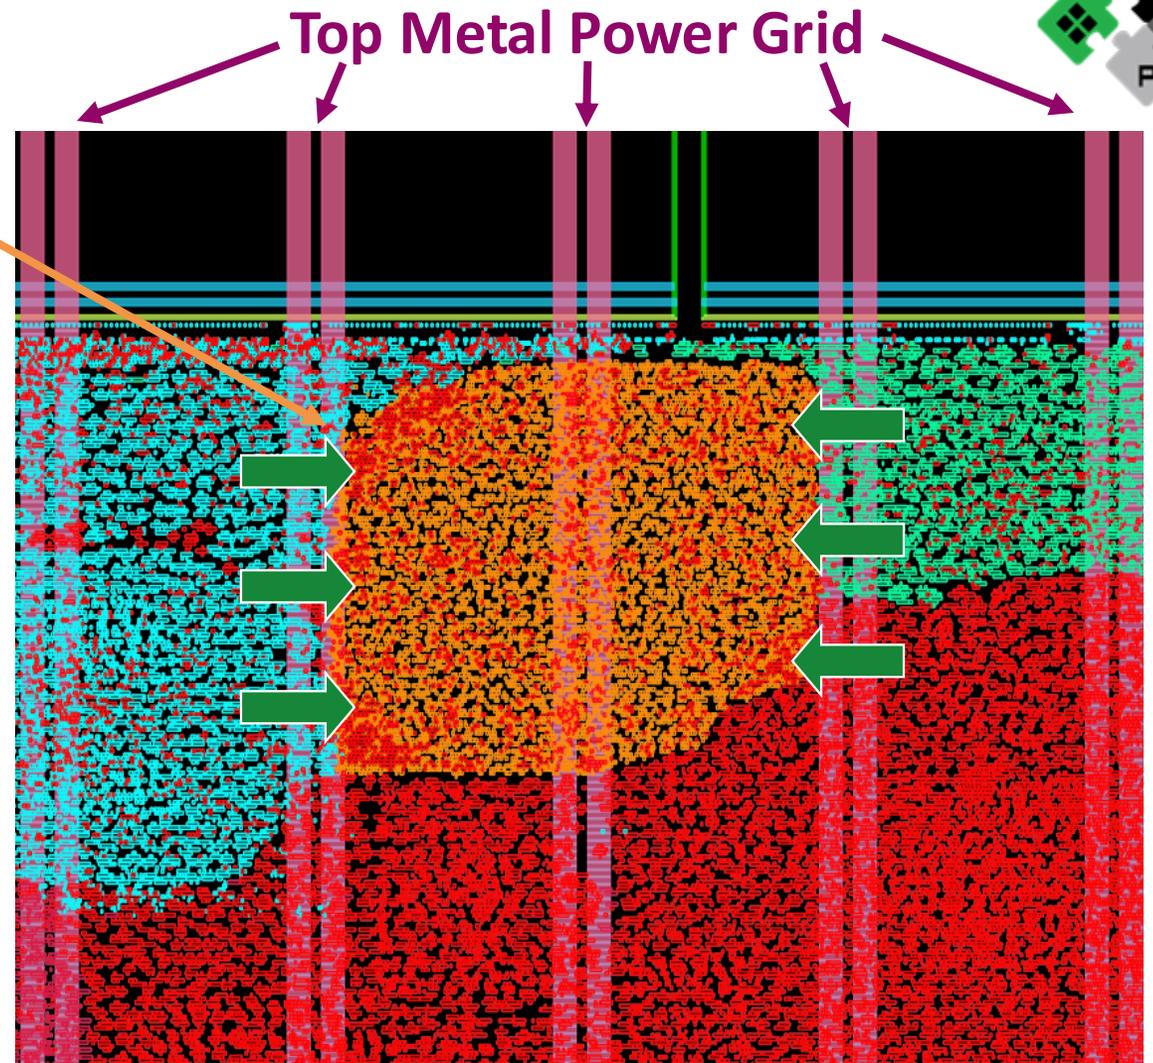
Bootrom



# Power Stripes are Fences

- **Initial power stripes**
  - Wide VSS/VDD in pairs
- **Wire-force pulls in**
  - High penalty for crossing stripes

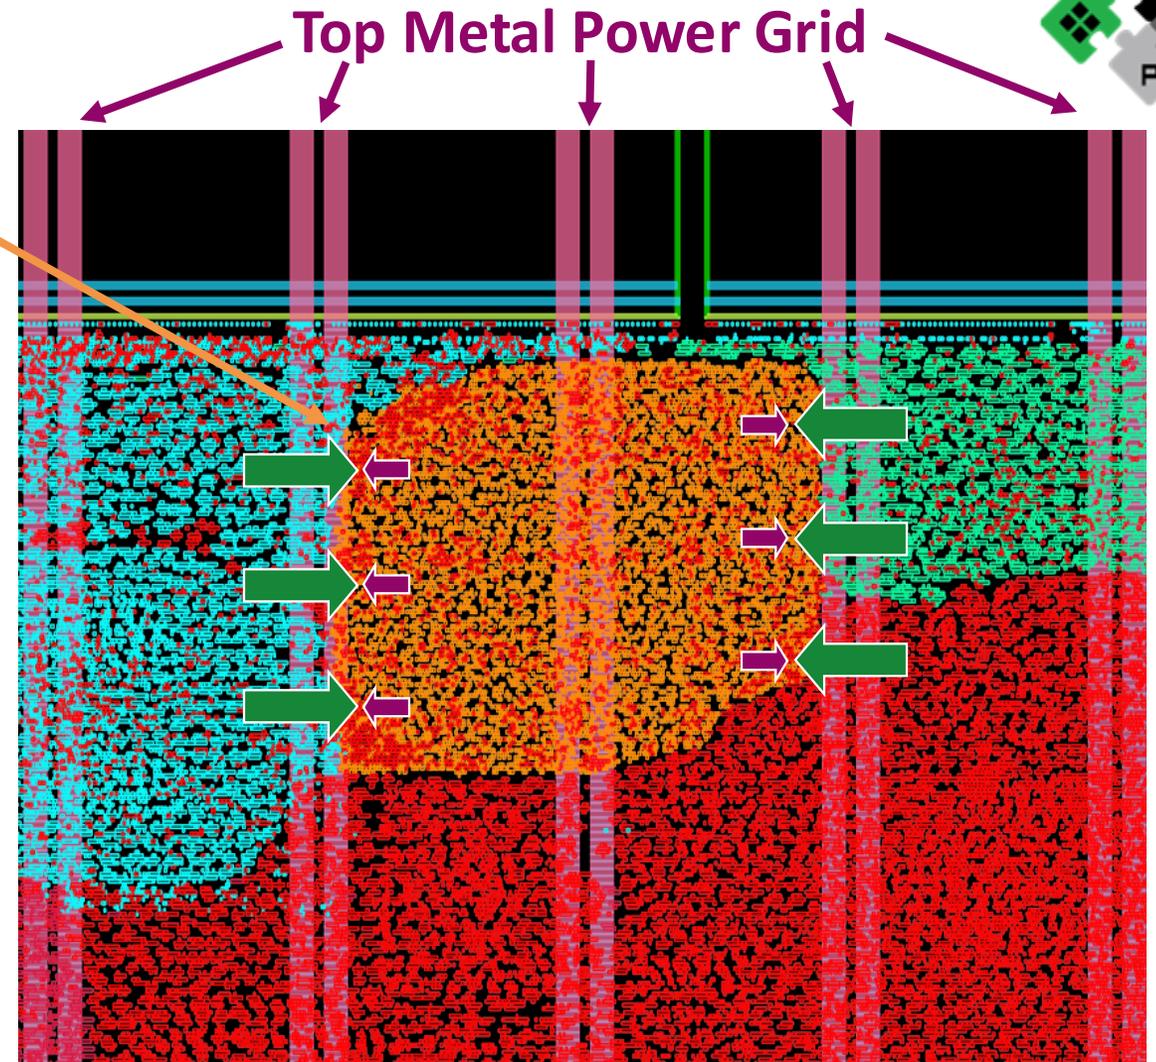
Bootrom



# Power Stripes are Fences

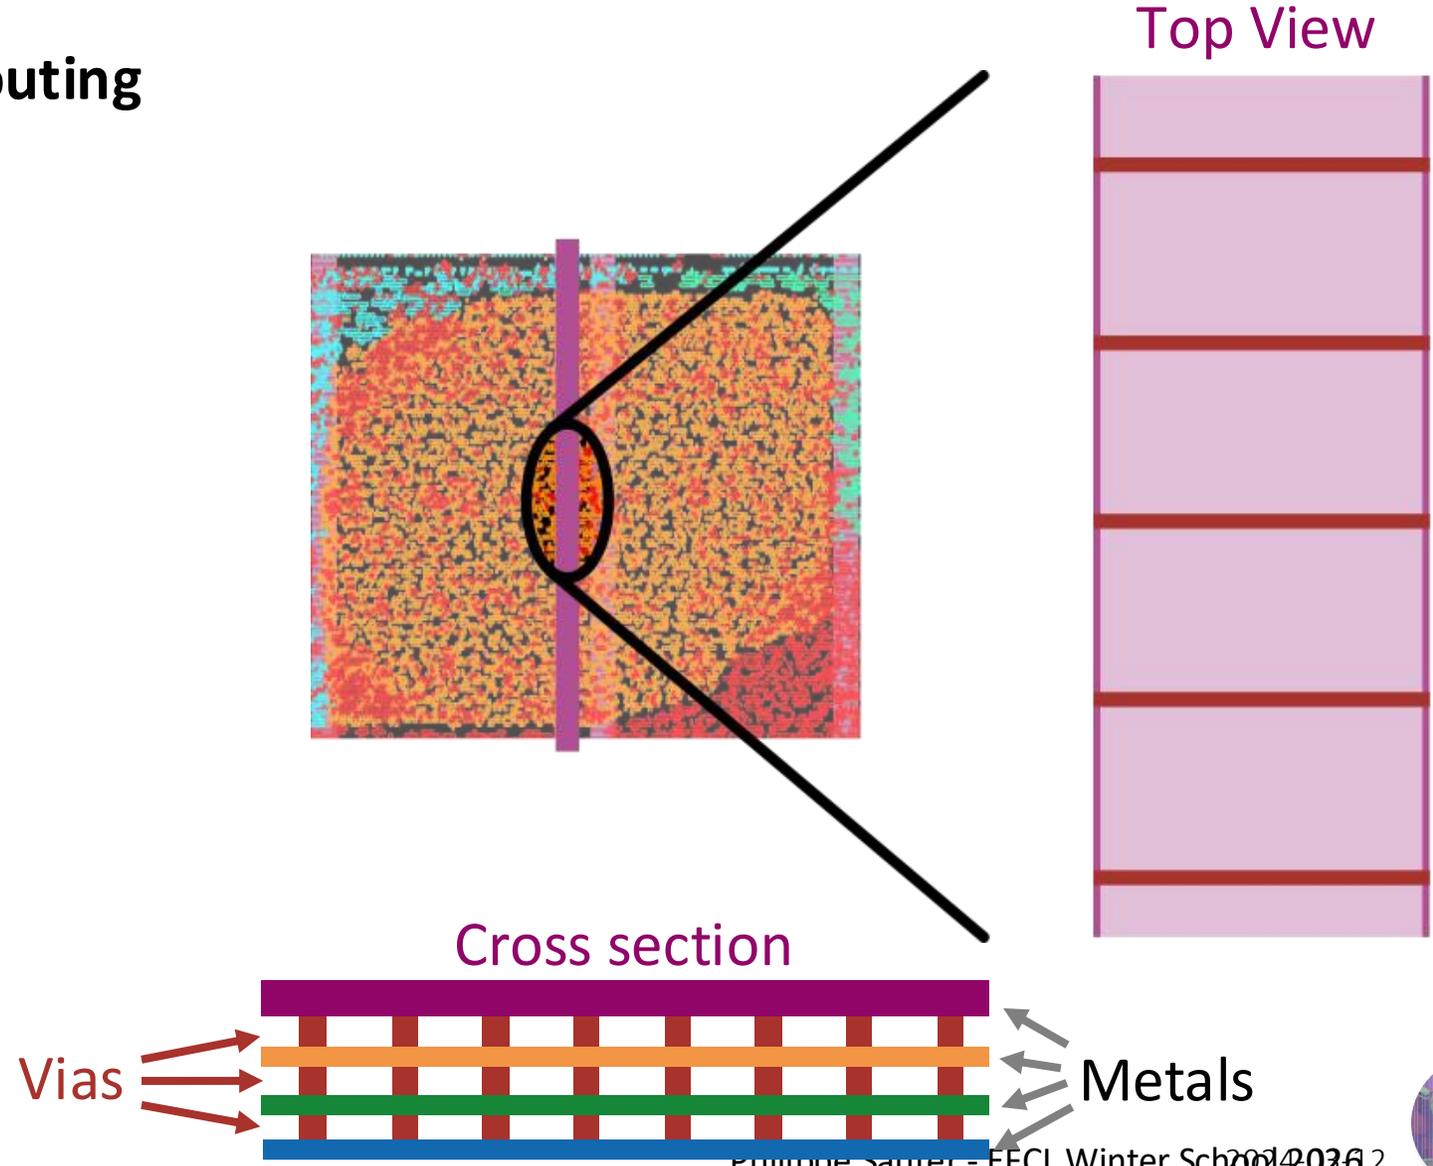
- **Initial power stripes**
  - Wide VSS/VDD in pairs
- **Wire-force pulls in**
  - High penalty for crossing stripes
- **Density-force too small**
  - Trying to decrease local density
  - Pushes cells below stripes
- **Solution: more even distribution**
  - Thinner, more frequent power stripes

