ETH zürich

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

From an Open-Source ISA to Open-Source HW to Open-Source Silicon

Integrated Systems Laboratory (ETH Zürich)

Luca Bertaccini lbertaccini@iis.ee.ethz.ch

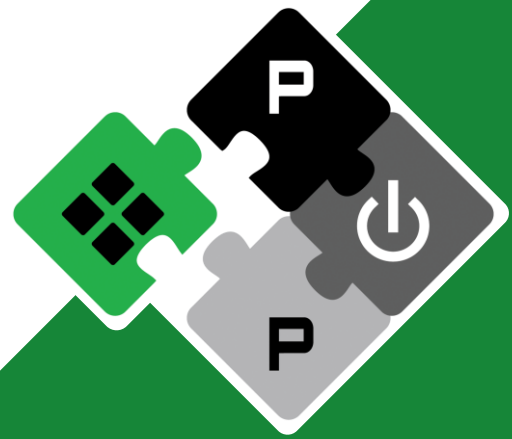

@pulp_platform

pulp-platform.org

youtube.com/pulp_platform

PULP Platform

Open Source Hardware, the way it should be!

Team of 100 people in ETH Zürich – University of Bologna

- **Research on energy-efficient computing architectures**
  - Started in 2013, celebrated 10 years of PULP last year
- **Led by Luca Benini**
  - Involves ETH Zürich (Switzerland) and University of Bologna (Italy)
  - Large group of almost 100 people

ETH zürich

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Team of 100 people in ETH Zürich – University of Bologna

- **Research on energy-efficient computing architectures**

We have designed and tested more than 60 PULP ICs

2013 (3)

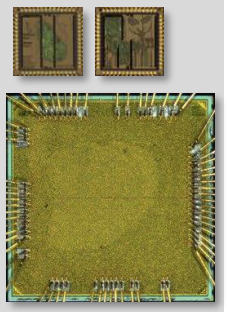

PULPv1
STM 28FDSOI
Multi-core processor

2014 (5)

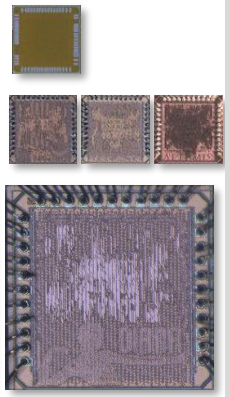

Diana
UMC 65
4-core system with approximate FPUs

2015 (10)

Fulmine
UMC 65
4-core system with ML and Crypto accelerators

2016 (3)

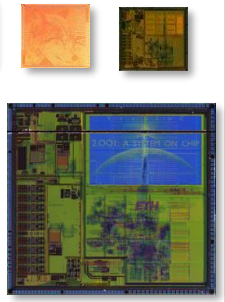

VivoSoC 2.001
SMIC 130
Mixed signal system for biosignal acquisiton

2017 (2)

Mr. Wolf
TSMC 40
8+1 core IoT processor

2018 (6)

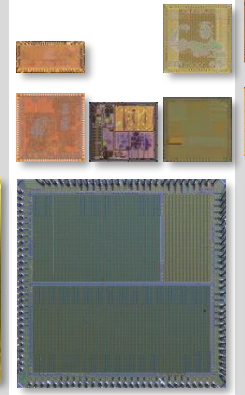

Poseidon
GF 22FDX
64bit RISC-V core, 32bit Microcontroller system, ML ...rator

2019 (7)

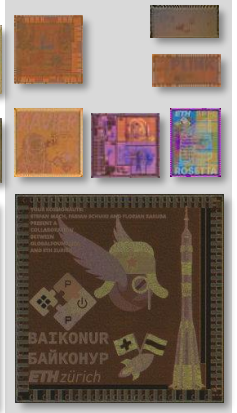

Baikonur
GF 22FDX
Dual 64bit RISC-V core, 3x 8core snitch clusters, Body biasing test vehicle

2020 (3)

Dustin
TSMC 65
IoT processor with 16 cores and QNN enhancements

2021 (7)

Kraken
GF 22FDX
IoT processor with Spiking Neural and Ternary Inference Engines

2022 (9)

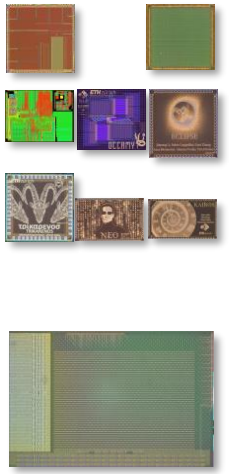

Occamy
GF 12LPP
ML accelerator with 216 + 1 cores and HBM interface

Check http://asic.ethz.ch for all our chips

We believe in open source

- **Collaboration with many different domain experts**
  - We cannot afford to do everything in-house
- **No long discussions on IP ownership and background IPs**
  - Everything in the open
- **We can start right away**
  - Time is spent on design not necessary paperwork
  - The licensing settles most of the needed discussions (who owns it, who can do what)
- **Friendly licensing for commercial purposes**
  - Permissive licensing allows commercial exploitation, foreground of partners can be closed
- **You can see what we have and evaluate us in advance**

Many benefits are enabled by open source

Managing Complex Designs

Faster Collaborations

Facilitates Industry/Academia Relationships

Auditable Designs, Reproducible Results

Starting from an open-source ISA: RISC-V

- To enable open-source full HW systems, we need an open-source ISA
- We could not have open-source processors otherwise

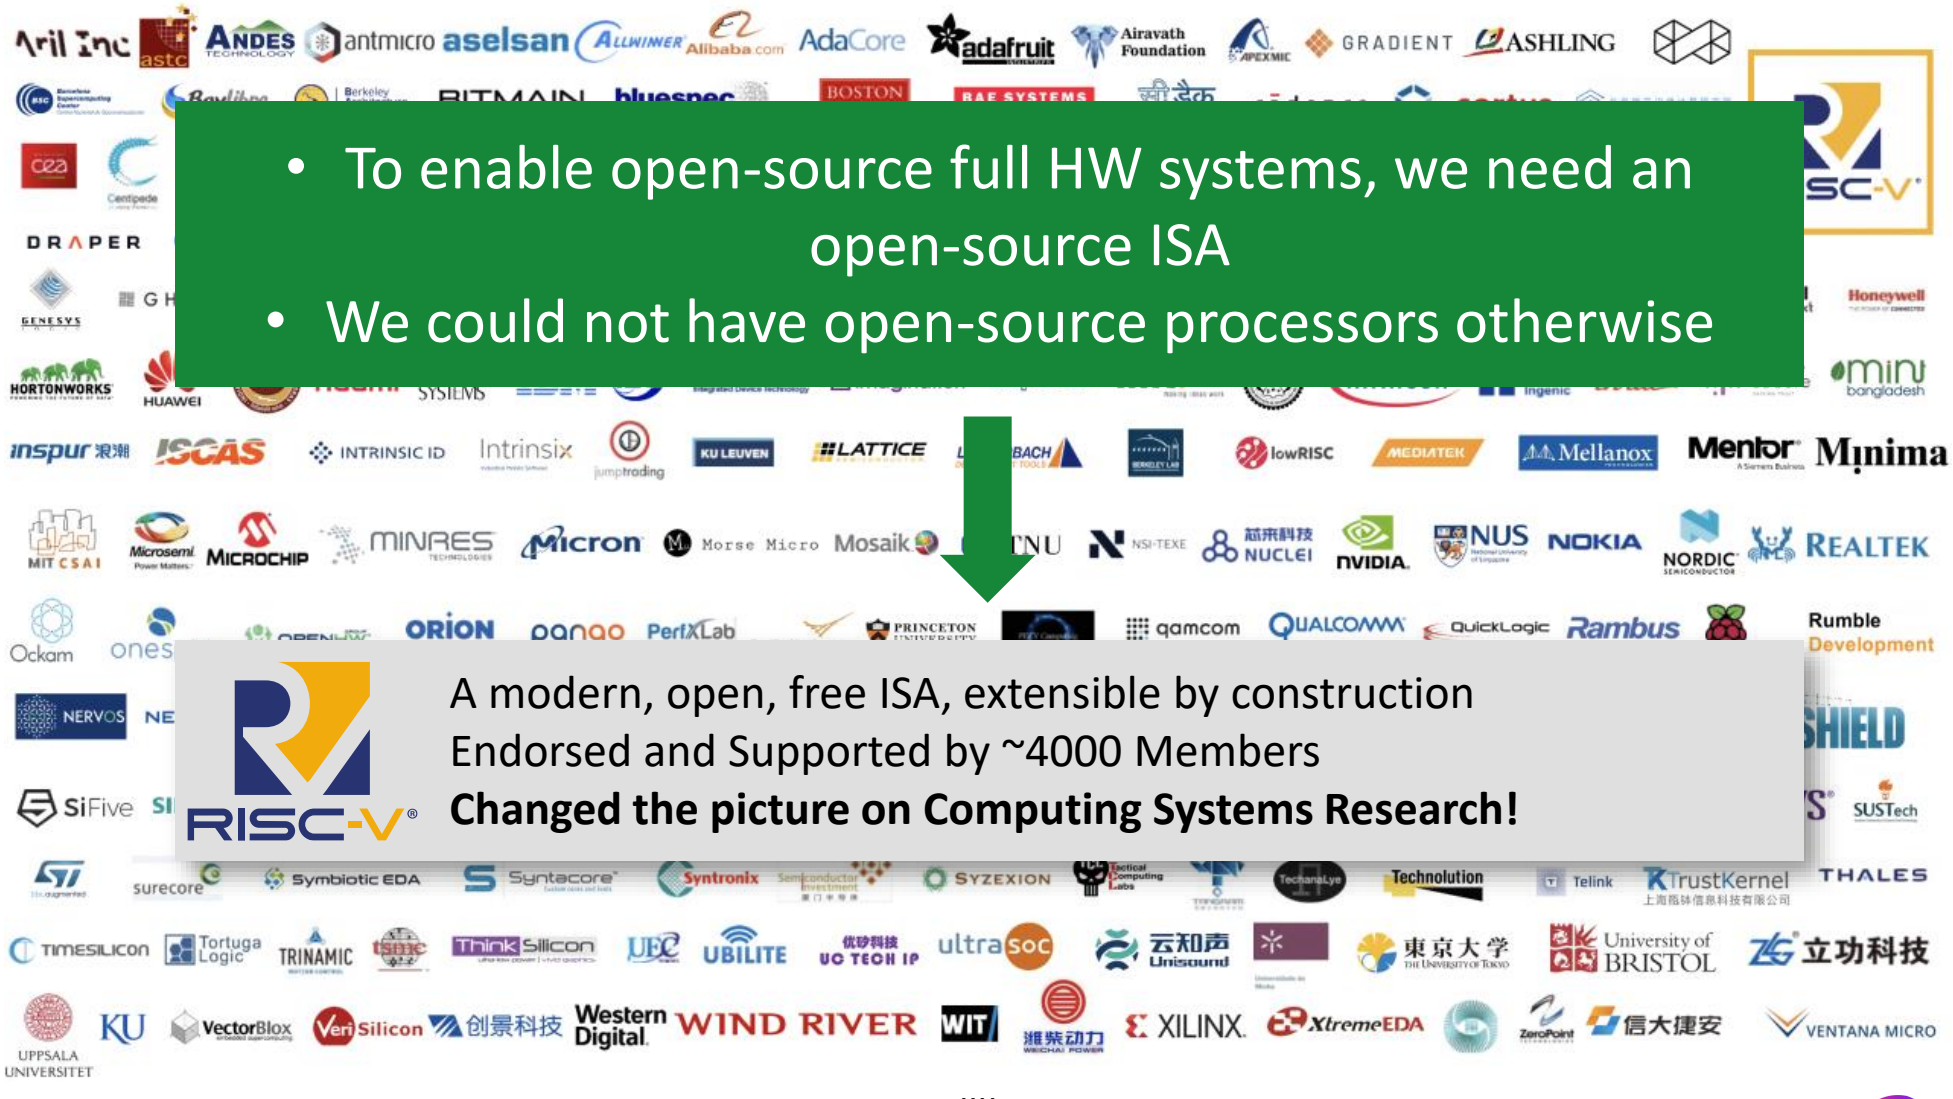

RISC-V Open-Source Software

Huge momentum, extensible, open-source, patent troll safe

- Toolchains → GCC, LLVM

Emulators: QEMU, TinyEMU, Spike, Renode

- System tools … and more

A wide SW ecosystem already available to evaluate and use RISC-V hardware systems

