50+ ASICs in 10 years: a visual history

Integrated Systems Laboratory (ETH Zürich)

Frank K. Gürkaynak kgf@iis.ee.ethz.ch

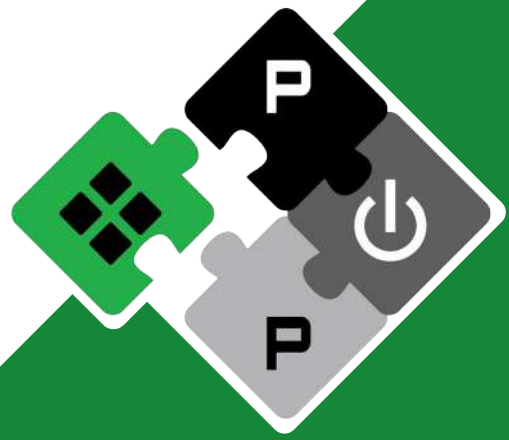

@pulp_platform

pulp-platform.org

youtube.com/pulp_platform

PULP Platform

Open Source Hardware, the way it should be!

We have used PULP to design and test 50+ ASICs

- **Our experience in ASIC design has been instrumental in our success**
  - We **learned from** this **experience** (more than we realize)
  - **Helped collaborations**, especially with Industry
  - PULP based solutions ended up in **actual products**