Bootrom



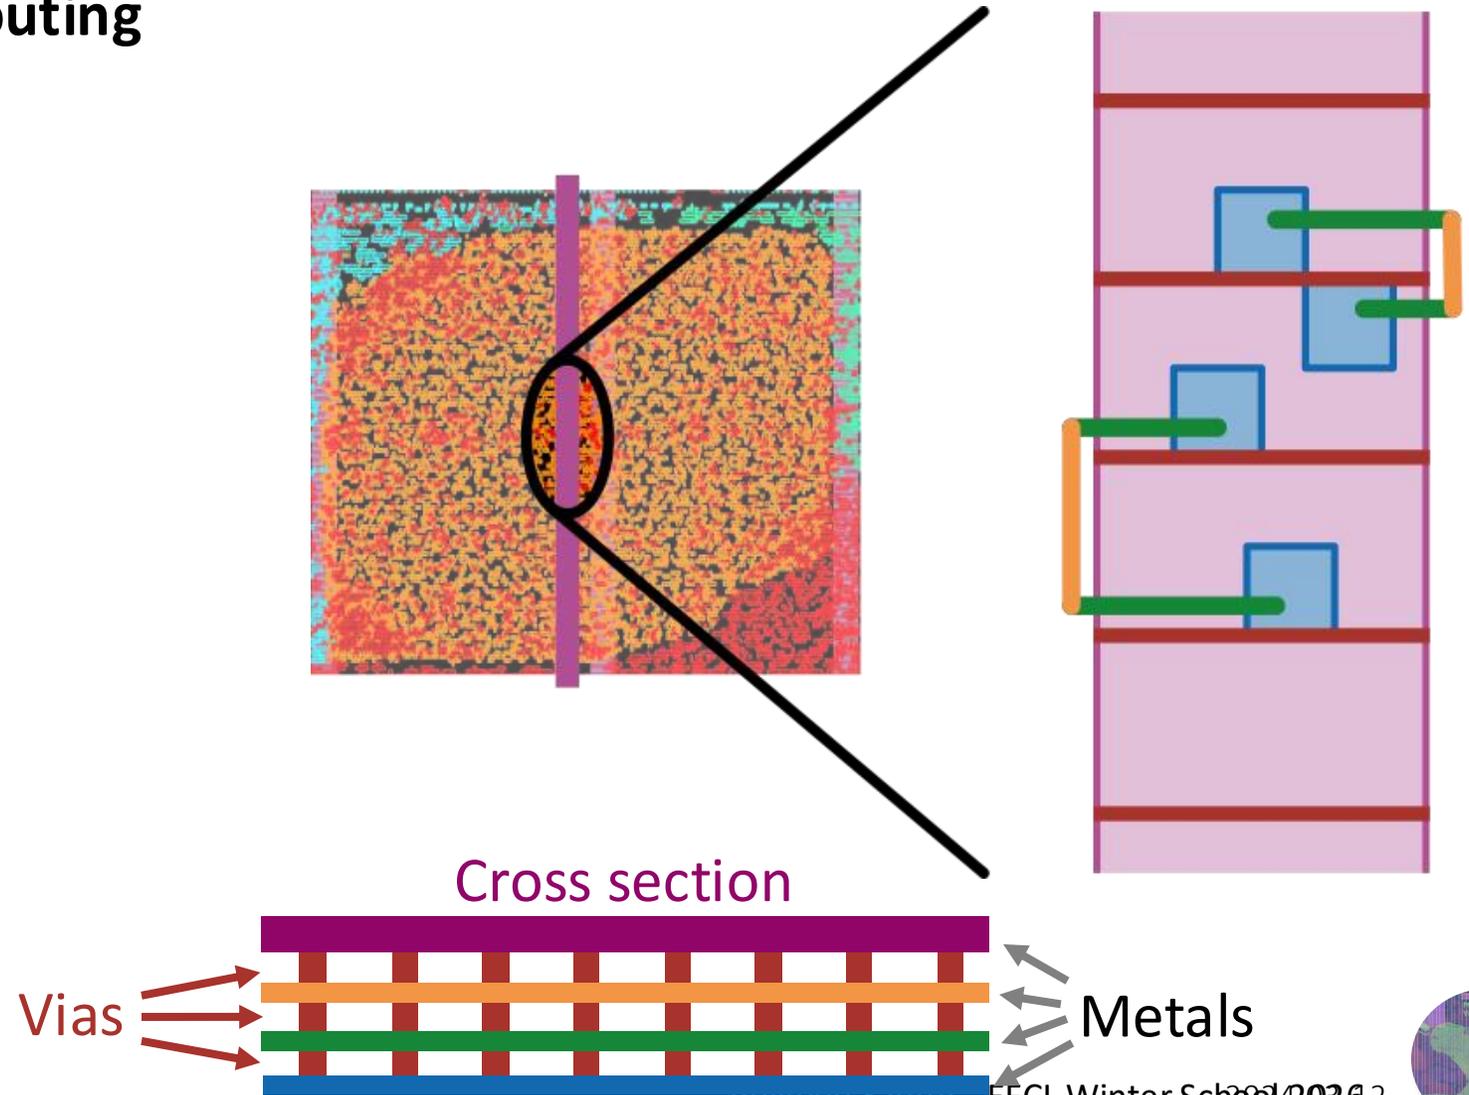
# Power Stripes Block Wires

- Power via-stack block routing



# Power Stripes Block Wires

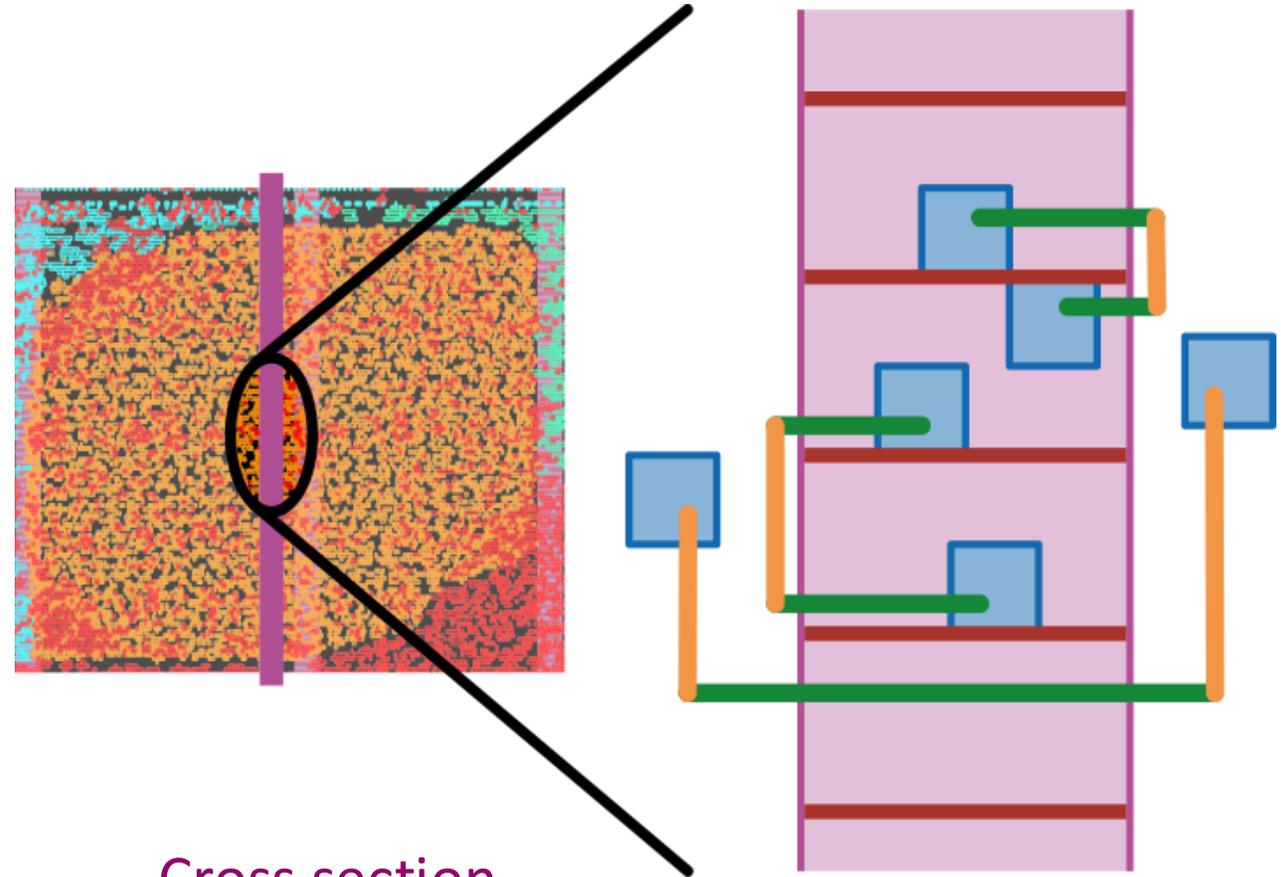
- Power via-stack block routing
- Cells under stripe
  - No vertical routing
  - Router needs to go around
  - Longer horizontal wires



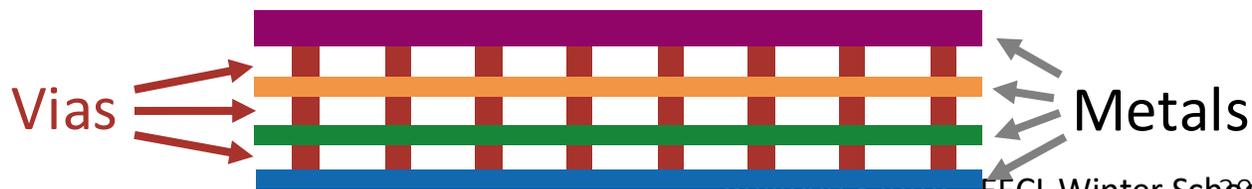
# Power Stripes Block Wires



- **Power via-stack block routing**
- **Cells under stripe**
  - No vertical routing
  - Router needs to go around
  - Longer horizontal wires
- **Cells near stripe**
  - Goes to next unused opening
  - Longer vertical wires