- Language Ru… more
- Operating Systems → Linux: Fedora, OpenSUSE, Gentoo, OpenEmbedded/Yocto, Buildroot, OpenWRT, FreeBSD
  FreeRTOS, Zephyr, RTEMS, Xv6, HelenOS, … and more

https://github.com/riscv/riscv-software-list

What PULP provides is a box of building blocks

RISC-V Cores

RI5CY	Ibex	Snitch	Ariane
32b	32b	32b	64b

Peripherals

JTAG	SPI
UART	I2S
DMA	GPIO

Interconnect

- Logarithmic interconnect
- APB – Peripheral Bus
- AXI4 – Interconnect

Platforms

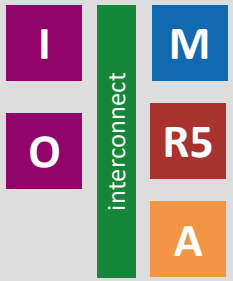

Single Core

- PULPissimo
- Cheshire

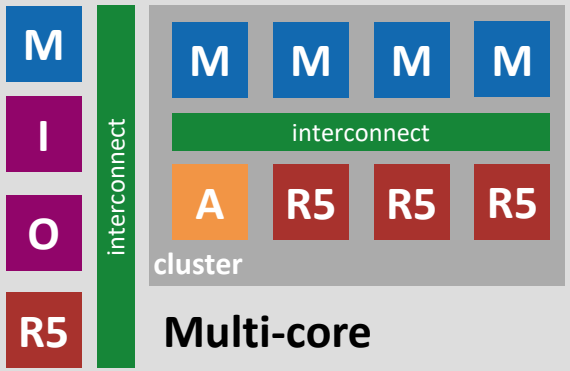

Multi-core

- PULPOpen
- Snitch Cluster

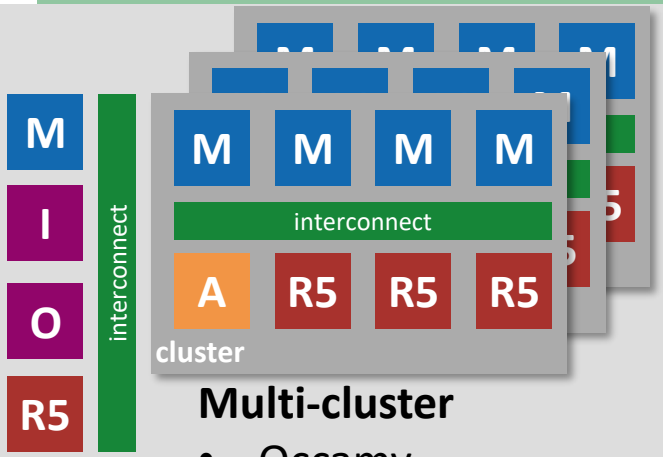

Multi-cluster

- Occamy
- Carfield

IOT → **HPC**

Accelerators

Neureka (Convolution)	RedMulE (TinyML Training)	CUTIE (Ternary DNN Inference)	SNE (Spiking NN)

All of our designs are open-source hardware

- **All our development is on GitHub using a permissive license**
  - HDL source code, testbenches, software development kit, virtual platform

https://github.com/pulp-platform

- Allows anyone to use, change, and make products without restrictions.

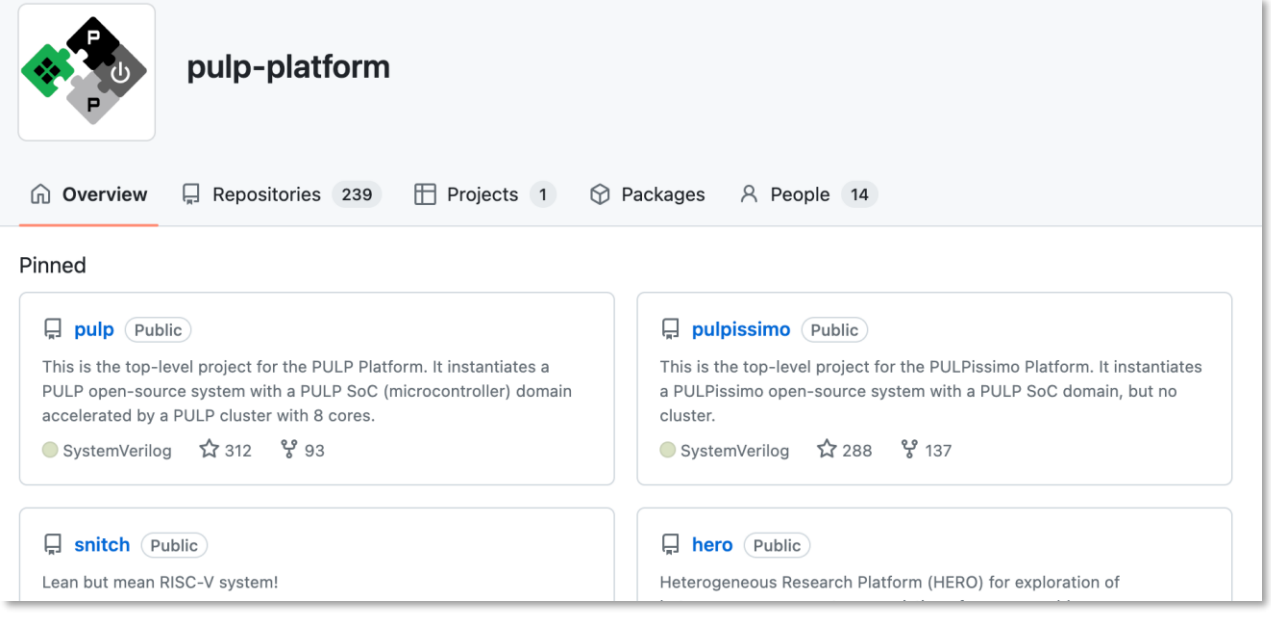

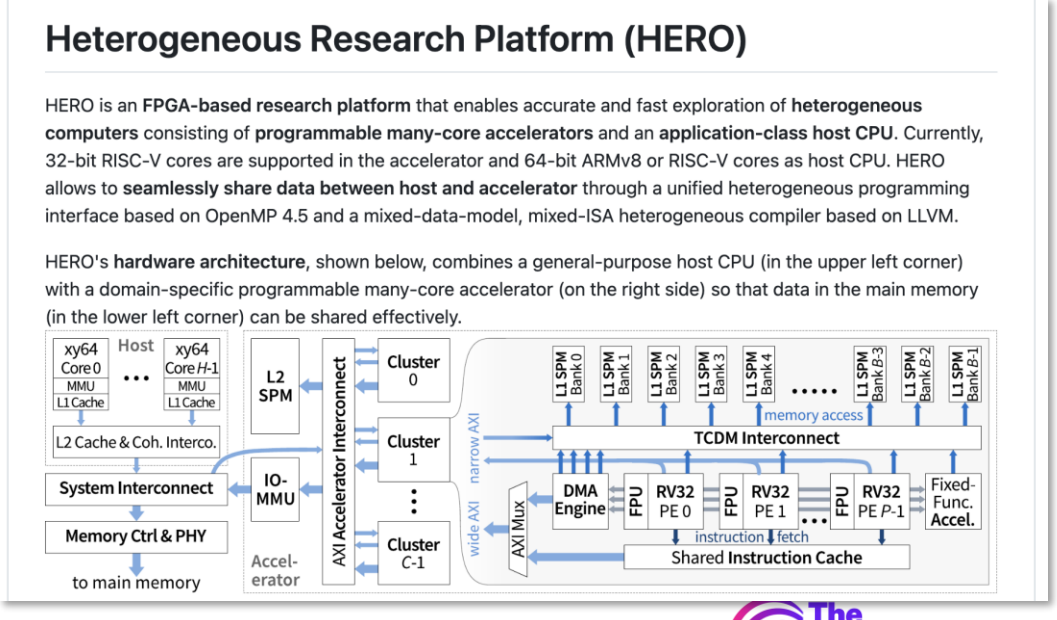

All of our designs are open-source hardware

- **All our development is on GitHub using a perr**
  - HDL source code, testbenches, software development

Many of the GitHub Trending SystemVerilog repositories for the past months are:

- either directly from our group
- or have originated in our group

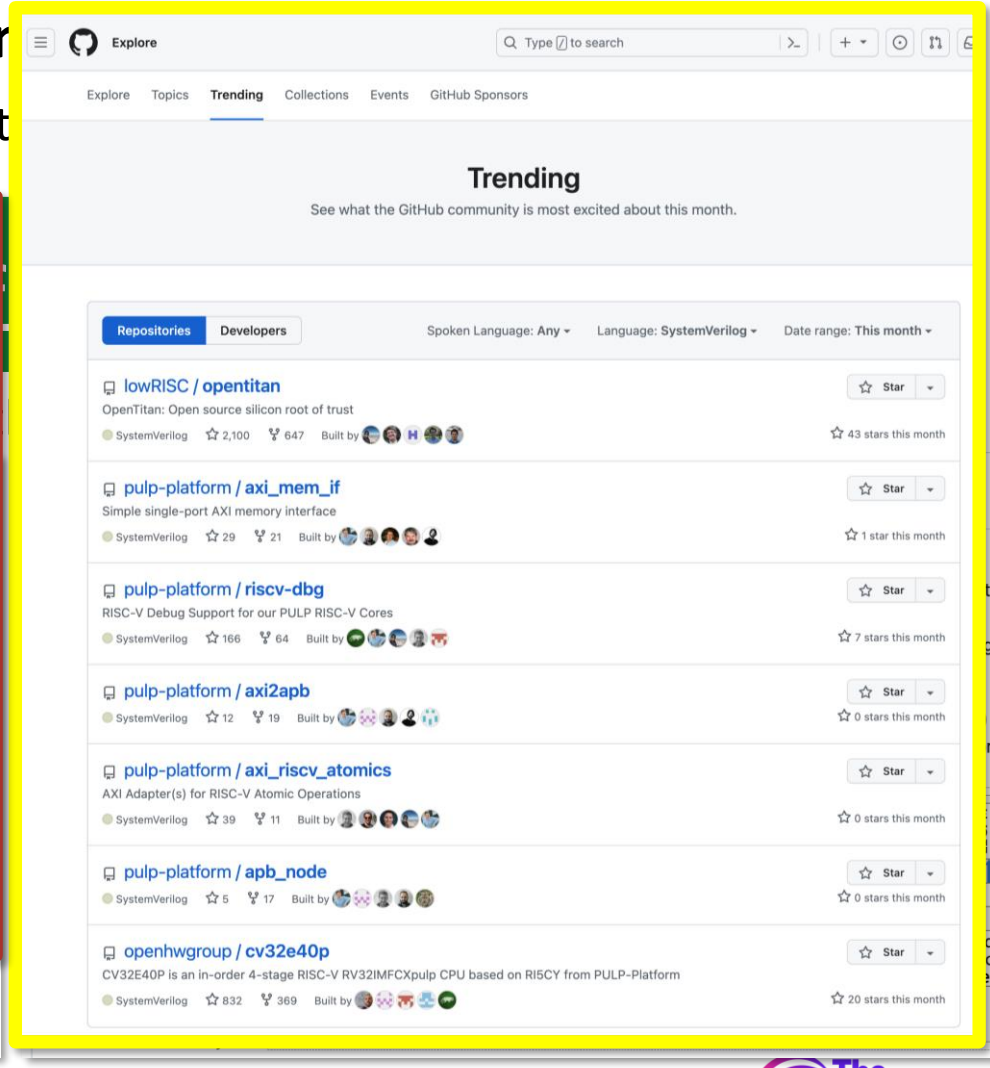

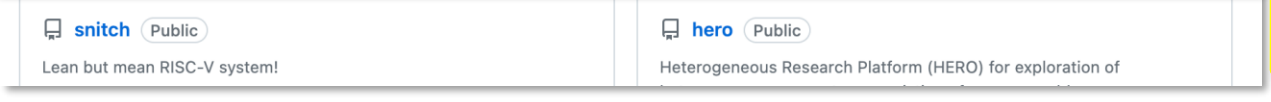

The open model led to successful industry collaborations

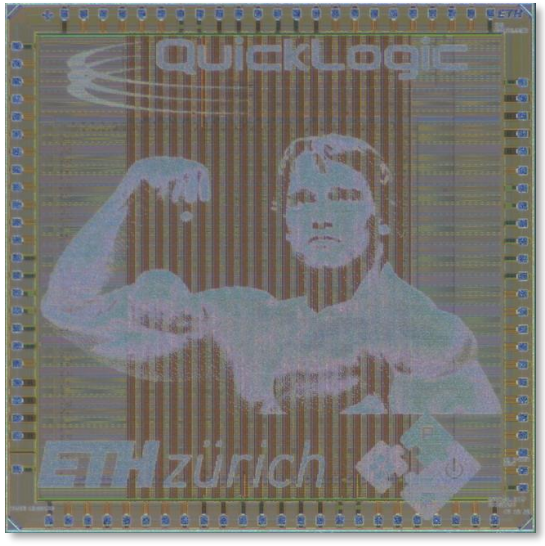

QuickLogic®

Arnold (GF22)

eFPGA with RISC-V core

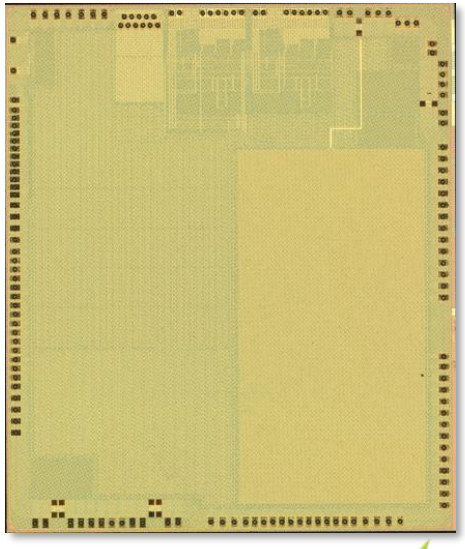

GREENWAVES TECHNOLOGIES

Vega (GF22)

IoT Processor with ML acceleration

Meta numem

Siracusa (TSMC16)

IoT Processor with NVM technology

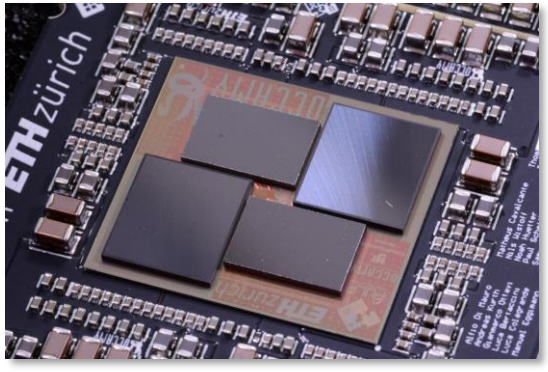

GlobalFoundries™

Occamy (GF12)

Chiplet based ML accelerator with 432 RISC-V cores

And many have used our work for their research

RISC-V week Barcelona 2018

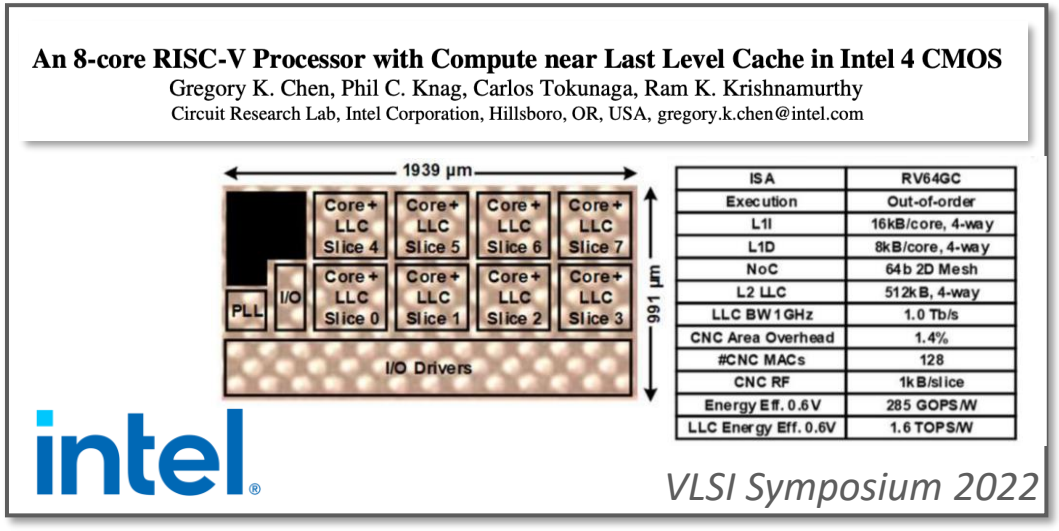

VLSI Symposium 2022

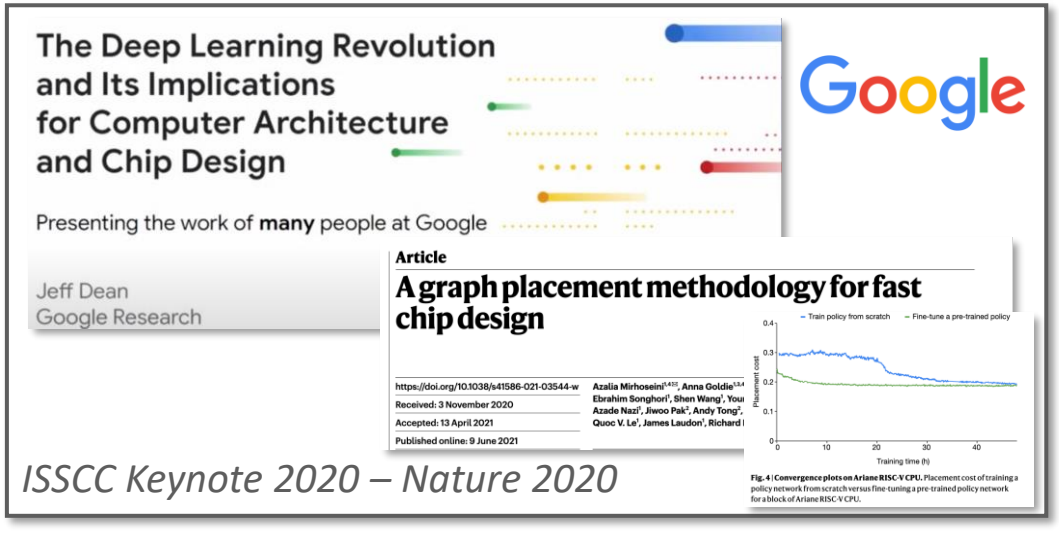

ISSCC Keynote 2020 – Nature 2020

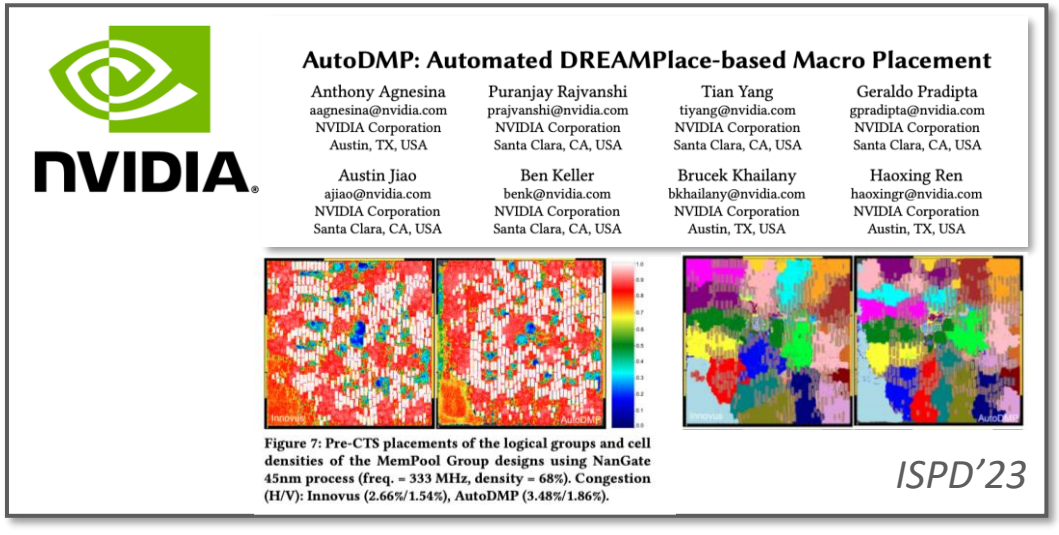

ISPD'23

Some (surprising) side effects of open-source hardware

- **There is a surprising amount of bureaucracy involved**
  - Code Ph.D. students/staff develop belongs to the university (they pay us)
  - Master/semester thesis students own the work they produce
  - Need to get proper approval for everyone involved

- **Most agreements with companies are not meant for open source**
  - Instead of paying for exclusive IP, we need sponsoring agreements
  - Important to make sure we do not sign anything that binds our open-source effort

From Open Hardware...