- **Check out our chip gallery: http://asic.ethz.ch/**

For a research group, this level of output is unprecedented

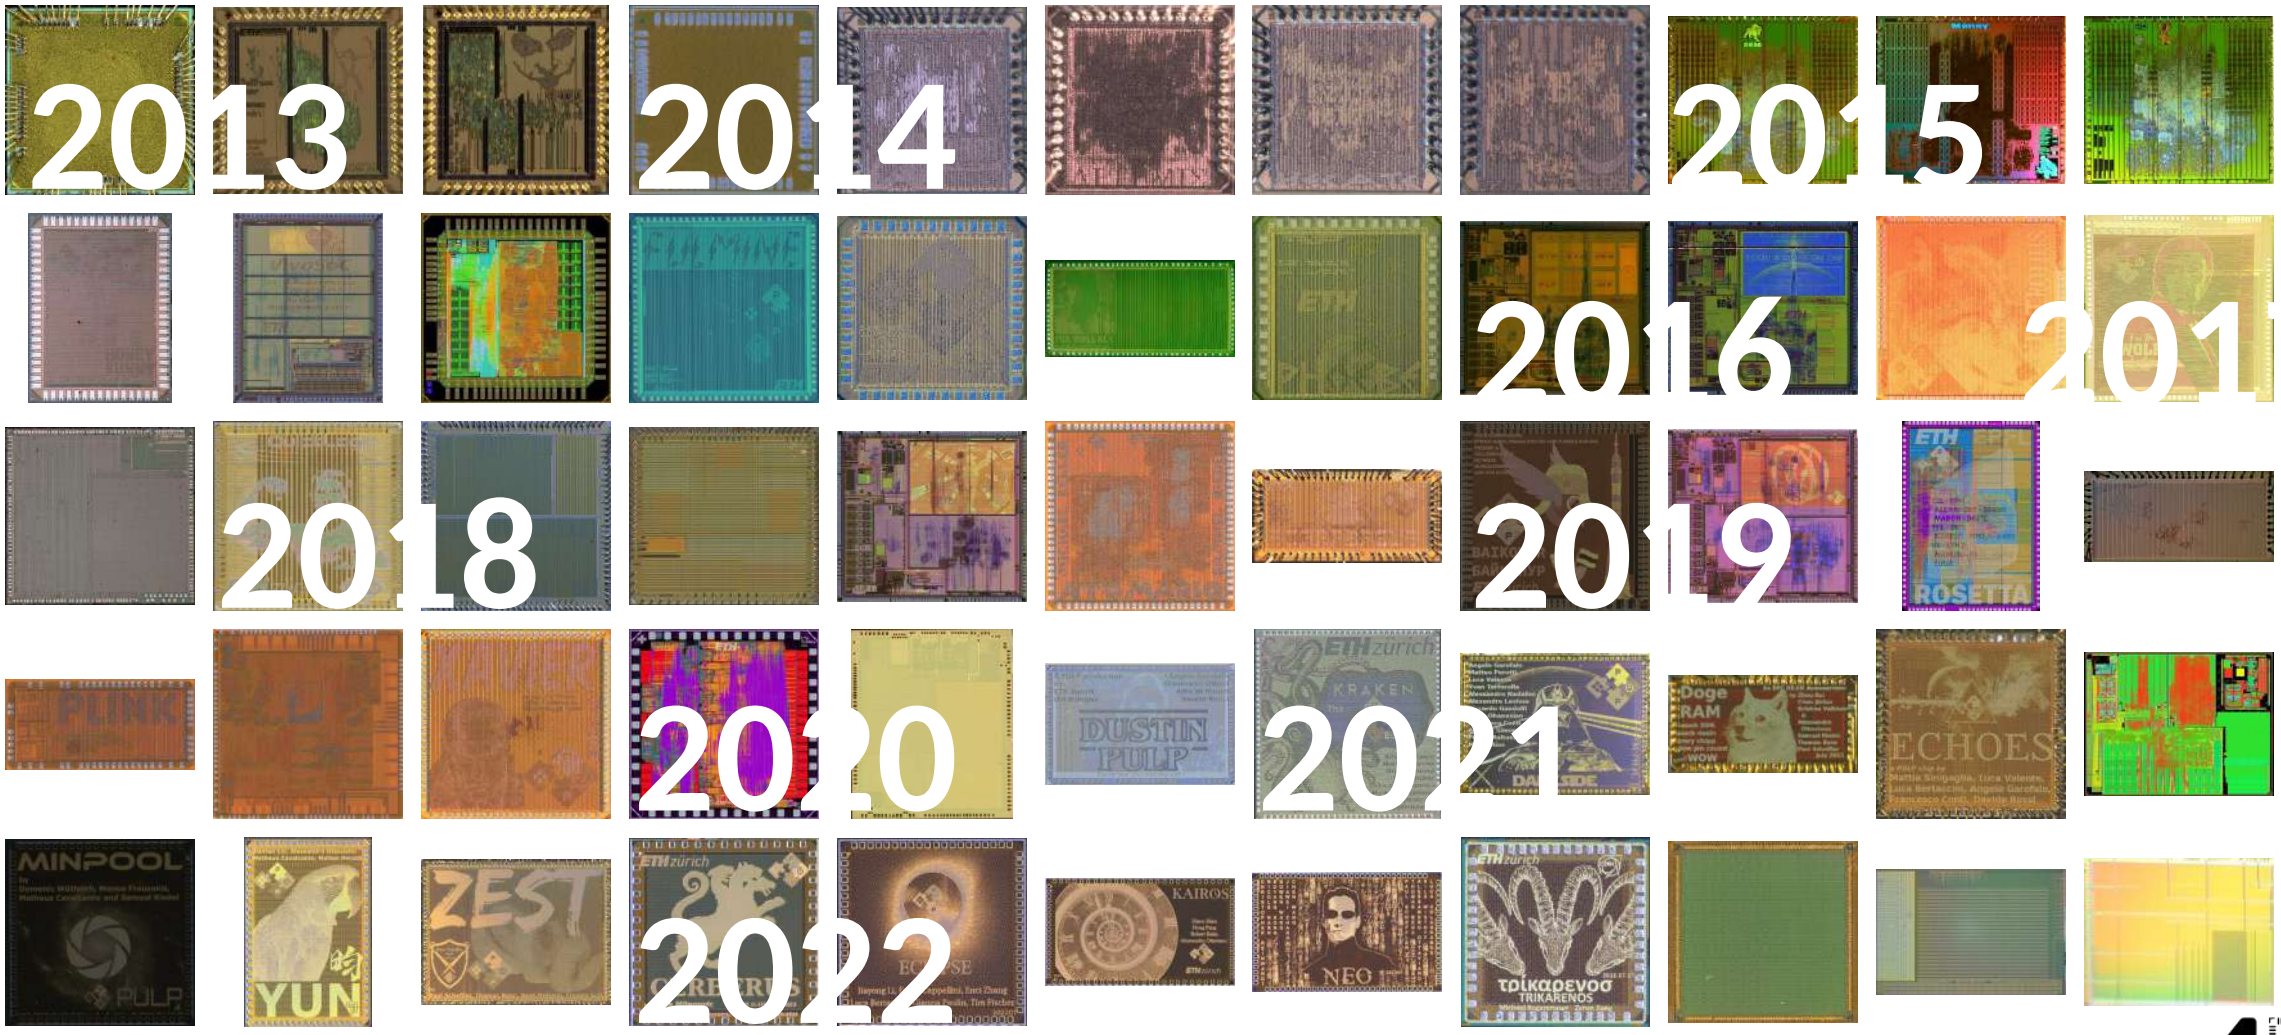

Most of these ASICs resulted in major publications

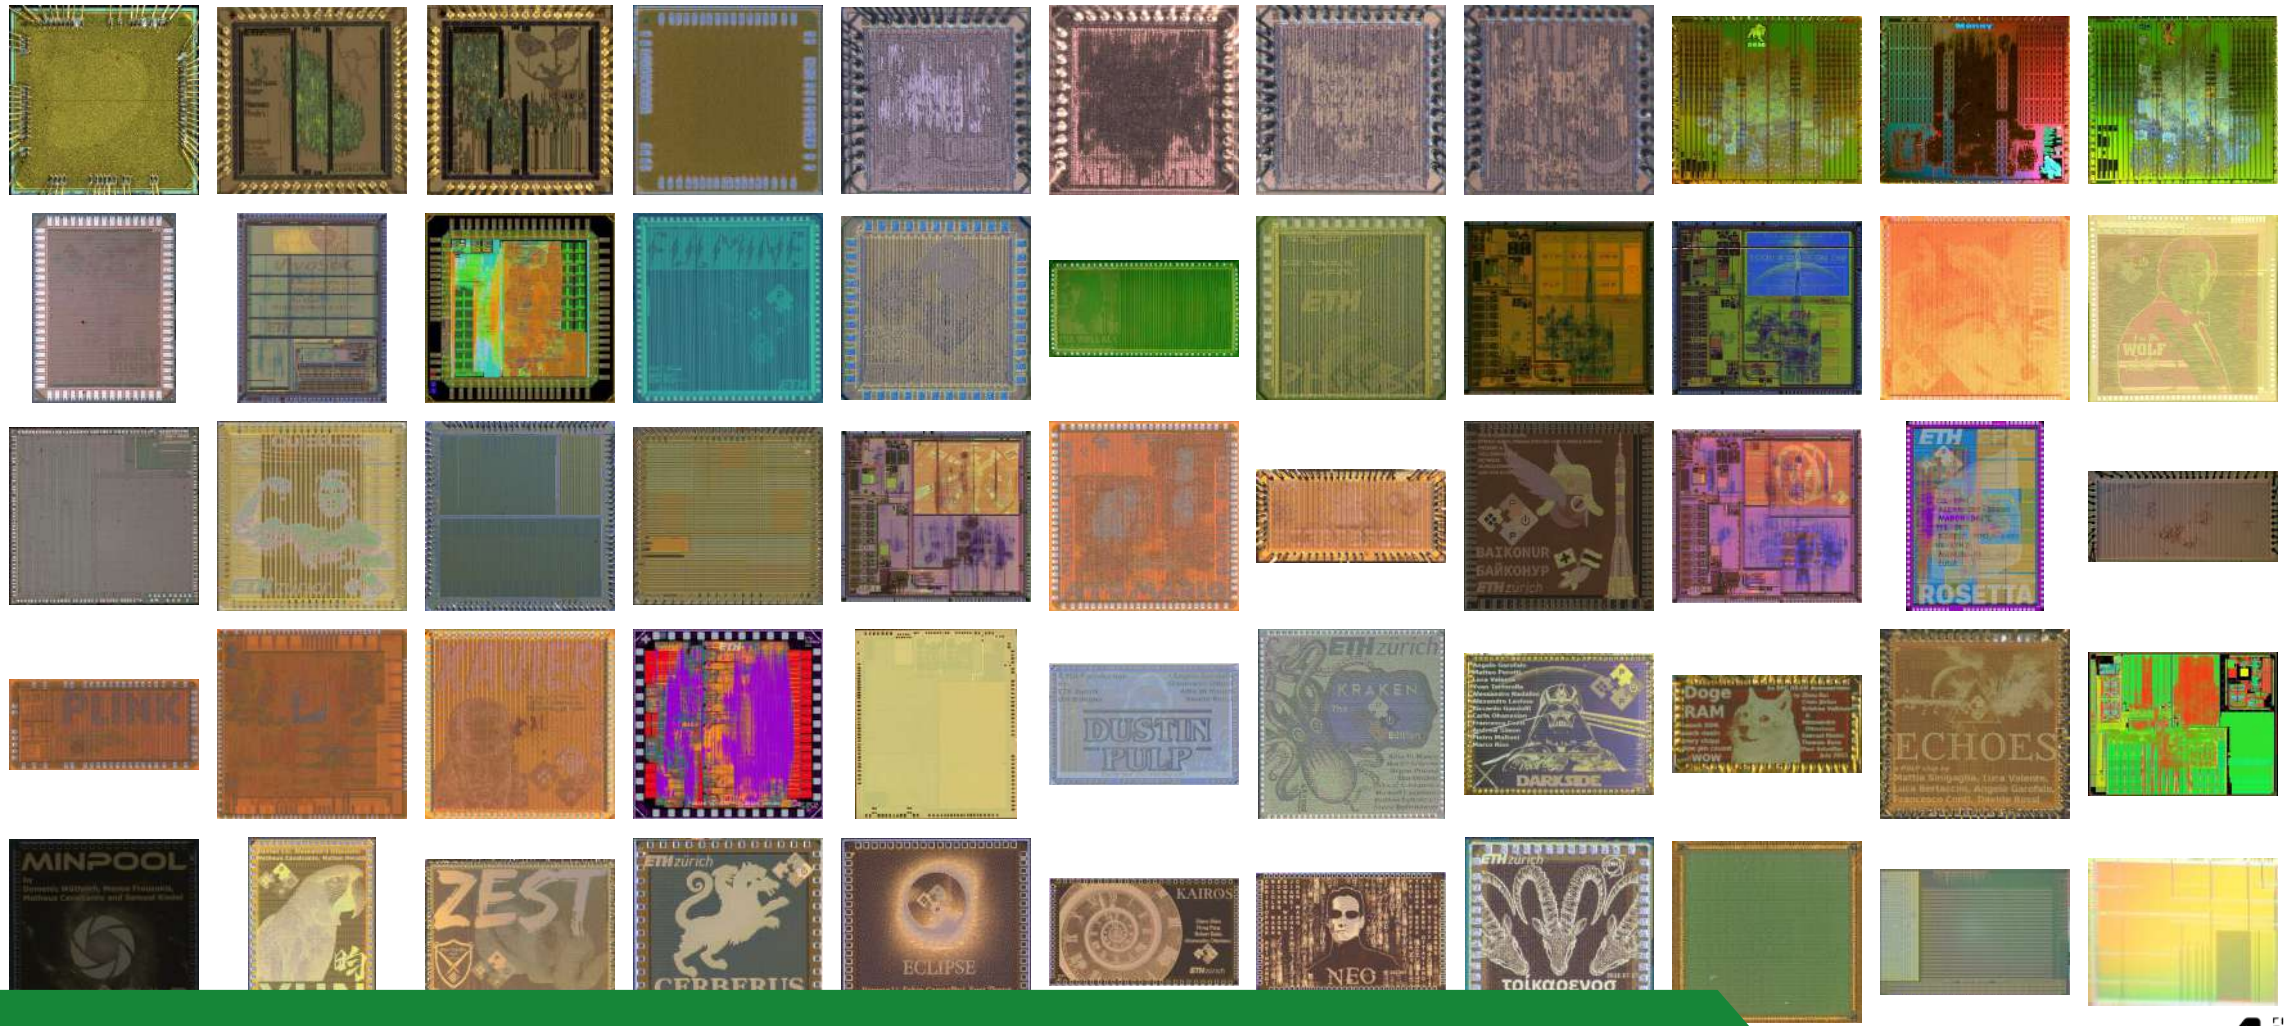

Watch this space, more publications from recent ASICs to come

These are not toy ASICs: (64b / 32b core count)