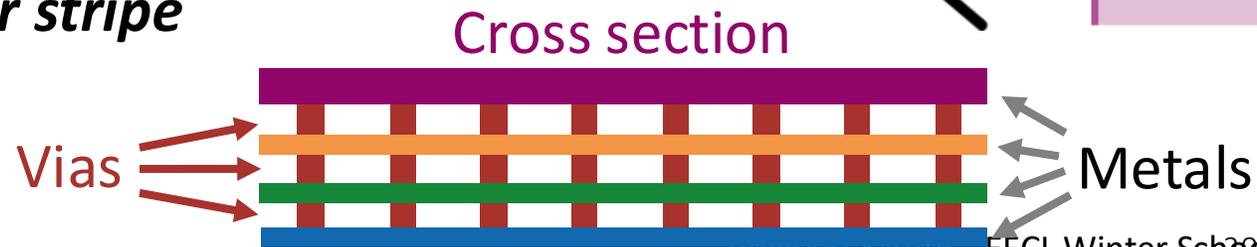
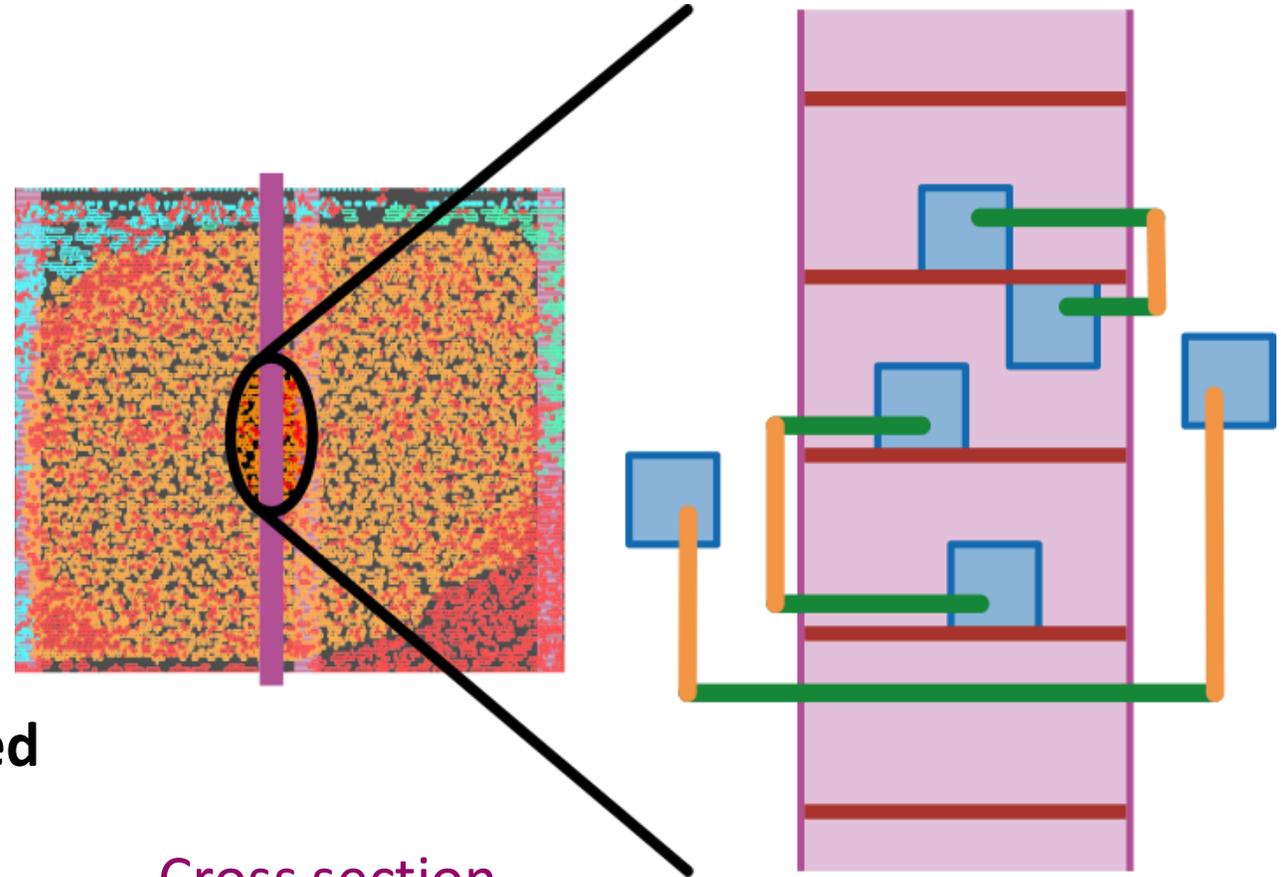
Cross section



# Power Stripes Block Wires

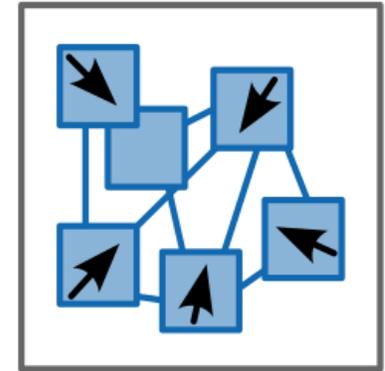
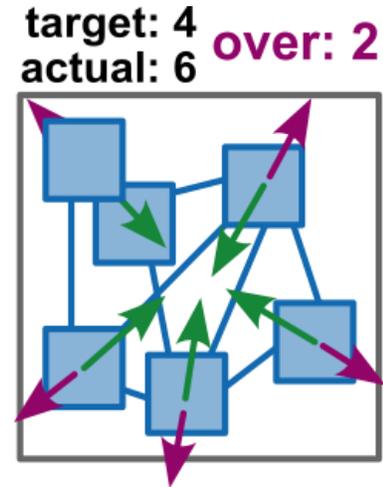


- **Power via-stack block routing**
- **Cells under stripe**
  - No vertical routing
  - Router needs to go around
  - Longer horizontal wires
- **Cells near stripe**
  - Goes to next unused opening
  - Longer vertical wires
- **Routing resources are overestimated**
- ***Solution: Blockage under stripe***



# Parameter Tuning Required

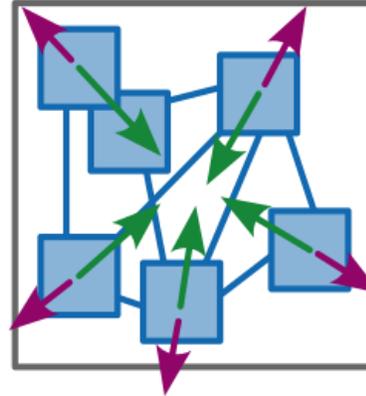
- **Highly interconnected modules**
  - More wires per cell
  - **Wire-force** dominates
  - Pulls cells tightly together



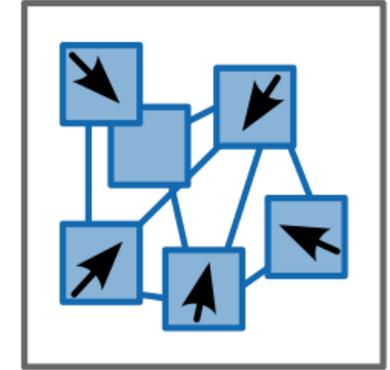
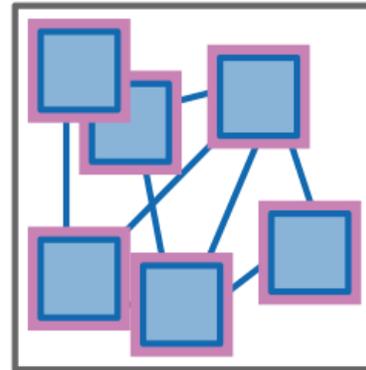
# Parameter Tuning Required

- **Highly interconnected modules**
  - More wires per cell
  - **Wire-force** dominates
  - Pulls cells tightly together
- **Solution: Routability driven**

target: 4  
actual: 6 **over: 2**

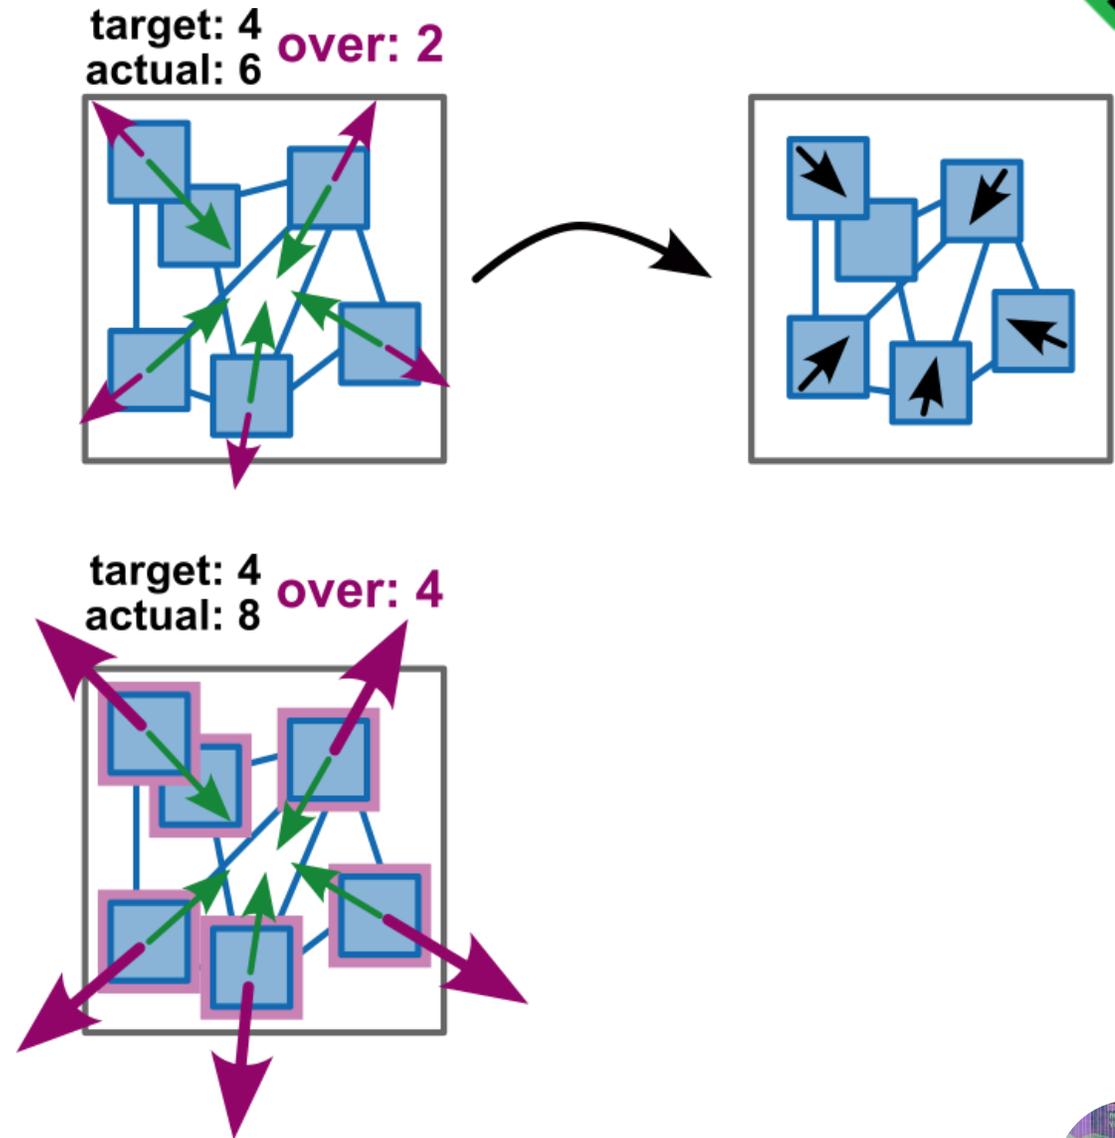


target: 4  
actual: 8 **over: 4**



# Parameter Tuning Required

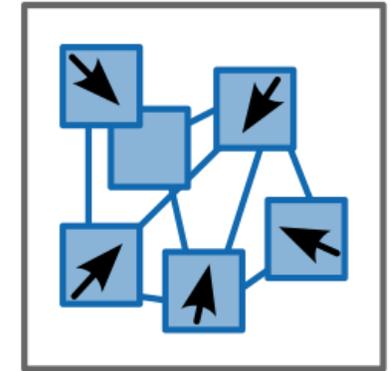
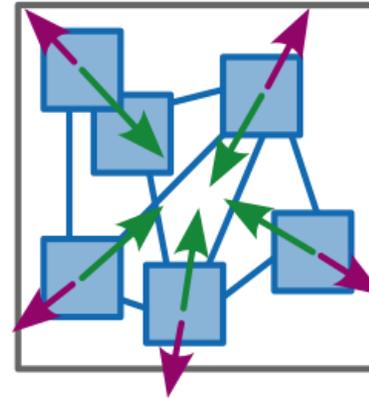
- **Highly interconnected modules**
  - More wires per cell
  - **Wire-force** dominates
  - Pulls cells tightly together
- **Solution: Routability driven**
  - Inflate cell area if hard to solve