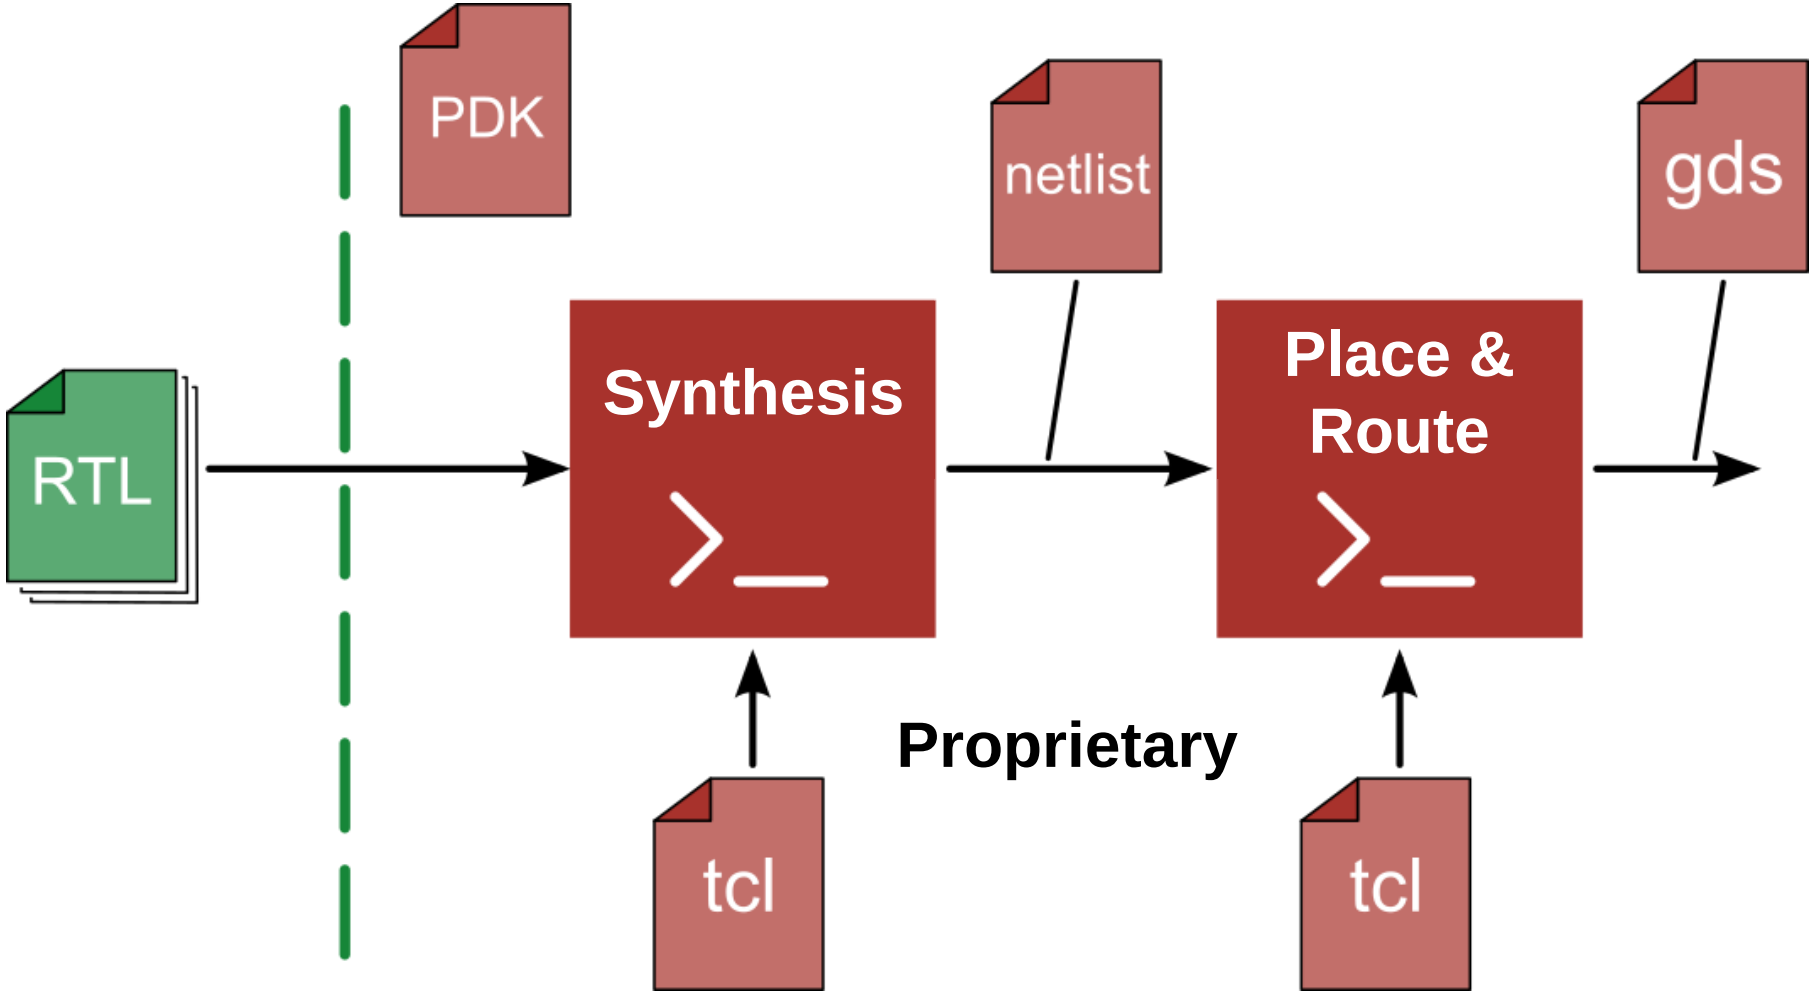

…to fully open-source IC design

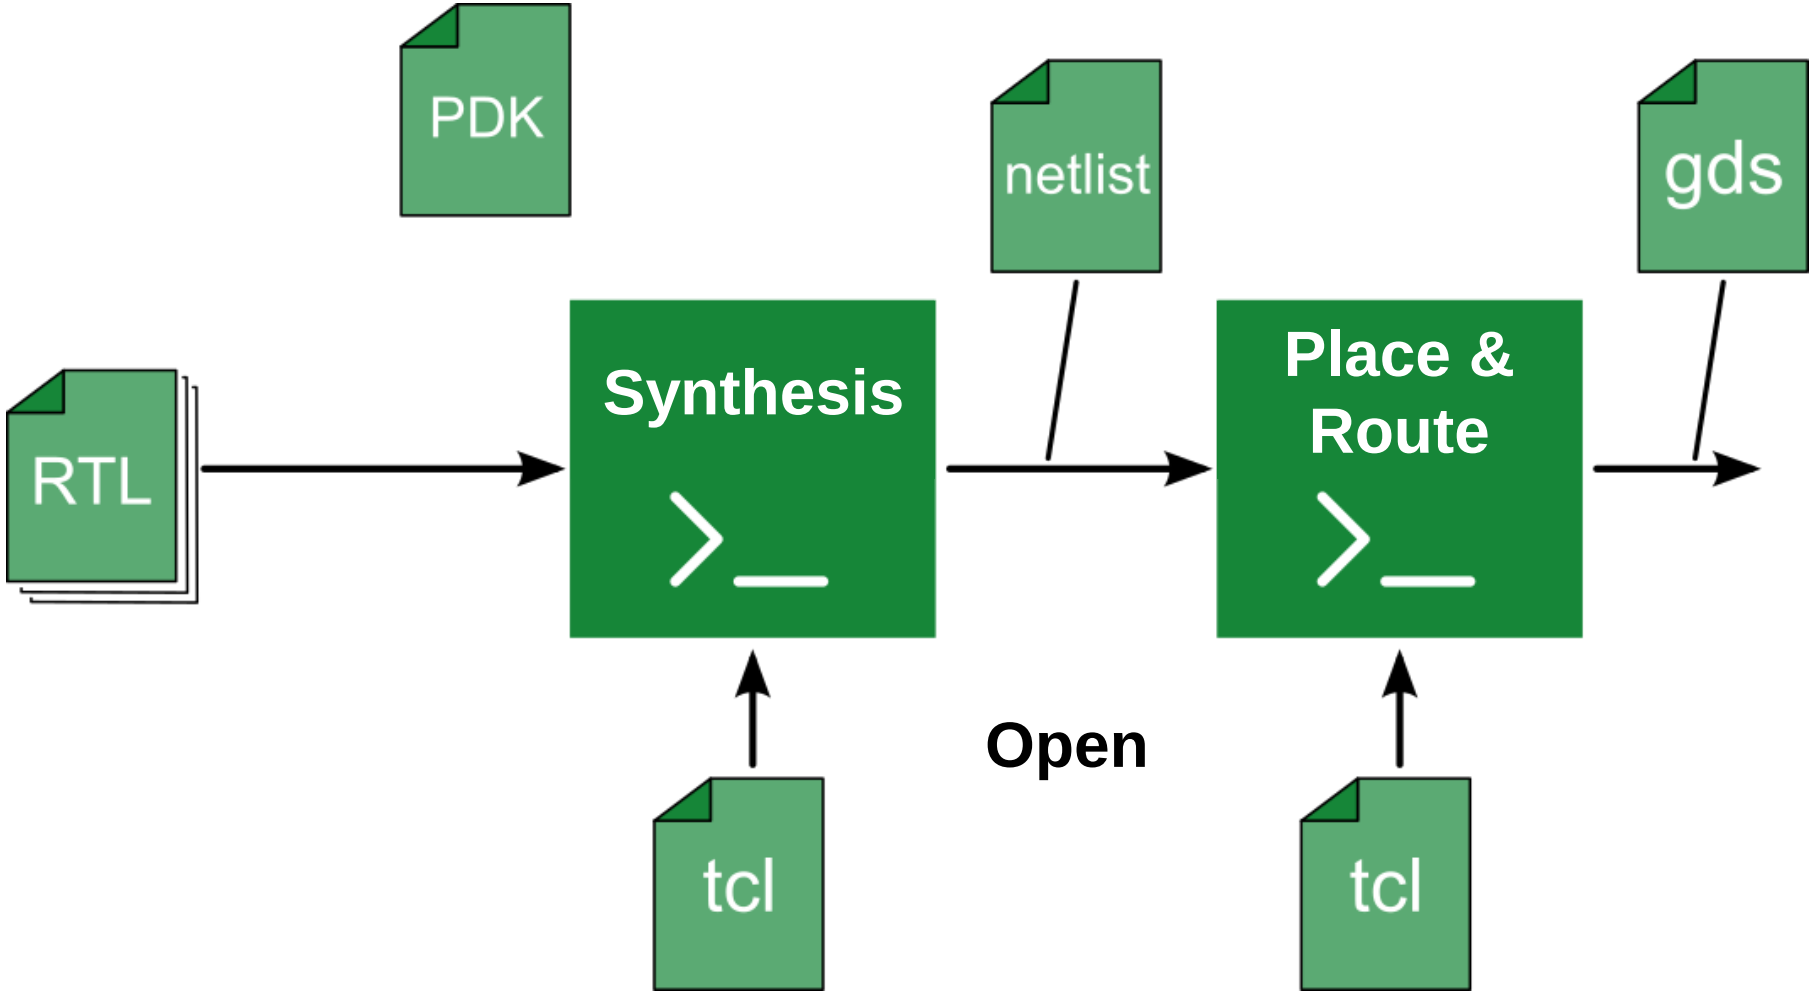

Benefits of end-to-end openness

- **Research**
  - Easier collaboration (no NDAs)
  - Reproducing results
  - New research (using tools or data)
- **Education**
  - Increased access
  - Experiment with flows and tools
  - No black boxes, full transparency
- **Industry**
  - Transparent chain of trust
  - Lower initial cost

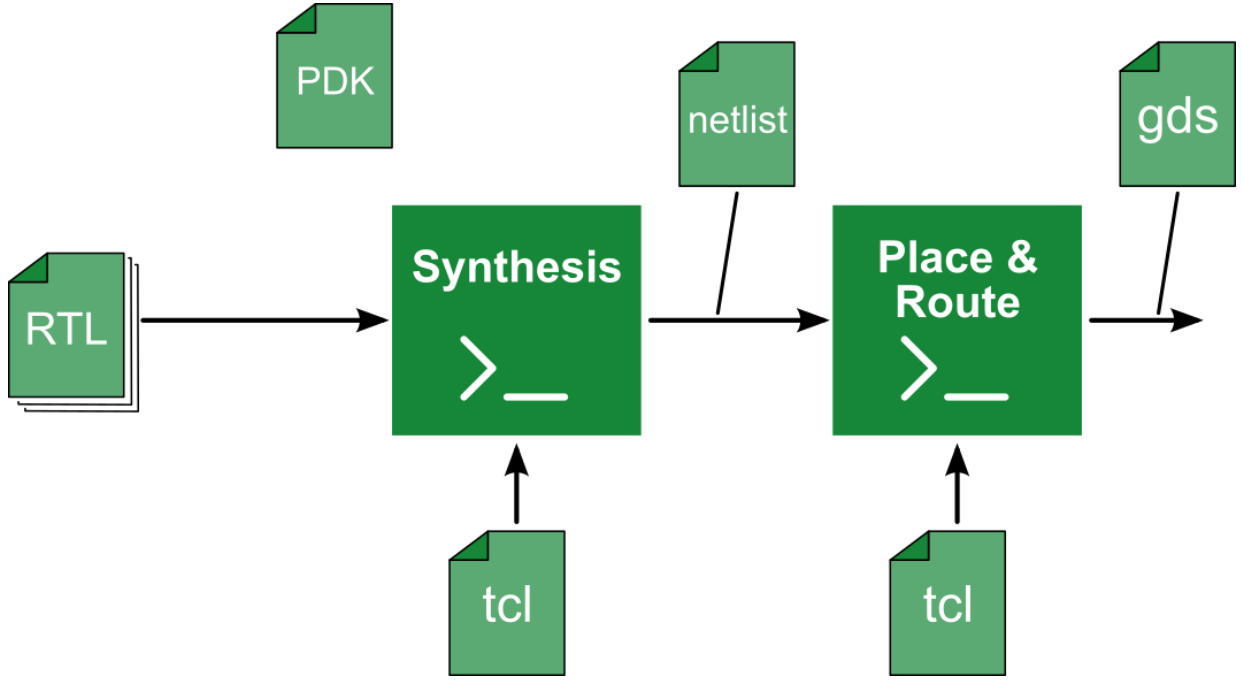

Open-source PDKs are the key

- **Physical designs will contain technology relevant information**
  - Nothing can be released (even if there were perfect open EDA tools) without technology info
  - So no open PDK no open source releases beyond RTL/HLS code

- **IPs containing proprietary information on technology cannot be released**
  - Open PDKs will allow us to change this

- **Some good progress, more is needed**