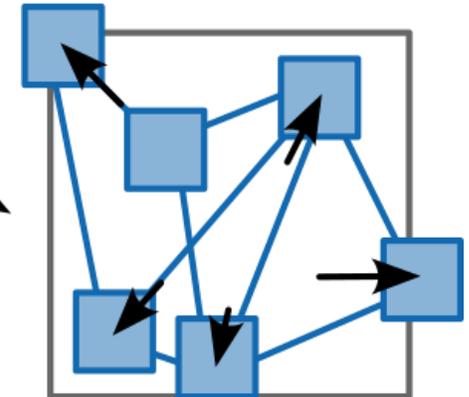
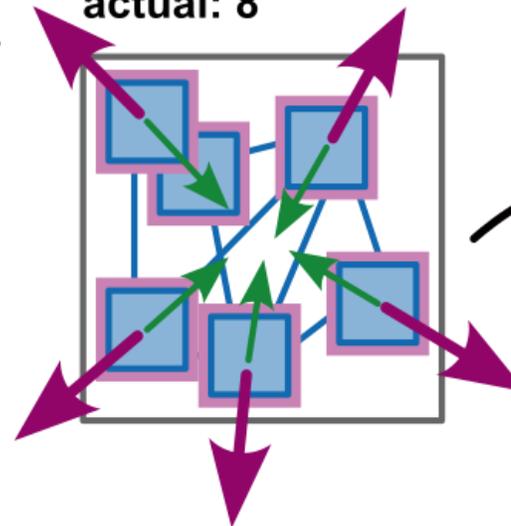
# Parameter Tuning Required

- **Highly interconnected modules**
  - More wires per cell
  - **Wire-force** dominates
  - Pulls cells tightly together
- **Solution: Routability driven**
  - Inflate cell area if hard to solve
  - Requires manually tuning hyper-parameters
  - `check_overflow`: How early OR uses it
  - `max_inflation_iter`: Number of iterations
  - `inflation_ratio_coef`: Virtual area increase on difficult cell