  - Skywater (130nm), IHP (130nm)
  - A 40nm/65nm open PDK would be a game changer, many viable products could be designed

Open-source EDA tools

- **Open EDA tools work well and are improving fast**
  - ETHZ is actively collaborating with UCSD
  - Currently working with Yosys (for synthesis) and OpenRoad (for PnR)
- **Will not replace commercial EDA**
  - A gap to PPA, service, support will remain
- **Open source EDA will allow you to make your own chips**
  - Good for many applications/needs
  - Will be great for teaching
- **Summer school at ETH Zürich**
  - June 3rd – 7th

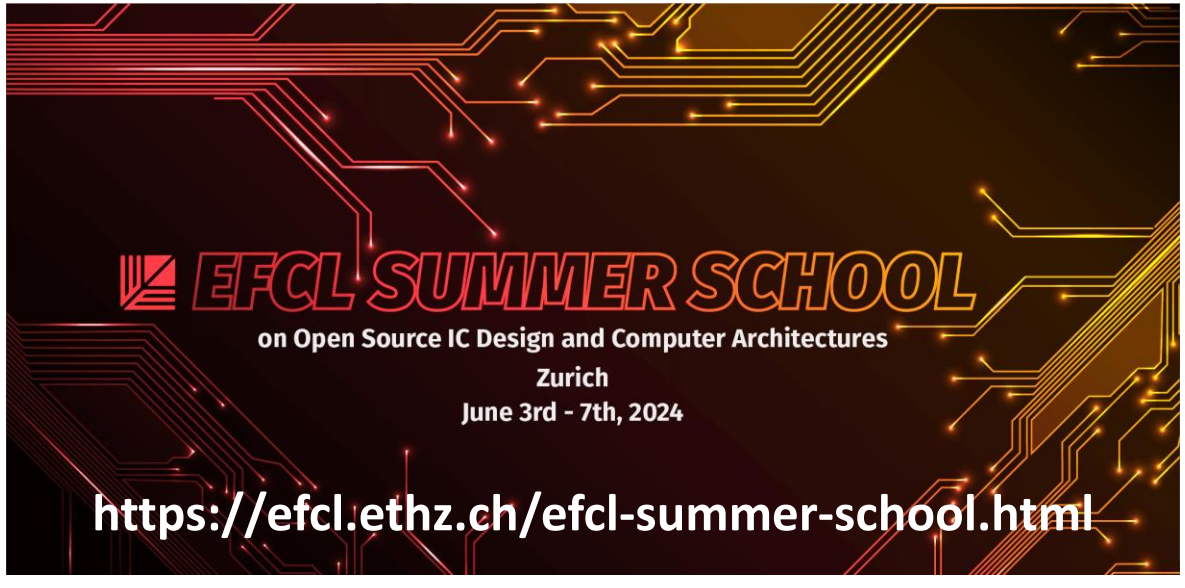

Fully open IC design

System architecture

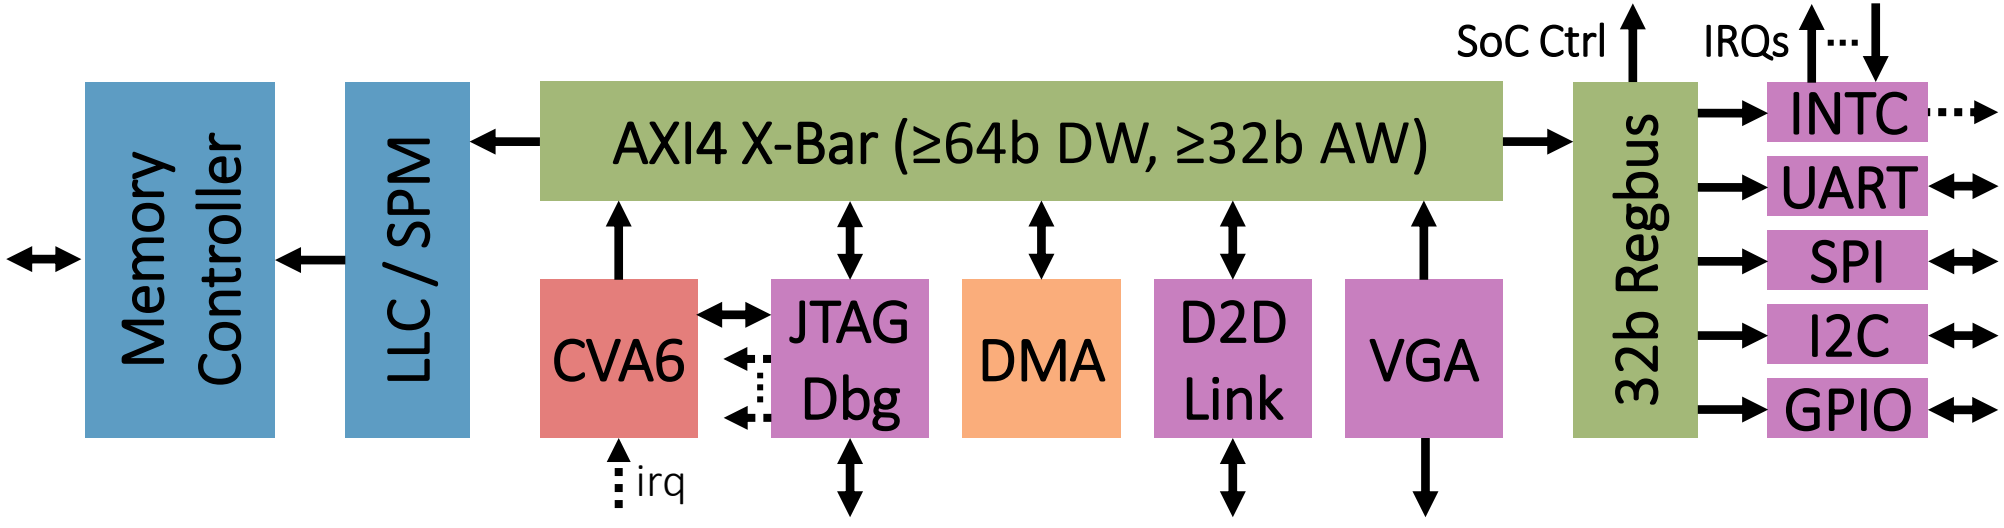

- An open-source Linux-capable RISC-V MPU (based on https://github.com/pulp-platform/cheshire)
- Including a 64-bit RISC-V core (CVA6 -https://github.com/openhwgroup/cva6)
- About 1 MGE of logic

Open-source tool-flow

- **Open-Source PDK (IHP 130nm)**

- **Pre-Process**
  - Simplify SystemVerilog
  - Convert to Verilog
- **Synthesis**
  - Yosys from RTL to generic-cells
  - Calls ABC for logic optimization and mapping
- **Place-and-Route**
  - Collection of research tools into OpenRoad

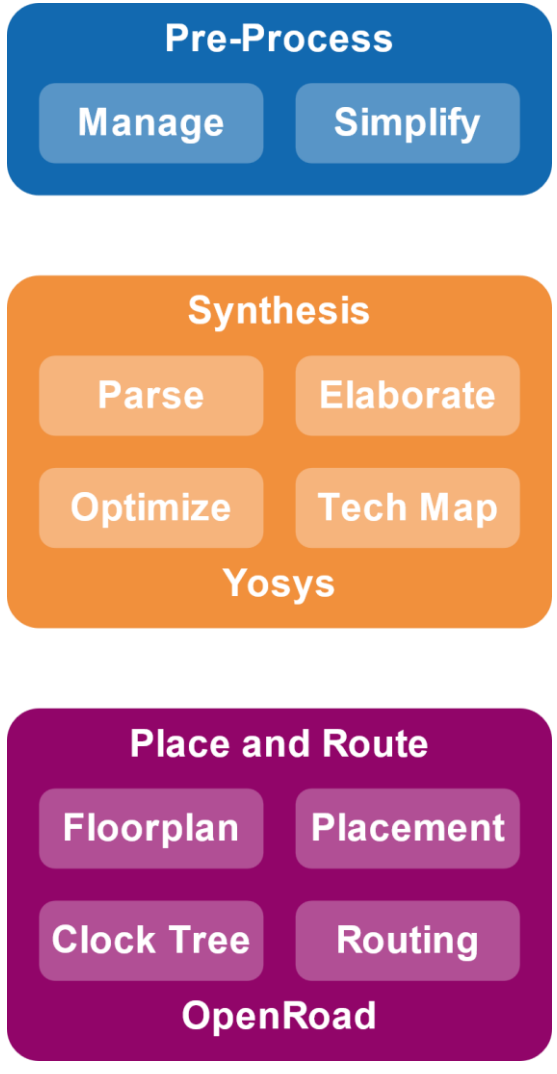

Massive PPA improvements in 6 months

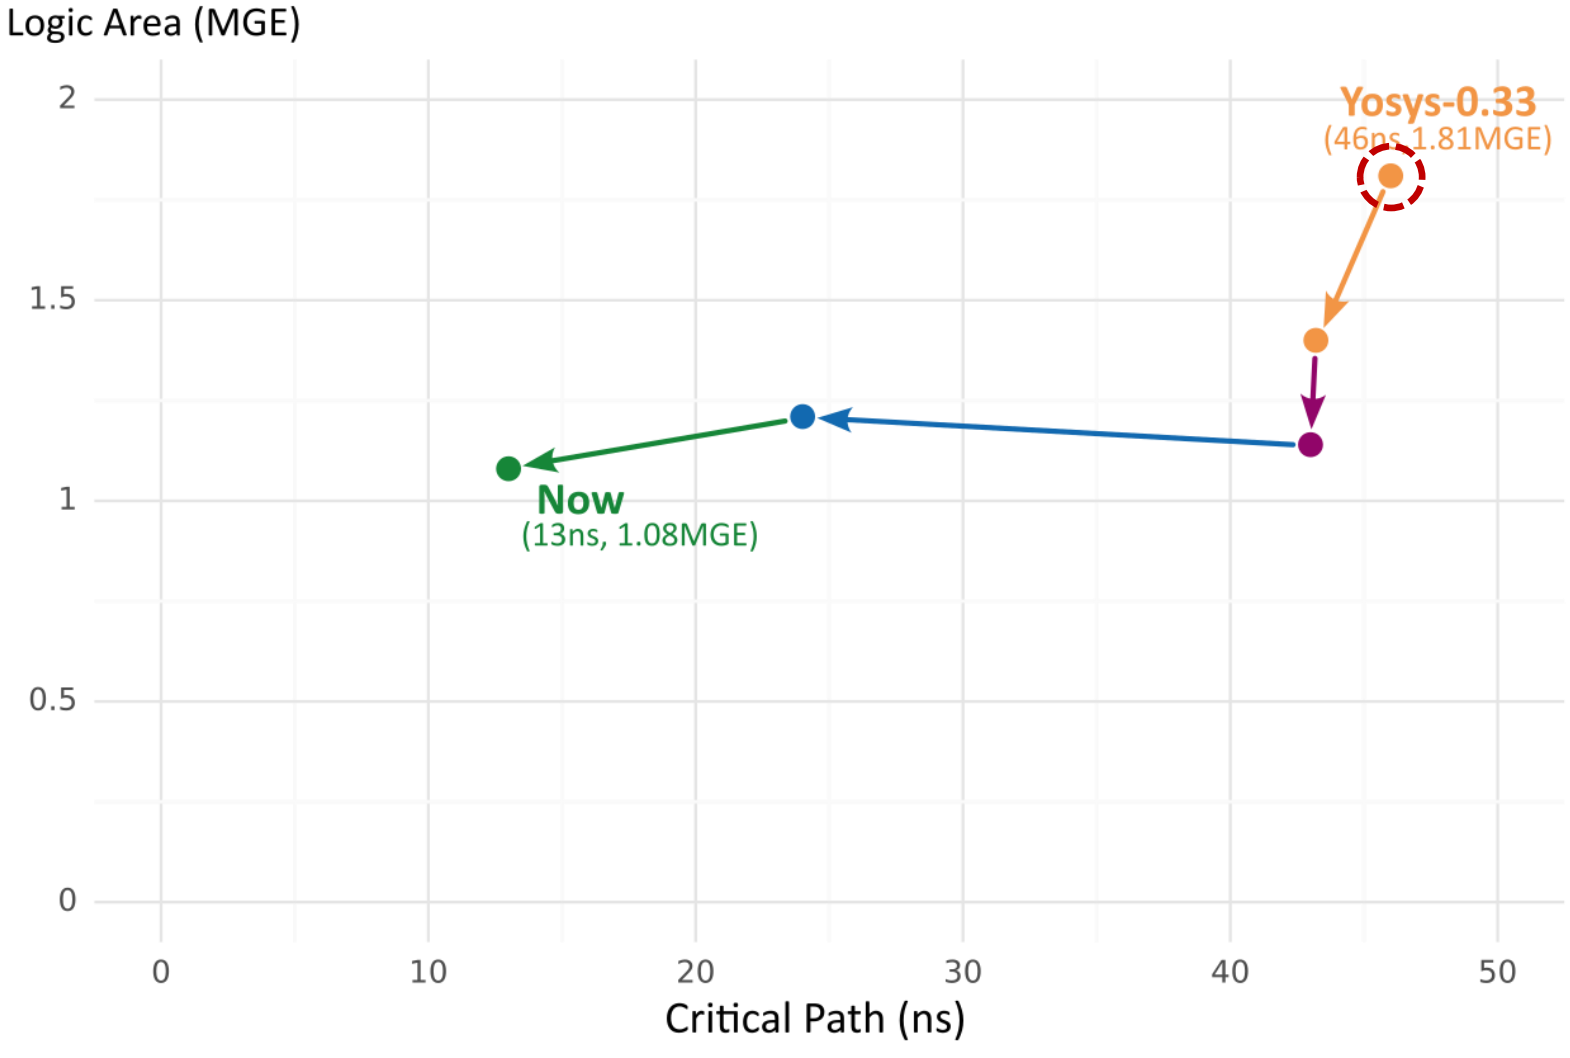

IHP130, tt-corner; Yosys-0.37 with optimized script if not mentioned otherwise

Massive PPA improvements in 6 months

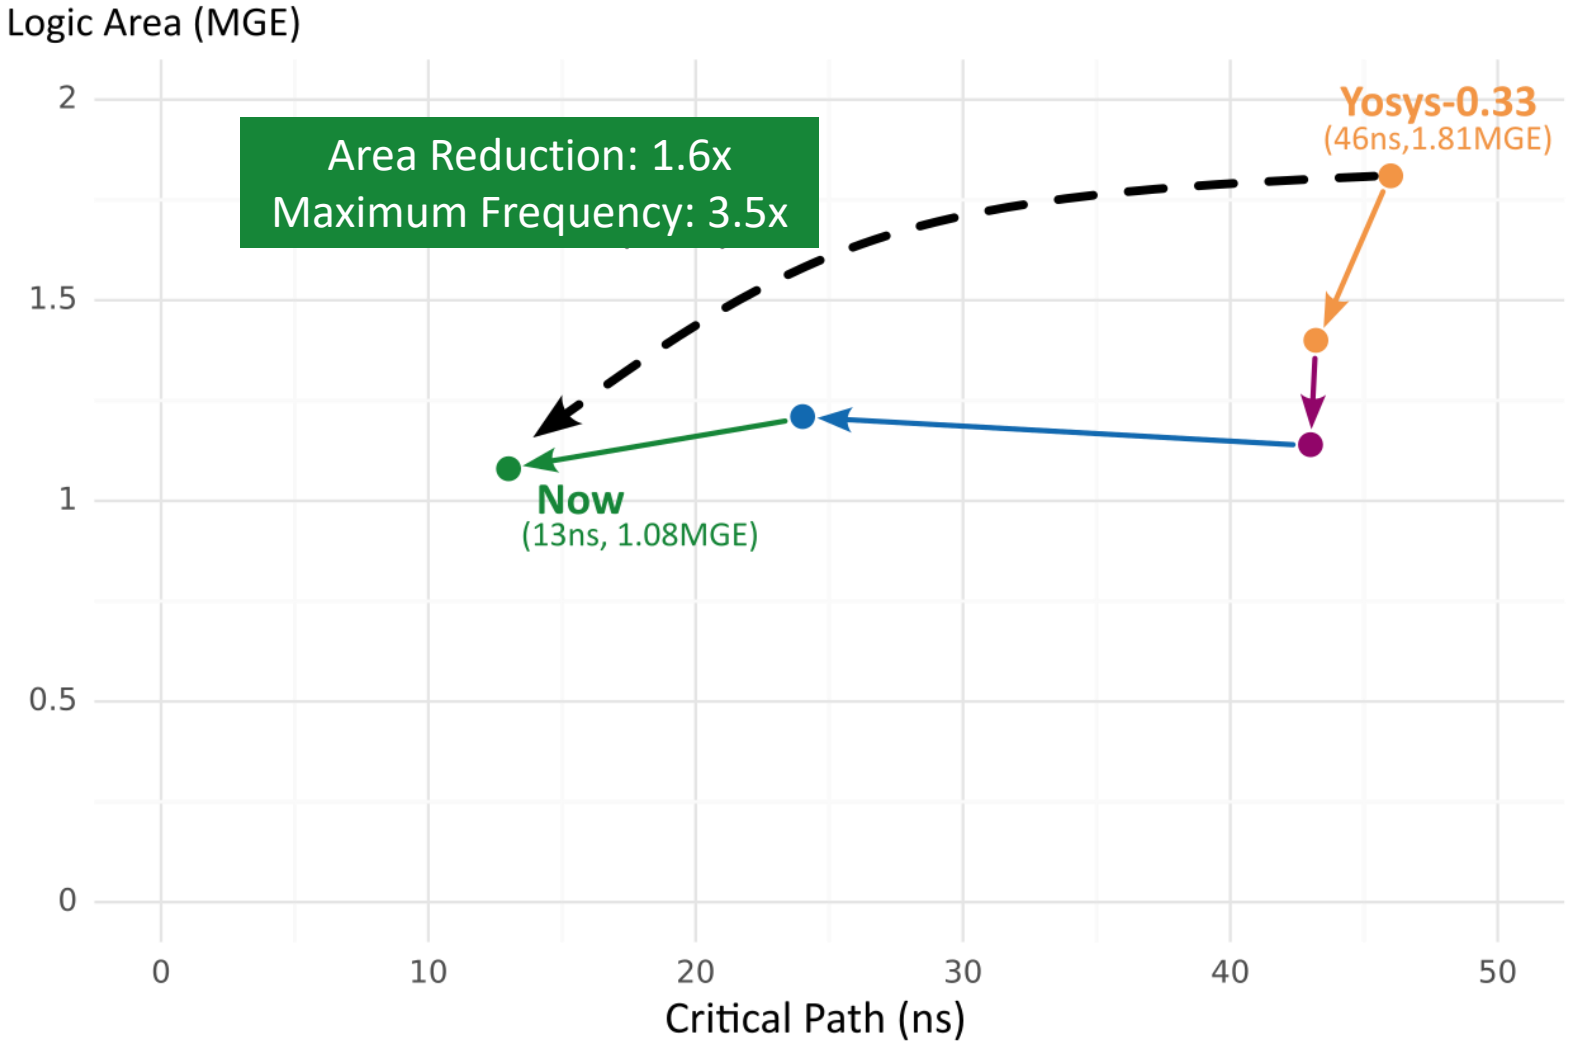

IHP130, tt-corner; Yosys-0.37 with optimized script if not mentioned otherwise