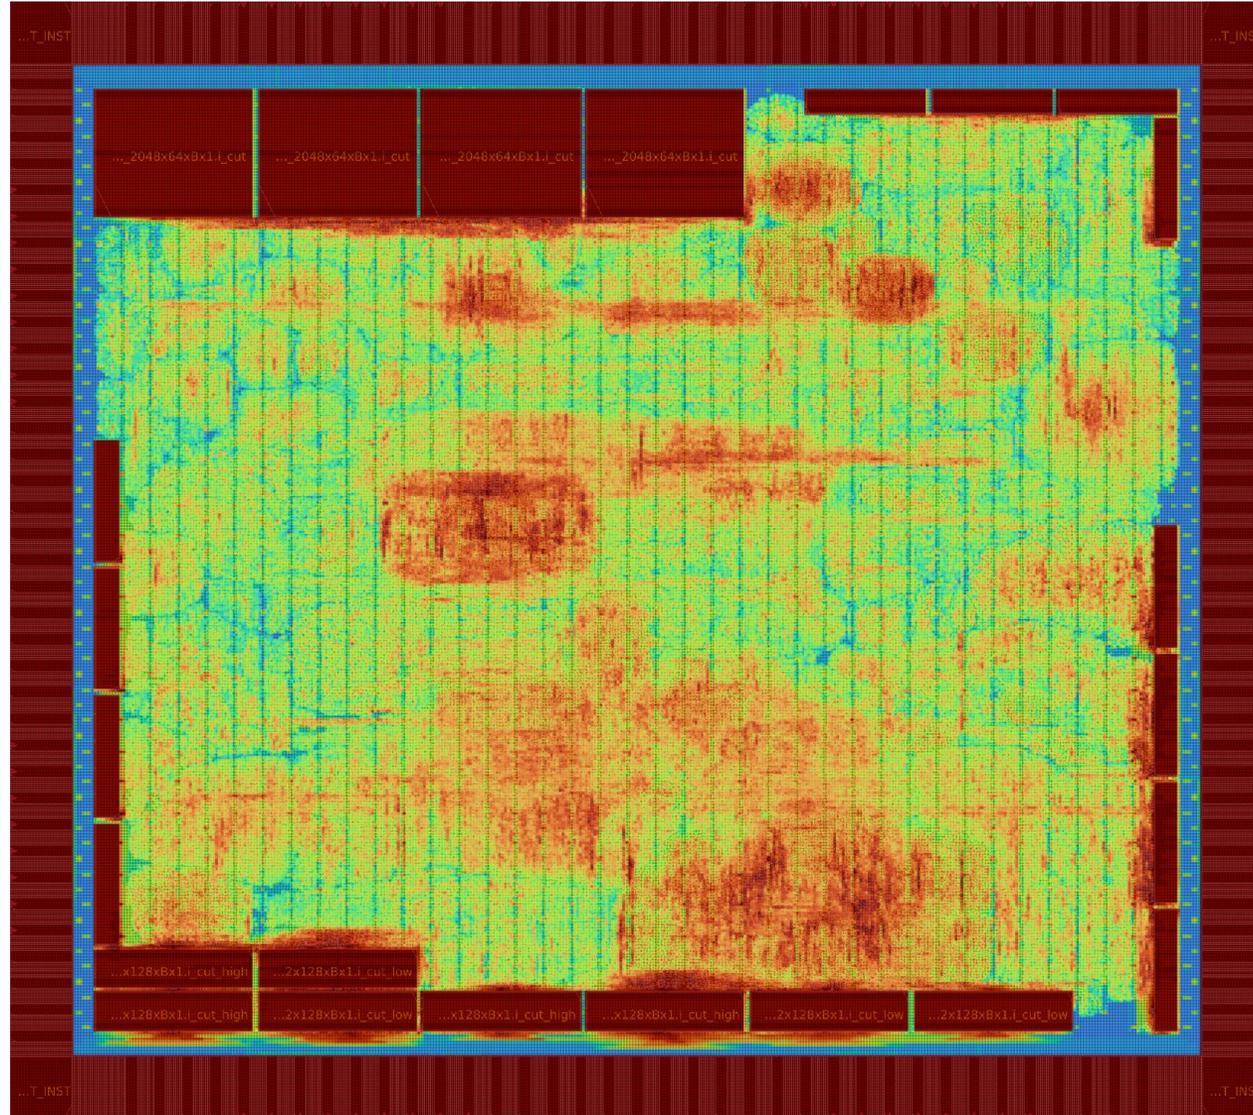
target: 4  
actual: 6 **over: 2**



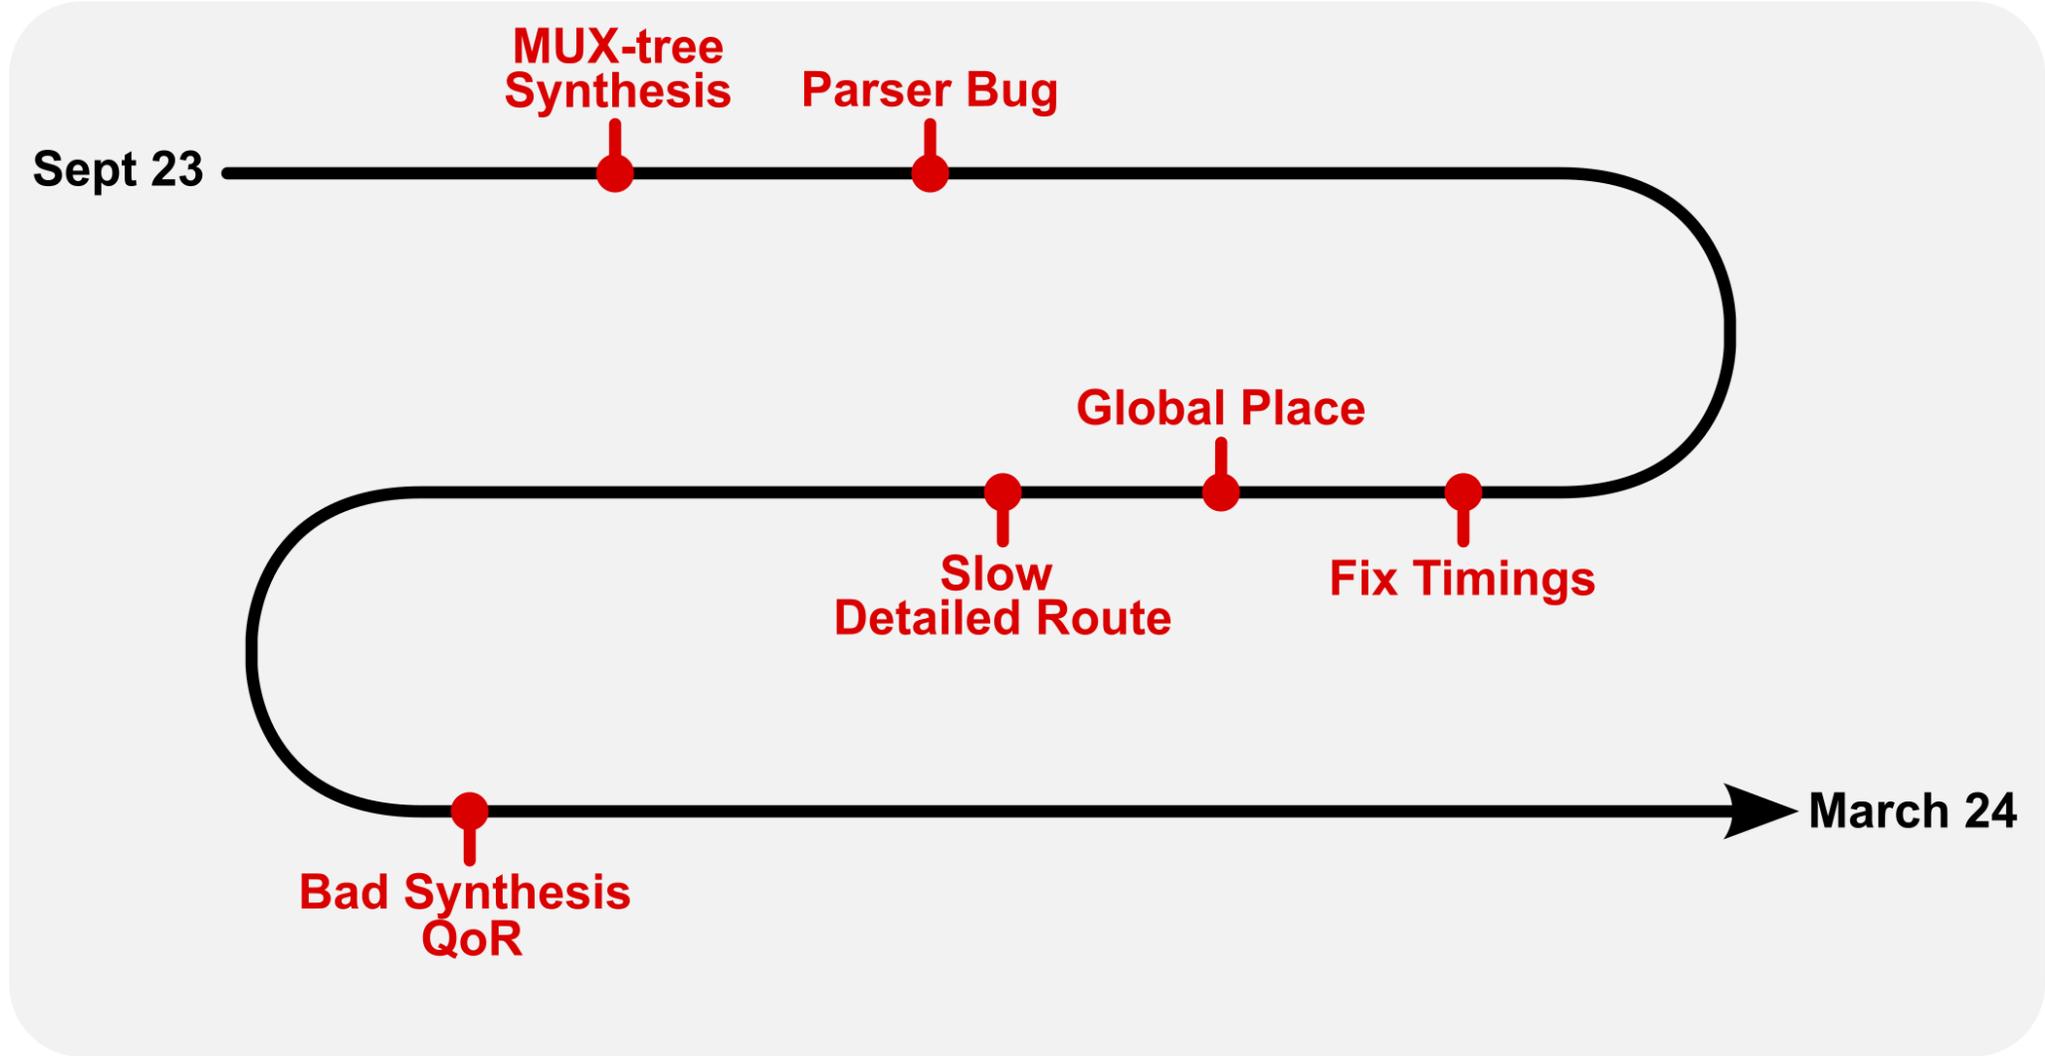
target: 4  
actual: 8 **over: 4**



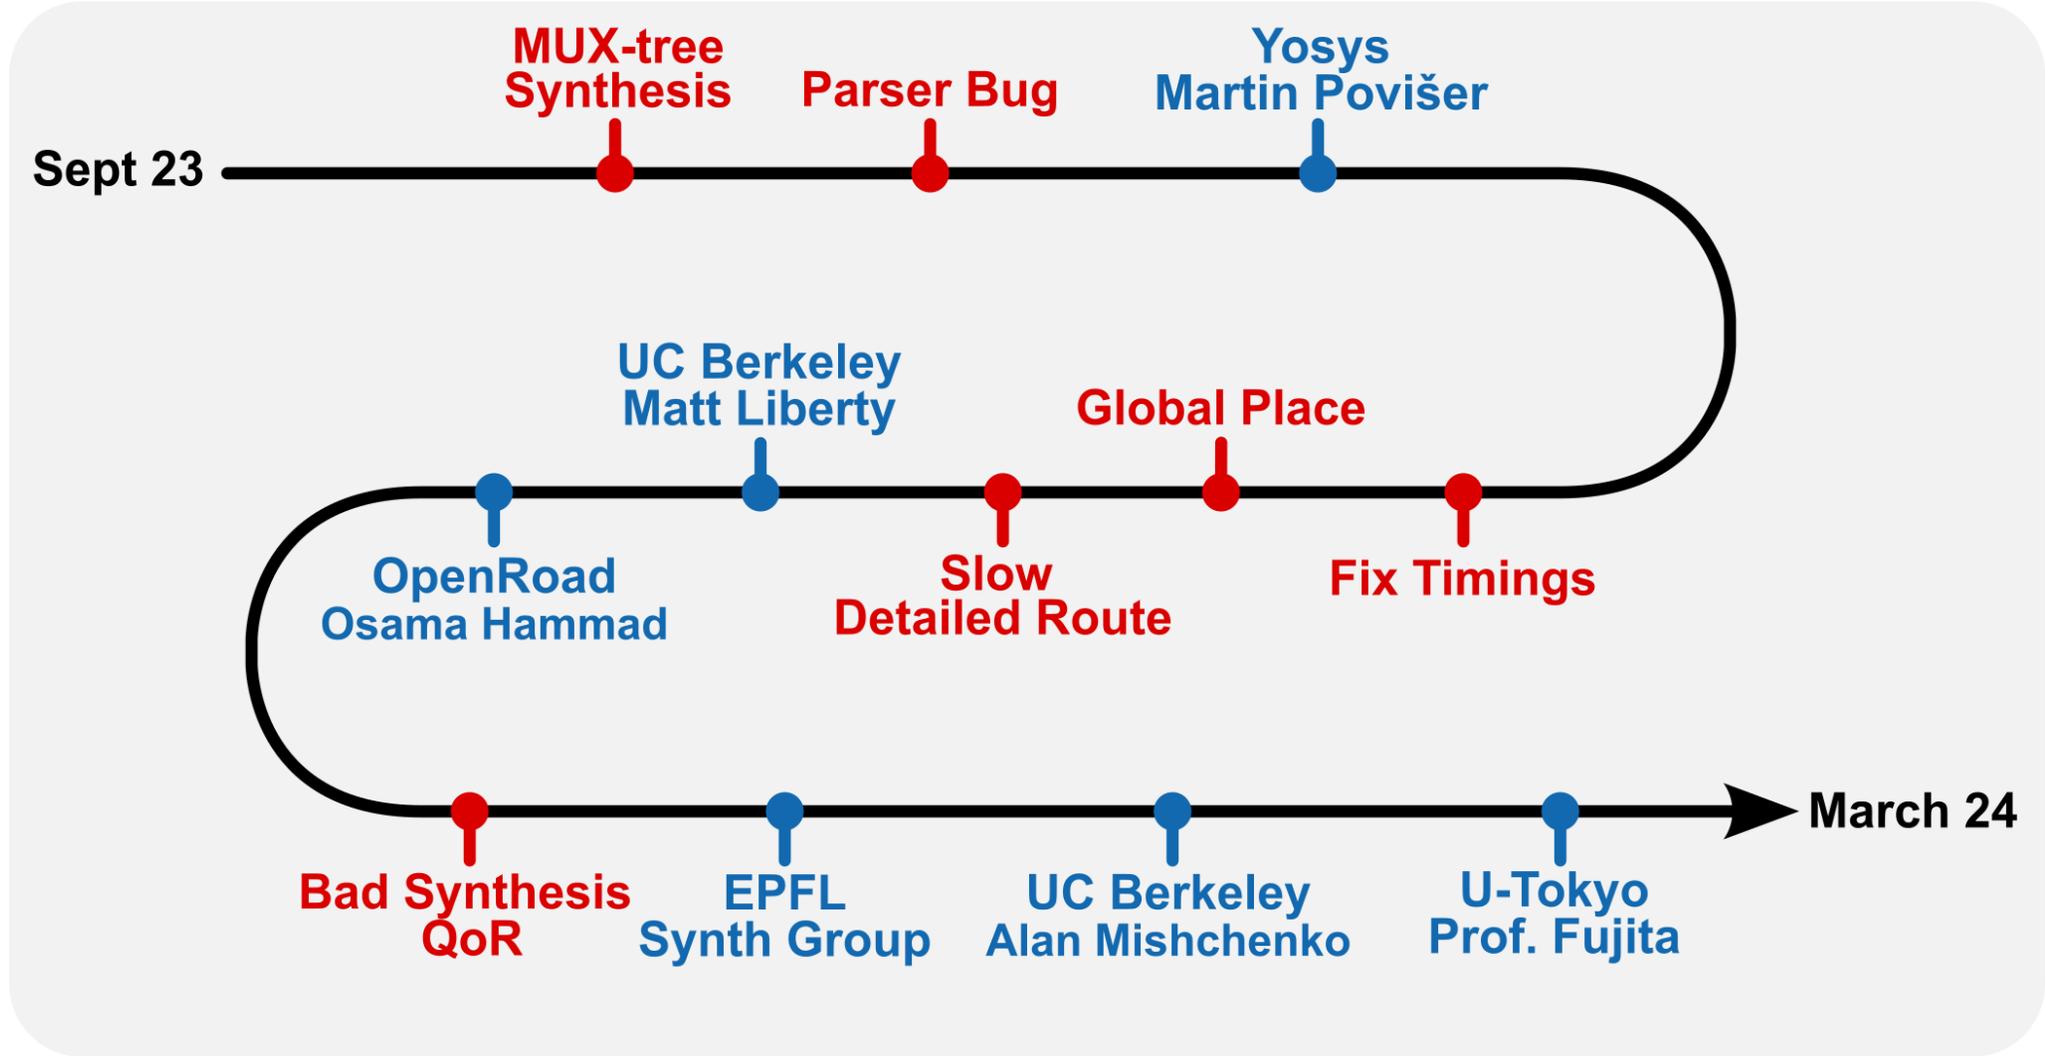
# Final Placement Result



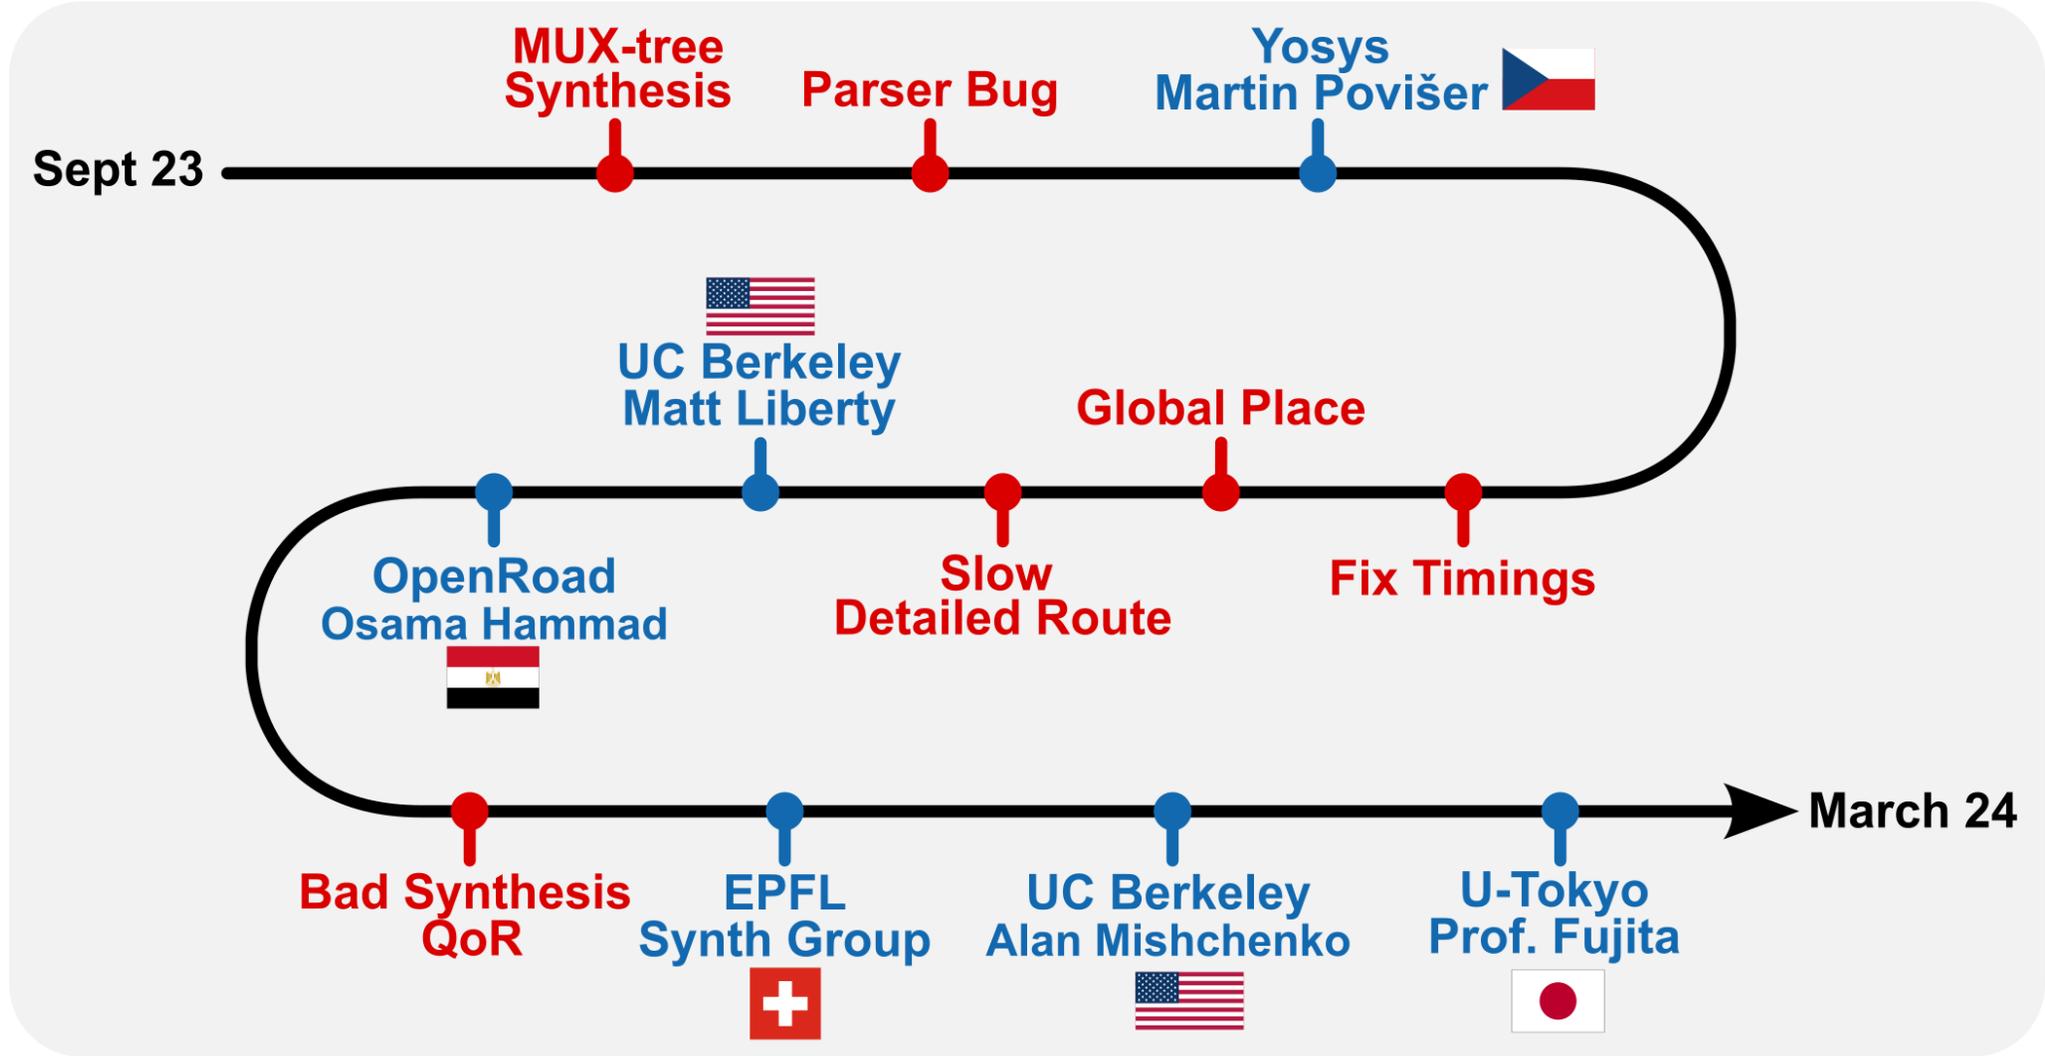
# Potholes and Speed-bumps



# Teamwork makes the Dream Work!



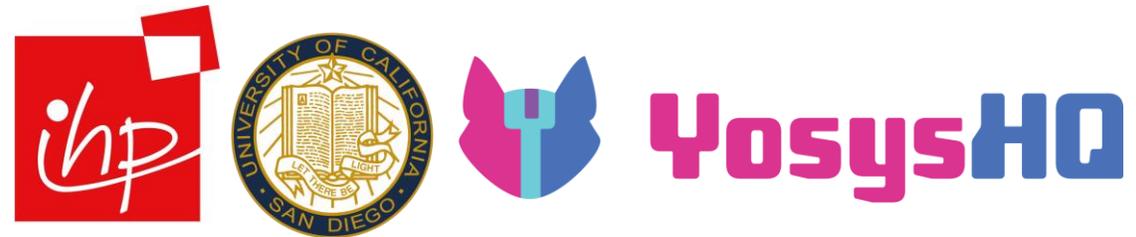
# Teamwork makes the Dream Work!



# Meet Basilisk: Open RTL, Open EDA, Open PDK

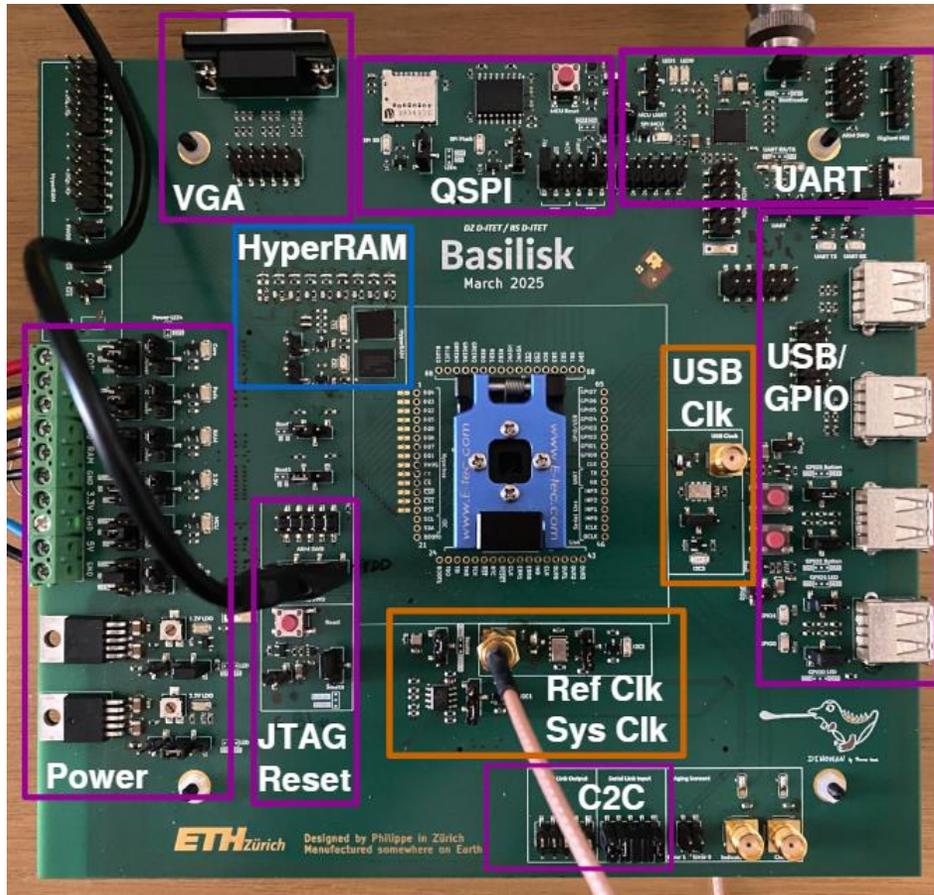


- **Designed in IHP 130nm OpenPDK**
  - **34mm<sup>2</sup>** (6.25mm x 5.50mm)
  - **~5× larger** than previous end-to-end OS designs
  - 2.7 MGE total, 1.14MGE logic
  - 24 SRAM macros (114 KiB)
  - **62MHz** at nominal voltage (1.2V)
- **RV64GC design runs Linux**
- **Active collaboration with**



Open-source EDA tools already enable complex designs, it will only get better

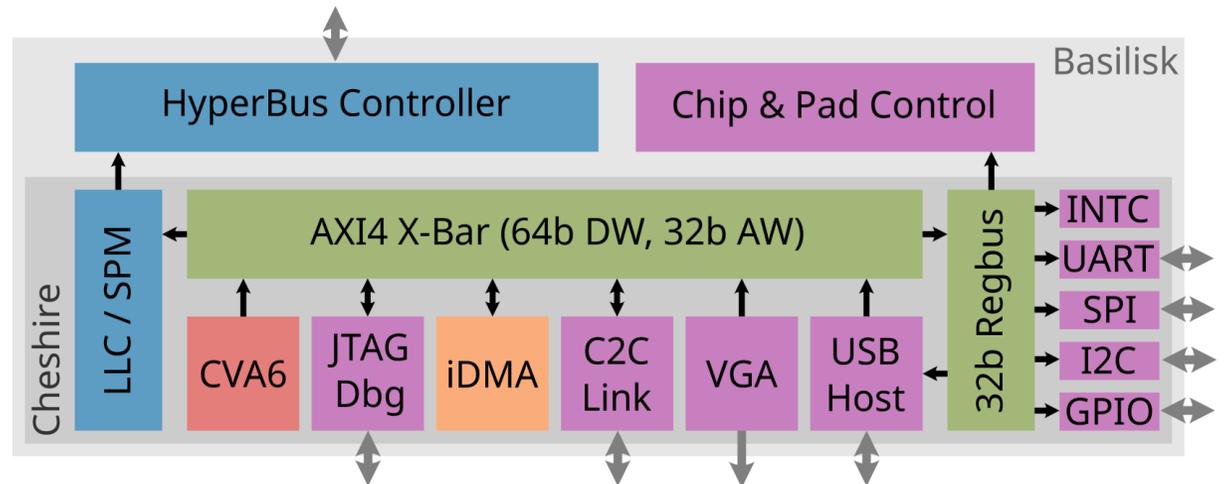
# Basilisk is a complete Linux-capable SoC



- 64-bit RISC-V core
- Rich peripherals:
  - HyperRAM controller @154MB/s
  - C2C AXI-Link @77MB/s
- Autonomous boot from SD-card



Watch the video



[arxiv.org/pdf/2505.10060](https://arxiv.org/pdf/2505.10060)



# We presented Basilisk at



**ETH zürich** **PULP**  
Parallel Ultra-Low Power

**34 mm<sup>2</sup> End-to-End Open-Source  
64-bit Linux-Capable RISC-V SoC in 130nm BiCMOS**

Authors: Philippe Sauter\*, Thomas Benz\*, Paul Scheffler, Martin Prader, Frank K. Gürkaynak\*, Luca Benini\*, ETH Zürich, University of Bologna

✉ psauter@ics.ee.ethz.ch

pulp-platform.org  
pulp-platform  
@pulp\_platform

**Our Design**

- 4-way 16KB L1 and L1D
- 64KB LLC/Scratchpad
- Hyperbus DRAM (124MB/s)