Massive PPA improvements in 6 months

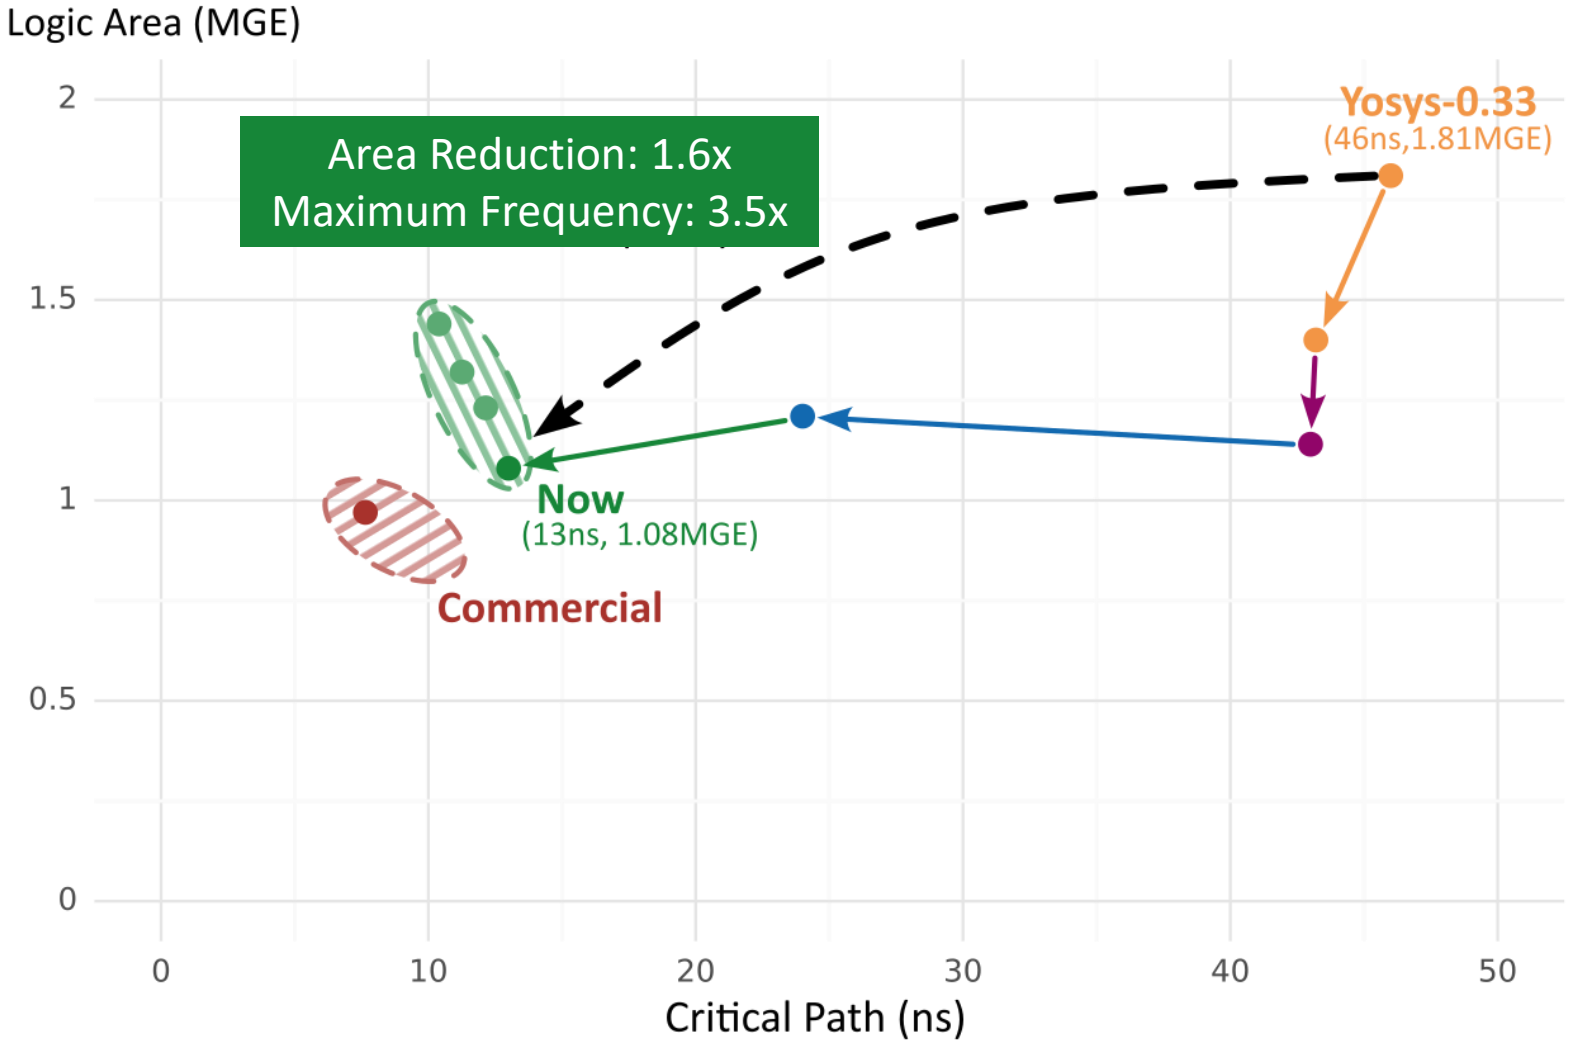

IHP130, tt-corner; Yosys-0.37 with optimized script if not mentioned otherwise

Basilisk: fully open SoC design (RTL to GDS)

- **Fully open flow**
  - GDS will be submitted in May 2024
- **Designed using IHP 130nm**
  - Open-source PDK
- **Collaboration with all main open-source EDA tool developers**
- **Key Metrics:**
  - 1.08MGE Logic
  - 13ns critical path
  - 6.25x5.5mm (34mm²)
  - **14.5h Runtime** (+50% vs commercial)*

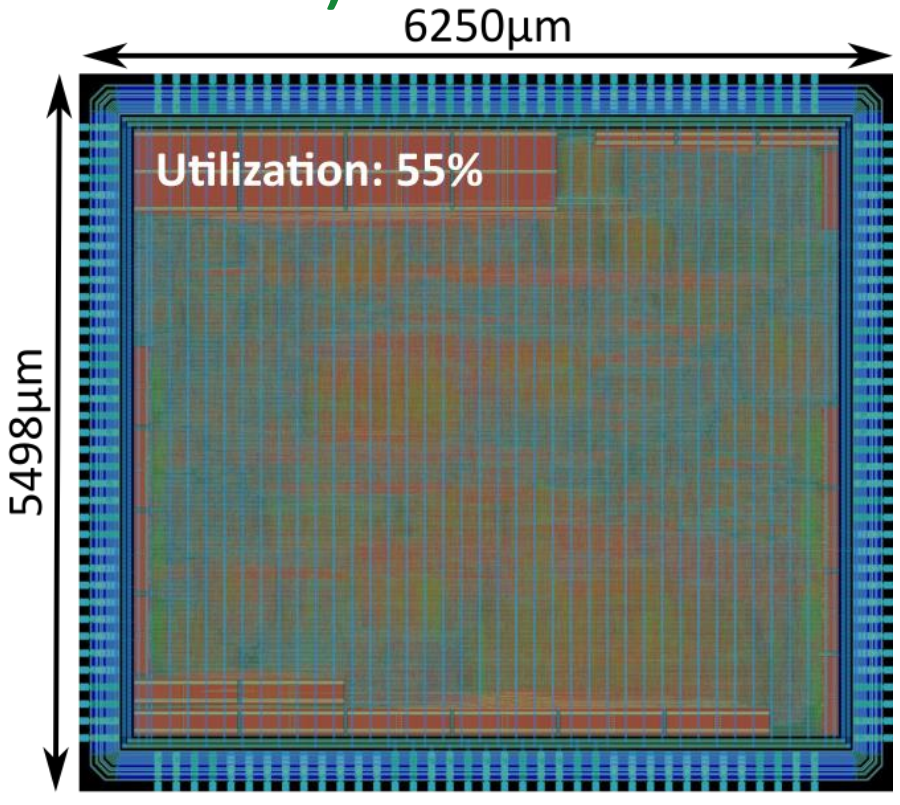

*2x 2.5GHz XeonE5-2670(10Core), max thread-count 20

In a nutshell: why open-source hardware & silicon?