- 2.7MGE design

**Motivation**

Open-source as enabler

- Academia: NDA-free collaboration on designs and tools
- Education: Widely accessible hands-on chip design
- Industry: Zero-trust verification, license-free deployment

Explore feasibility for larger Linux-capable SoCs

- Previously: OS EDA tools used for tapeout of small designs
- Now: 4.8x larger than largest previously published design
- Novelty: Significant improvements in QoR and performance of open-source EDA tools and flow

**End-to-End Open-Source EDA flow**

RTL written by PULP (Cheshire, AXI...), OpenTitan (SPI, I2C) and OpenHW (CVA8 core)

**Yosys-Slang**: Newly developed frontend

- Supports industry-grade SystemVerilog

**Improved Yosys synthesis**

- Lazy man's synthesis for high effort optimization
- Improved bit-select operator to multiplexer mapping
- Optimized critical multiply-add implementation
- Improved timing (2.3x), area (1.6x) and runtime (2.5x) vs reference open-source flow
- 51 logic level critical path: Competitive with 49 LL in previous commercial implementations

**Tuned backend**: -12% die area vs open reference flow

- Based on OpenROAD-flow-scripts flow
- Flow and Hyperparameters tuned to design

**Conclusion**

- First end-to-end permissive, free and open-source Linux-capable SoC with an application-class core and rich peripherals
- Newly developed Yosys-Slang enables synthesis of complex, industry-grade SystemVerilog RTL
- Reproducible and sharable high-quality designs for collaboration and research

**2025 Hot Chips Best Student Poster Award**

*Basilisk: A 34 mm<sup>2</sup> End-to-End Open-Source 64-bit Linux-Capable RISC-V SoC in 130nm BiCMOS*

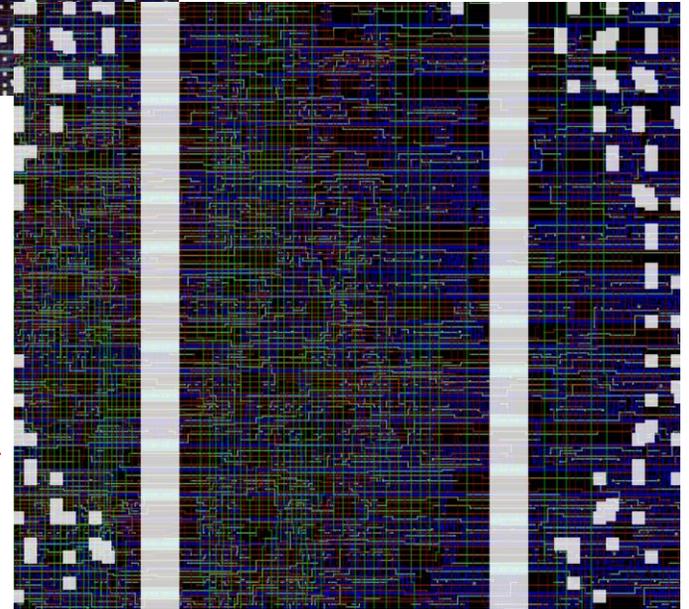
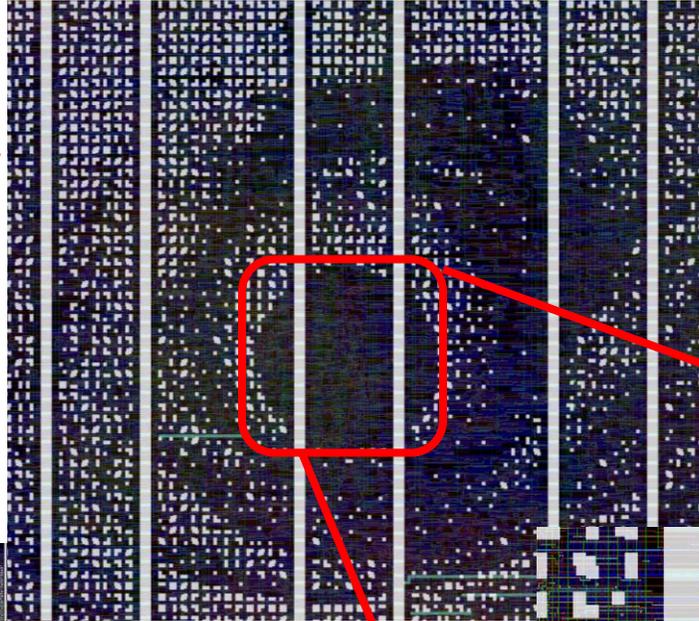
**Industry Noticed!**

"Basilisk at Hot Chips 2025 Presented Ominous Challenge to IP/EDA Status Quo"\*

Poster: [lnkd.in/d/aB6HskB](https://lnkd.in/d/aB6HskB)



# Its Open-Source down to the Transistors





Where do we go next?



# The PULP Open-Source Ecosystem



## RISC-V Cores and Vector Units

|                       |                       |        |       |                       |     |
|-----------------------|-----------------------|--------|-------|-----------------------|-----|
| RI5CY<br><i>CV32E</i> | Zero R<br><i>lbex</i> | Snitch | Spatz | Ariane<br><i>CVA6</i> | ARA |
| RV32                  | RV32                  | RV32   | RVV   | RV64                  | RVV |

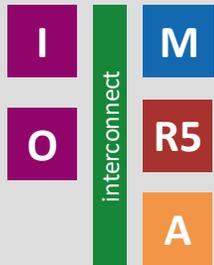
## Peripherals

|      |      |
|------|------|
| JTAG | SPI  |
| UART | I2S  |
| DMA  | GPIO |

## Interconnects

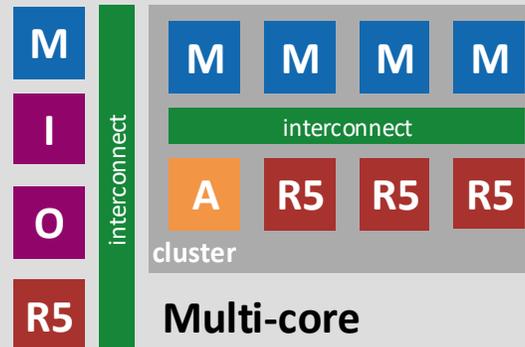
|      |         |
|------|---------|
| LIC  | HCI     |
| APB  | FlooNoC |
| AXI4 |         |

## Platforms



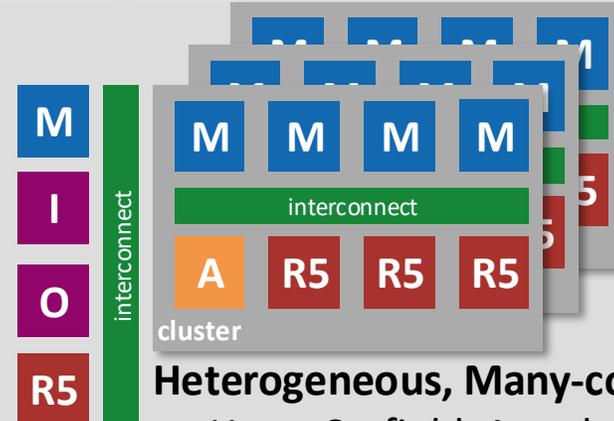
### Single core

- PULPino, PULPissimo
- **Cheshire**



### Multi-core

- OpenPULP



### Heterogeneous, Many-core

- Hero, Carfield, Astral
- Occamy, Mempoal

IOT

HPC

## Accelerators and ISA extensions

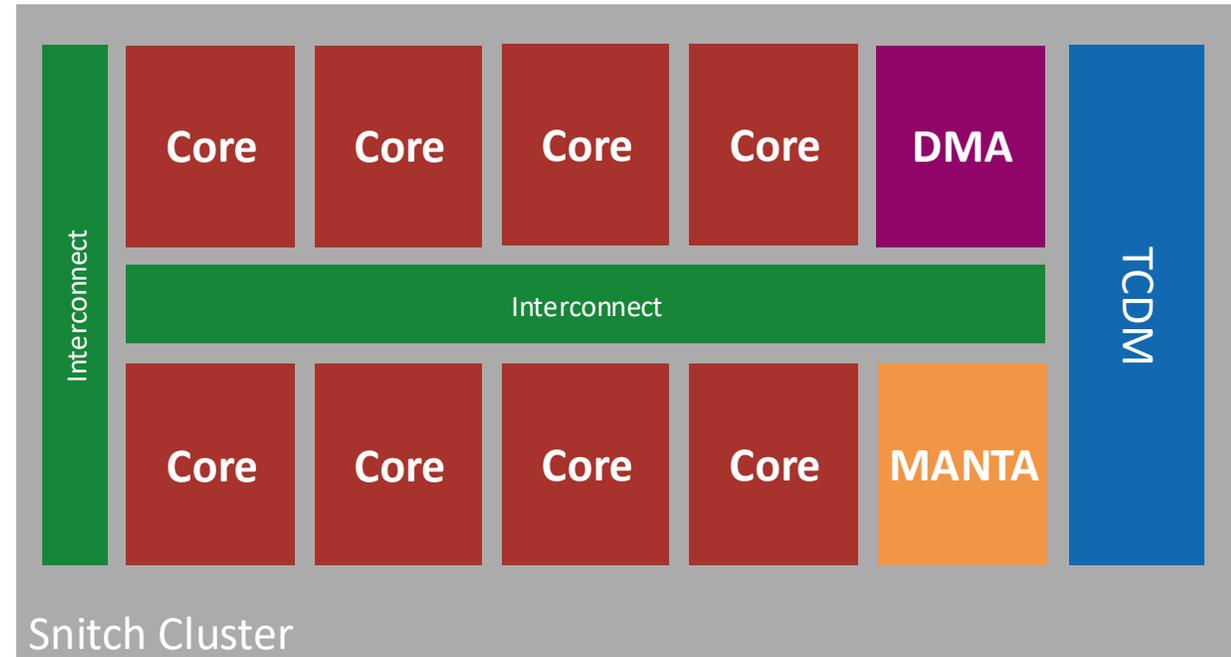
|                      |                       |                        |              |                        |
|----------------------|-----------------------|------------------------|--------------|------------------------|
| XpulpNN,<br>XpulpTNN | ITA<br>(Transformers) | RBE, NEUREKA<br>(QNNs) | FFT<br>(DSP) | REDMULE<br>(FP-Tensor) |
|----------------------|-----------------------|------------------------|--------------|------------------------|



# Exploring Scientific Designs with Open-Source Tools



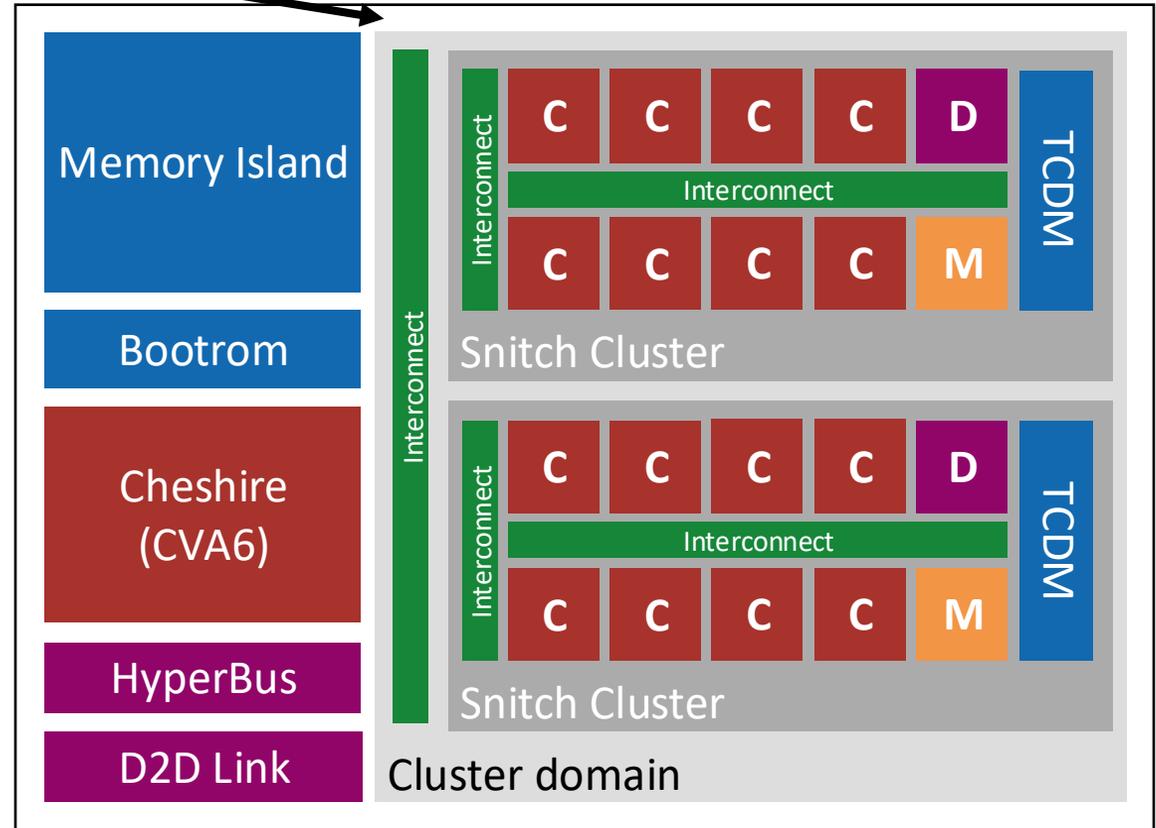
- **Tightly Coupled L1 Memory** ■
  - Hierarchical organisation of banks
- **MANTA Accelerator** ■
  - Accelerates GEMM in MXINT8 format
  - **Main scientific contribution and focus**
- **Snitch Cores** ■
  - Optimized for data processing
  - Scalable number of cores
  - **Various extensions optimized for ML**
- **Dedicated DMA Core** ■
  - Connected to a wide AXI interface
  - Moves data to/from L2 and clusters