- **It is a necessity**
  - We can not afford to make everything ourselves, we need to collaborate
  - Makes it possible to work together quickly
  - Your results are more trustworthy, anybody can verify it!
- **It works**
  - We have more projects, and more funding due to our open-source activities
  - We were able to start many interesting and fruitful collaborations
- **It helps others as well**
  - Many companies, universities, individuals are using pieces of PULP
  - There is already significant commercial use, a lot we don’t even know about

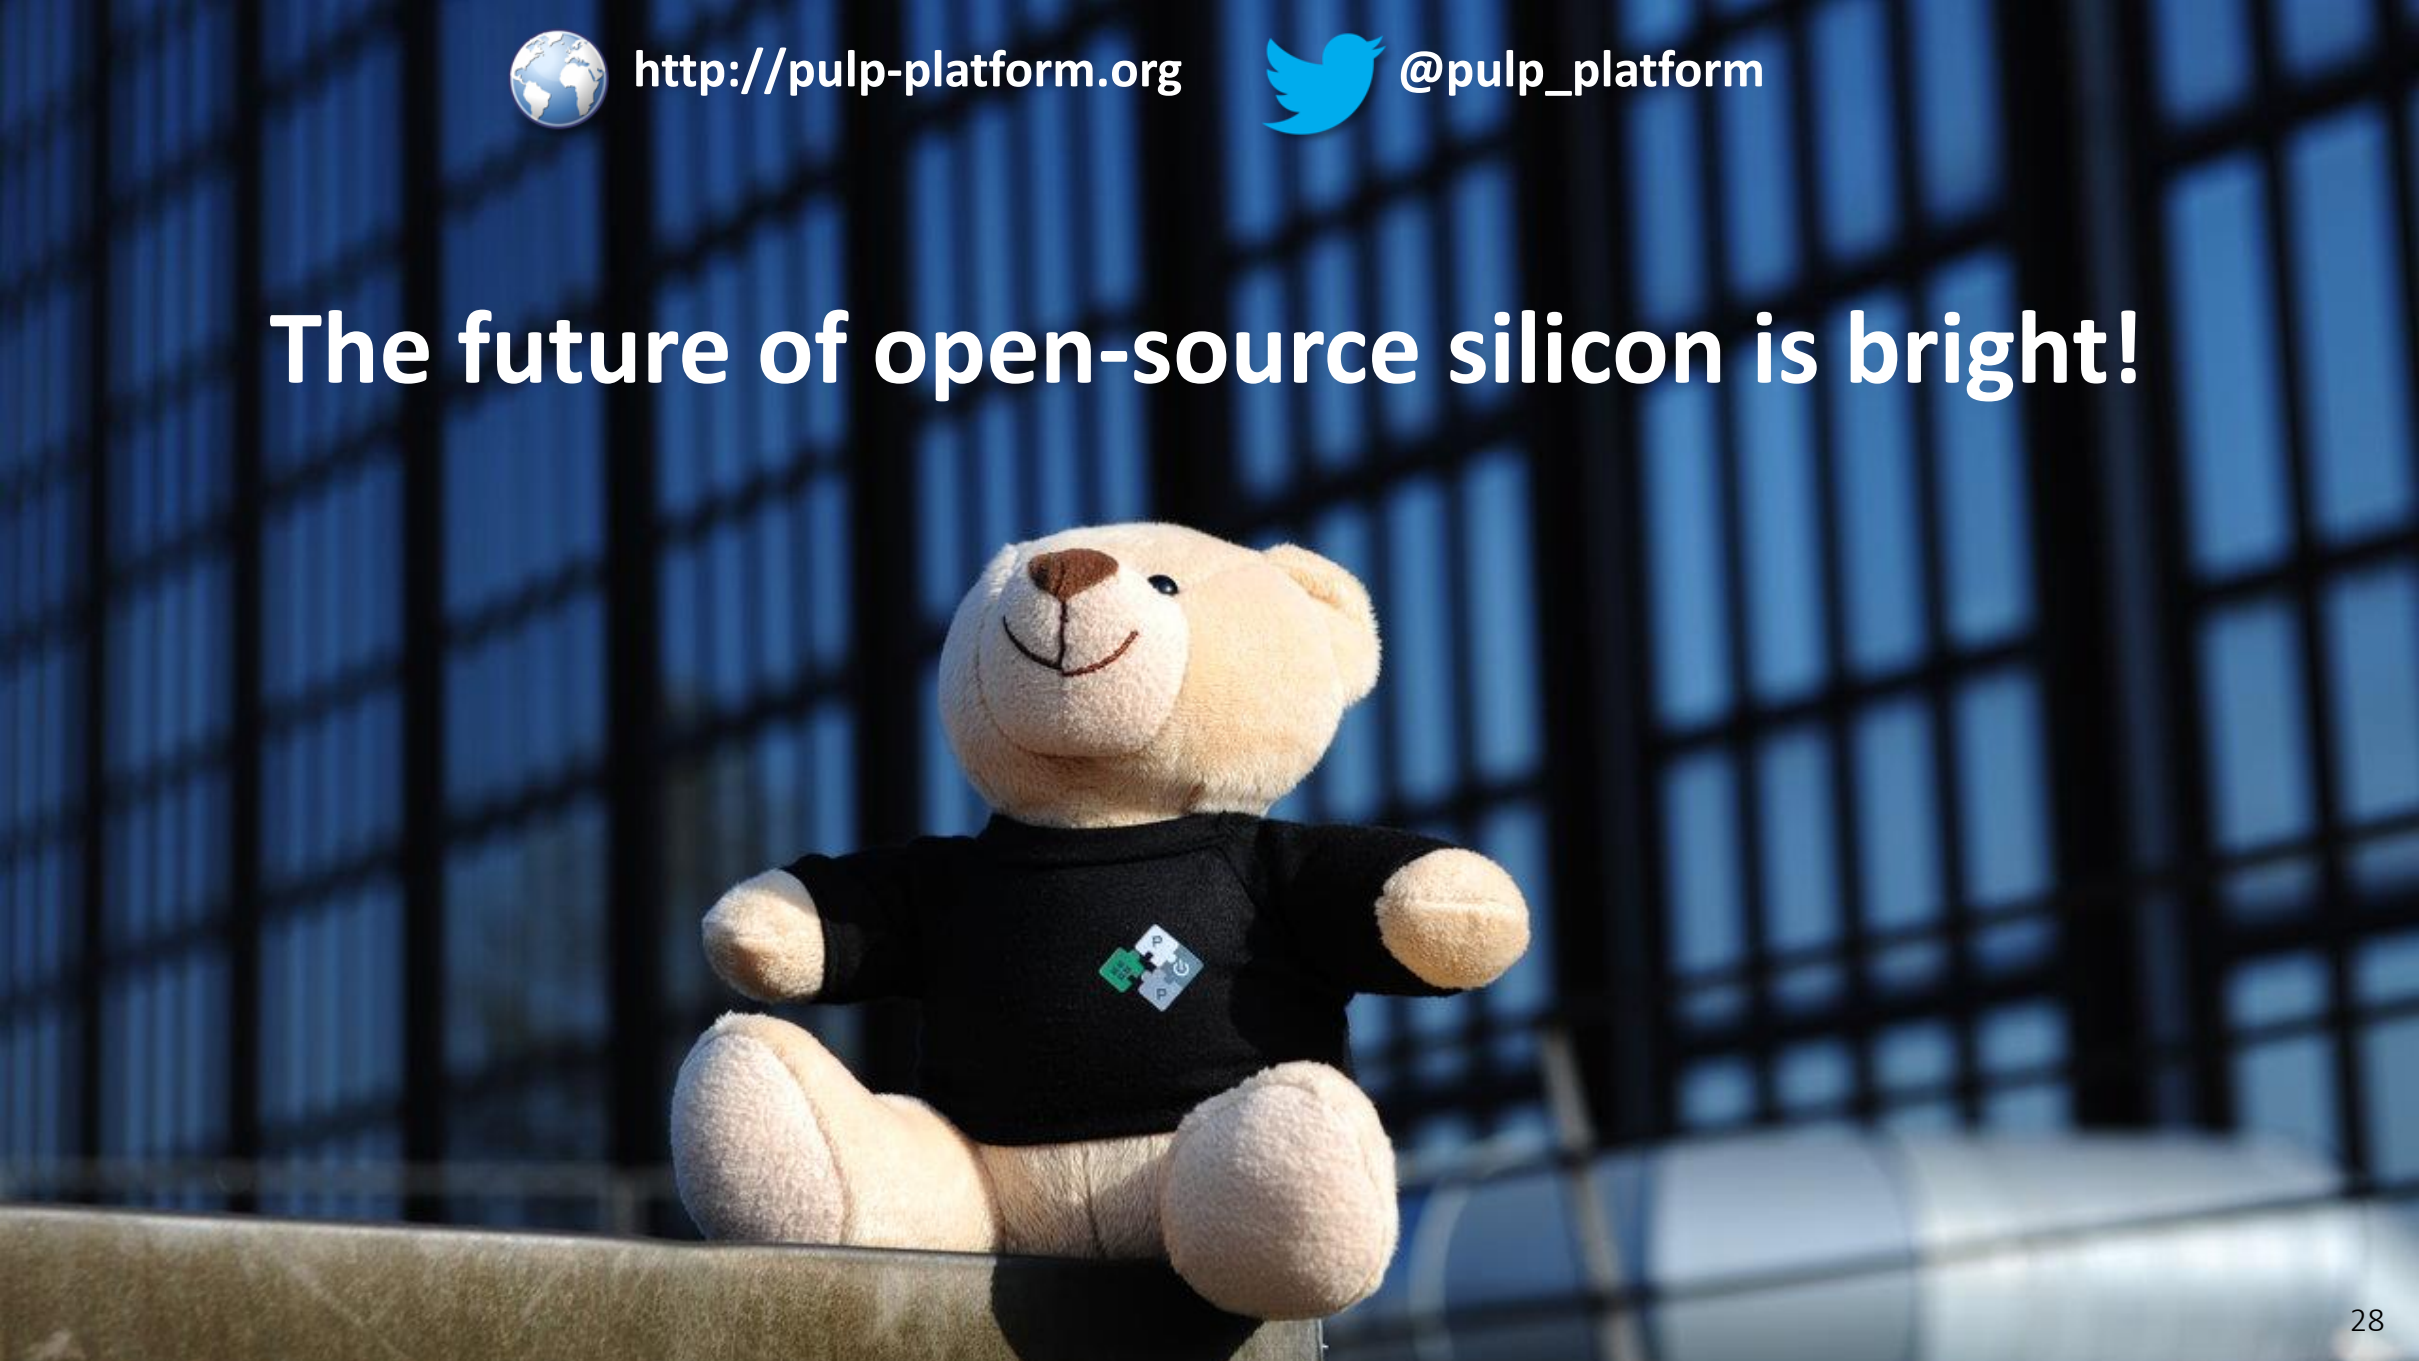
http://pulp-platform.org
@pulp_platform
The future of open-source silicon is bright!