# SeyrITA: Open-Source Tools with GF22



- **Transformer acceleration**
  - DeiT Image Transformer (5M Params)
  - BERT language model (4-40M Params)
  - MobileBERT-Tiny (15M Params)
  - 2-4 compute clusters
- **Memory Island**
  - ~2MB SRAM per die
- **Linux Capable Host (Cheshire)**
  - Hyperbus for DRAM
  - SD-card for persistent Flash storage
- **Die-to-die Chiplet Interconnect**
  - 4Gbit/s per lane, 2-4 lanes



# SeyrITA: Open-Source Tools with GF22



- **Transformer acceleration**

- DeiT Image Transformer (5M Params)
- BERT language model (4-40M Params)
- MobileBERT-Tiny (15M Params)
- 2-4 compute clusters

- Memory Island

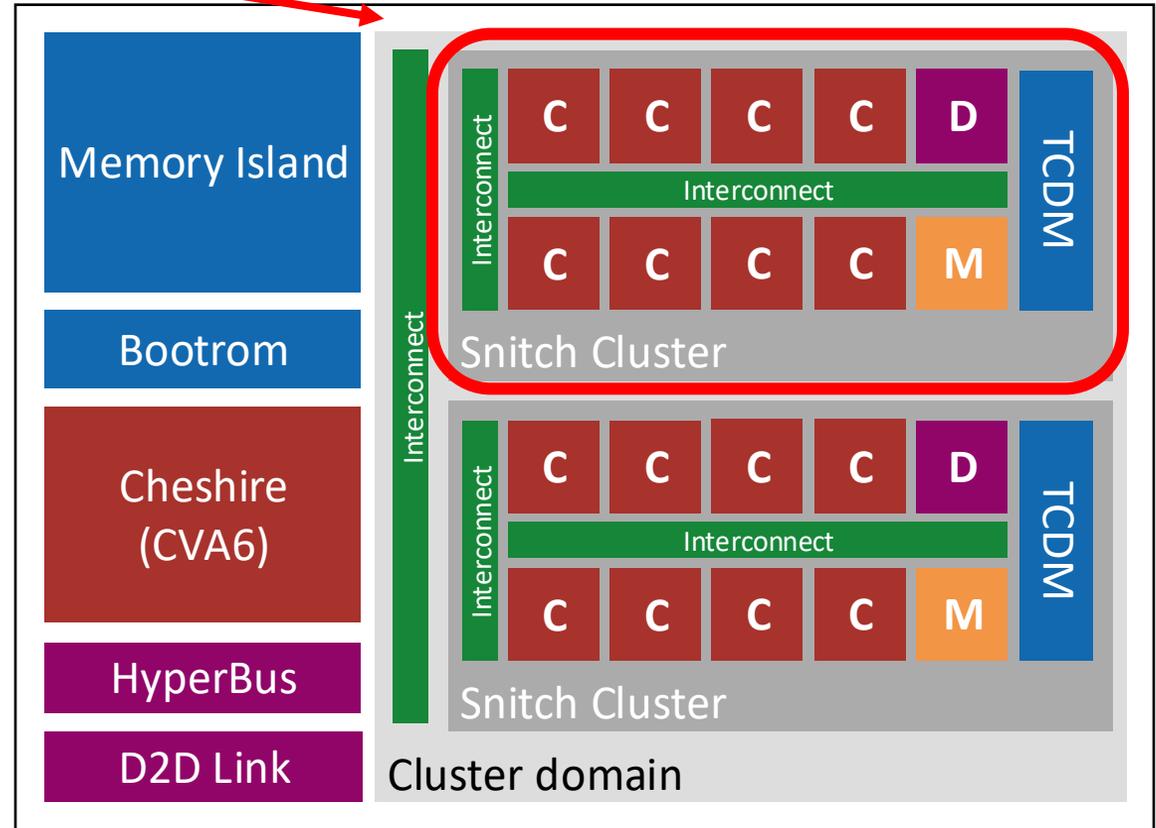
**The focus is on the Transformer Acceleration**

- **Linux Capable Host (Cheshire)**

- Hyperbus for DRAM
- SD-card for persistent Flash storage

- **Die-to-die Chiplet Interconnect**

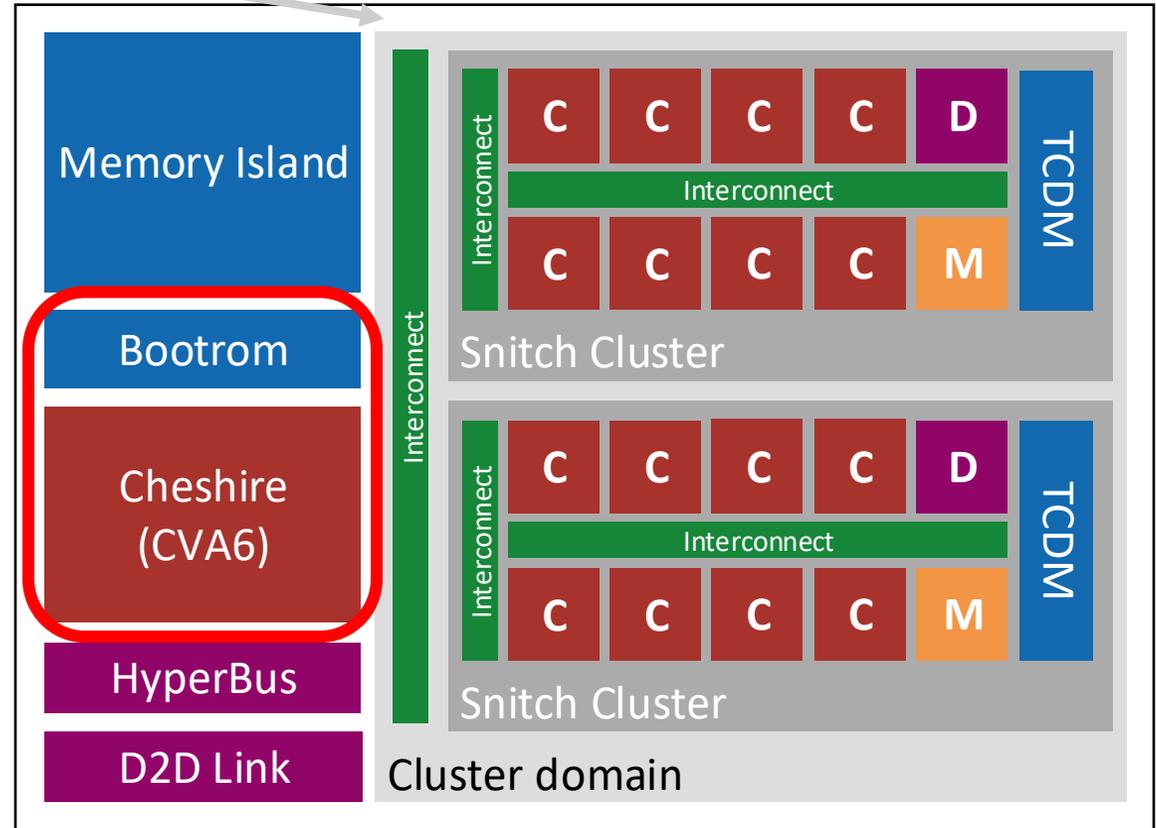
- 4Gbit/s per lane, 2-4 lanes



# SeyrITA: Open-Source Tools with GF22



- Transformer acceleration
  - DeiT Image Transformer (5M Params)
  - BERT language model (4-40M Params)
  - MobileBERT Tiny (15M Params)
  - 2-4 compute clusters
- Memory Island
  - Basilisk system as a small part
  - **SeyrITA is 10-15x Basilisk!**
- ~2MB SRAM per die
- Linux Capable Host (Cheshire)
  - Hyperbus for DRAM
  - SD-card for persistent Flash storage
- Die-to-die Chiplet Interconnect
  - 4Gbit/s per lane, 2-4 lanes



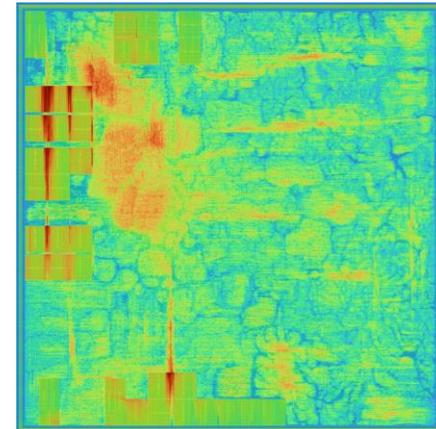
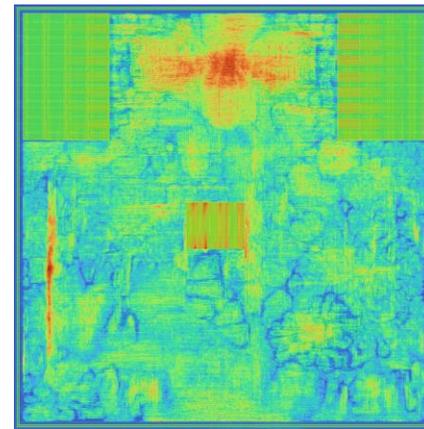
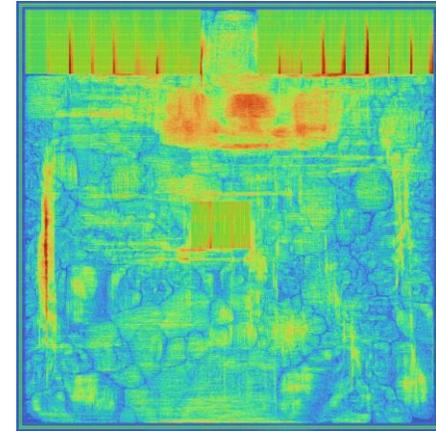
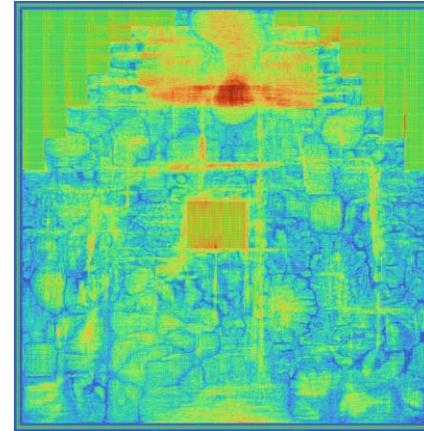
# On the Horizon: SeyrITA – GF22 with Open-Source Tools



- **Demonstrate** a large **22nm tapeout** with open-source tools
- **Improve tools** and close the **performance gap**
- Identify and **implement missing features** along the way
- **Active Collaboration with**



**YosysHQ**



Snitch cluster floorplan exploration



# Benefits of end to end openness



## Research

- Easier collaboration (no NDAs)
- Reproducible results, benchmarking
- **Combined impact of design and design automation**

## Industry

- Transparent chain of trust, sovereignty
- Lower initial cost
- **Accelerate research to product pipeline**

## Education

- Increased accessibility
- No black boxes, full visibility
- **Experiment with flows and tools**



# Any Questions?

- **We use open source because it works**
  - Allows us to manage complex designs
  - Facilitates Industry/Academia Relationships
  - Creates Auditable Designs, Reproducible Results
  - Enables research into new directions

There is still  
more to come



Helps us and others concentrate work where it matters

- **Open Source sees no borders**
  - There is **no** 'European/Chinese/American Open Source',
  - There **can be** 'European/Chinese/American support for Open Source'

Open Source is global, it just can have more or less support in a region/country



# Next up: Intro to Synthesis



There is still  
more to come

For you that's the  
Synthesis exercise 😊



**ETH**zürich Institut für Integrierte Systeme  
Integrated Systems Laboratory

Department of Information Technology and Electrical Engineering  
**EFCL Winter School 2026**

---

Exercise 03

---

**Synthesis with Yosys**

---

F. K. Gürkaynak  
Prof. L. Benini

Last Changed: 2026.02.09