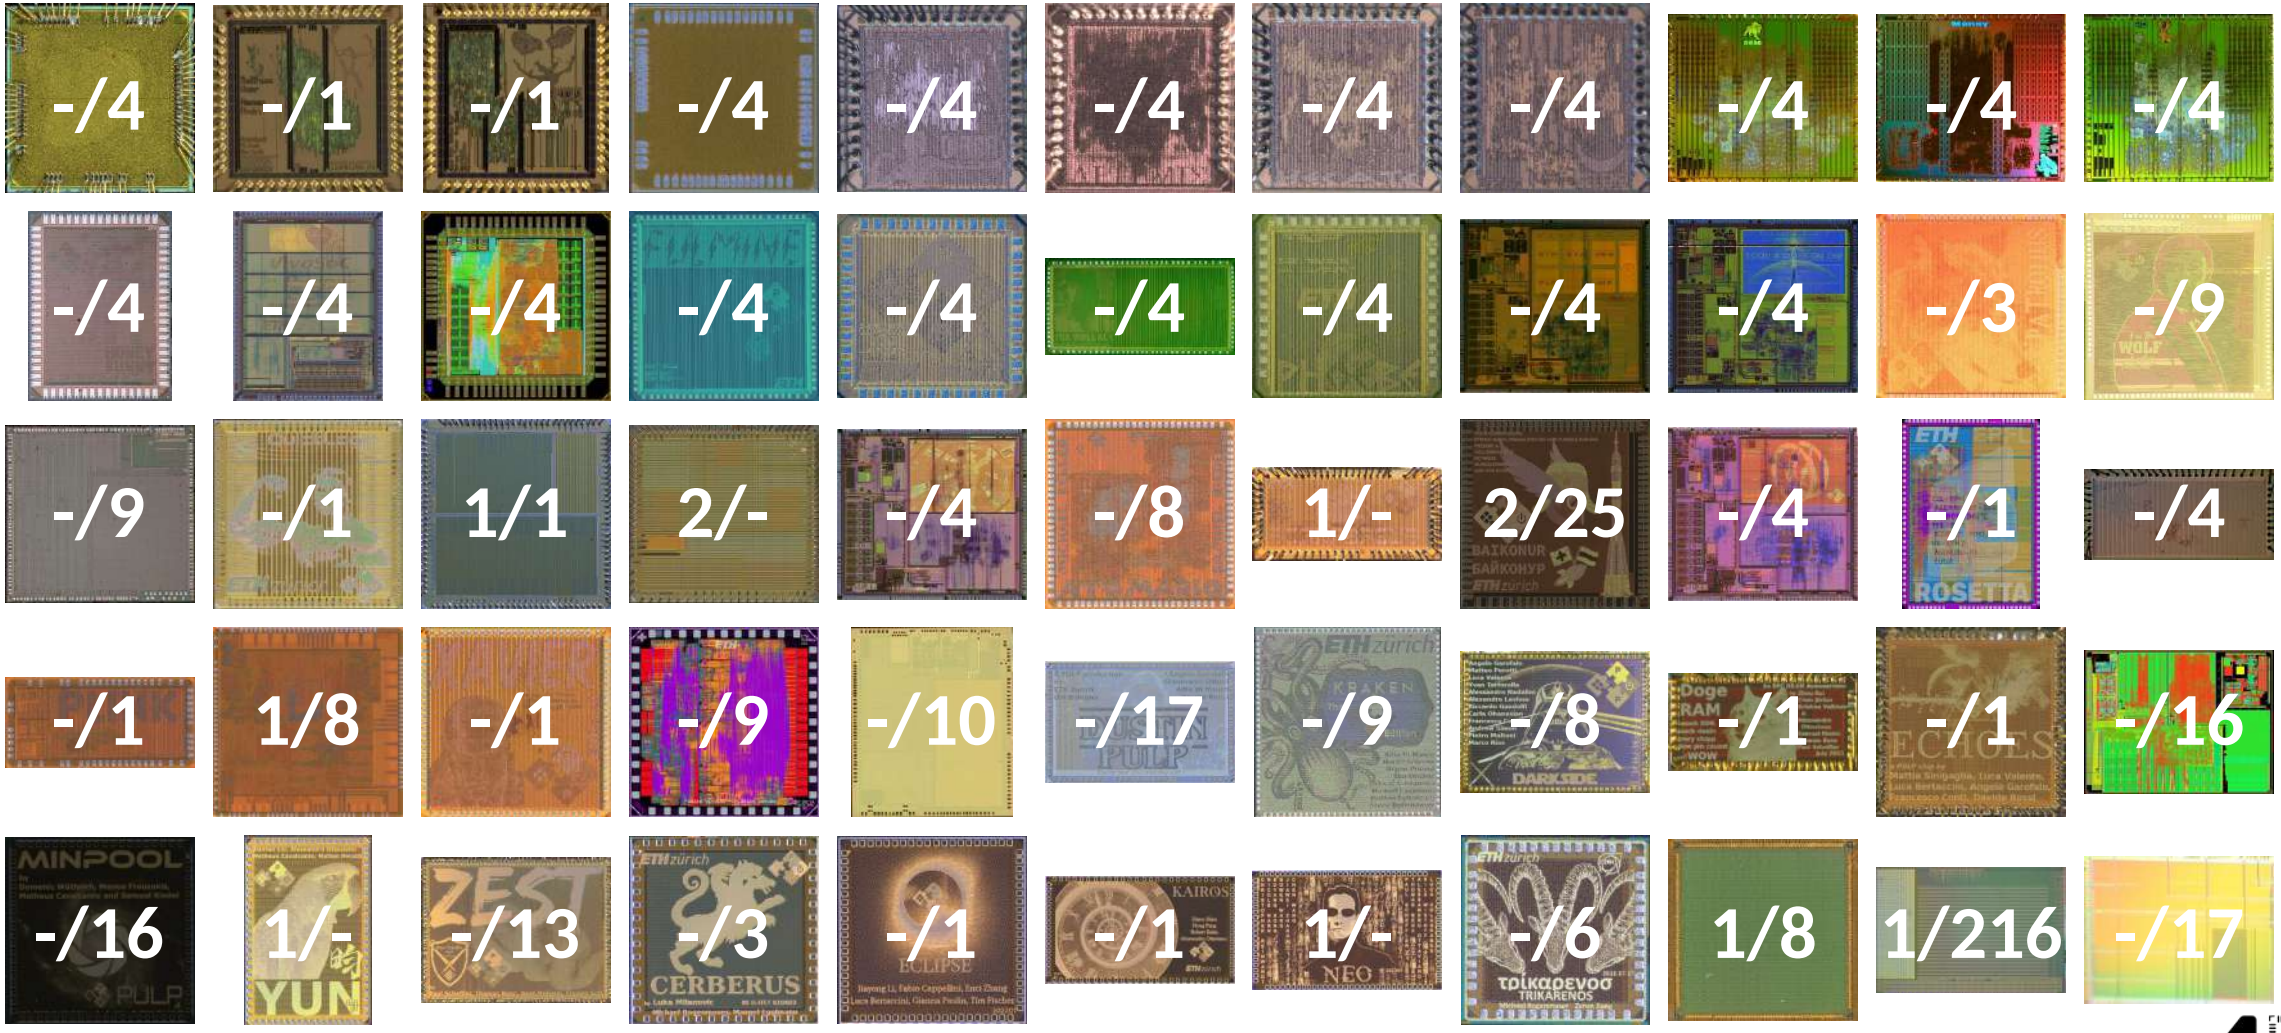

Life cycle of an IC Design project

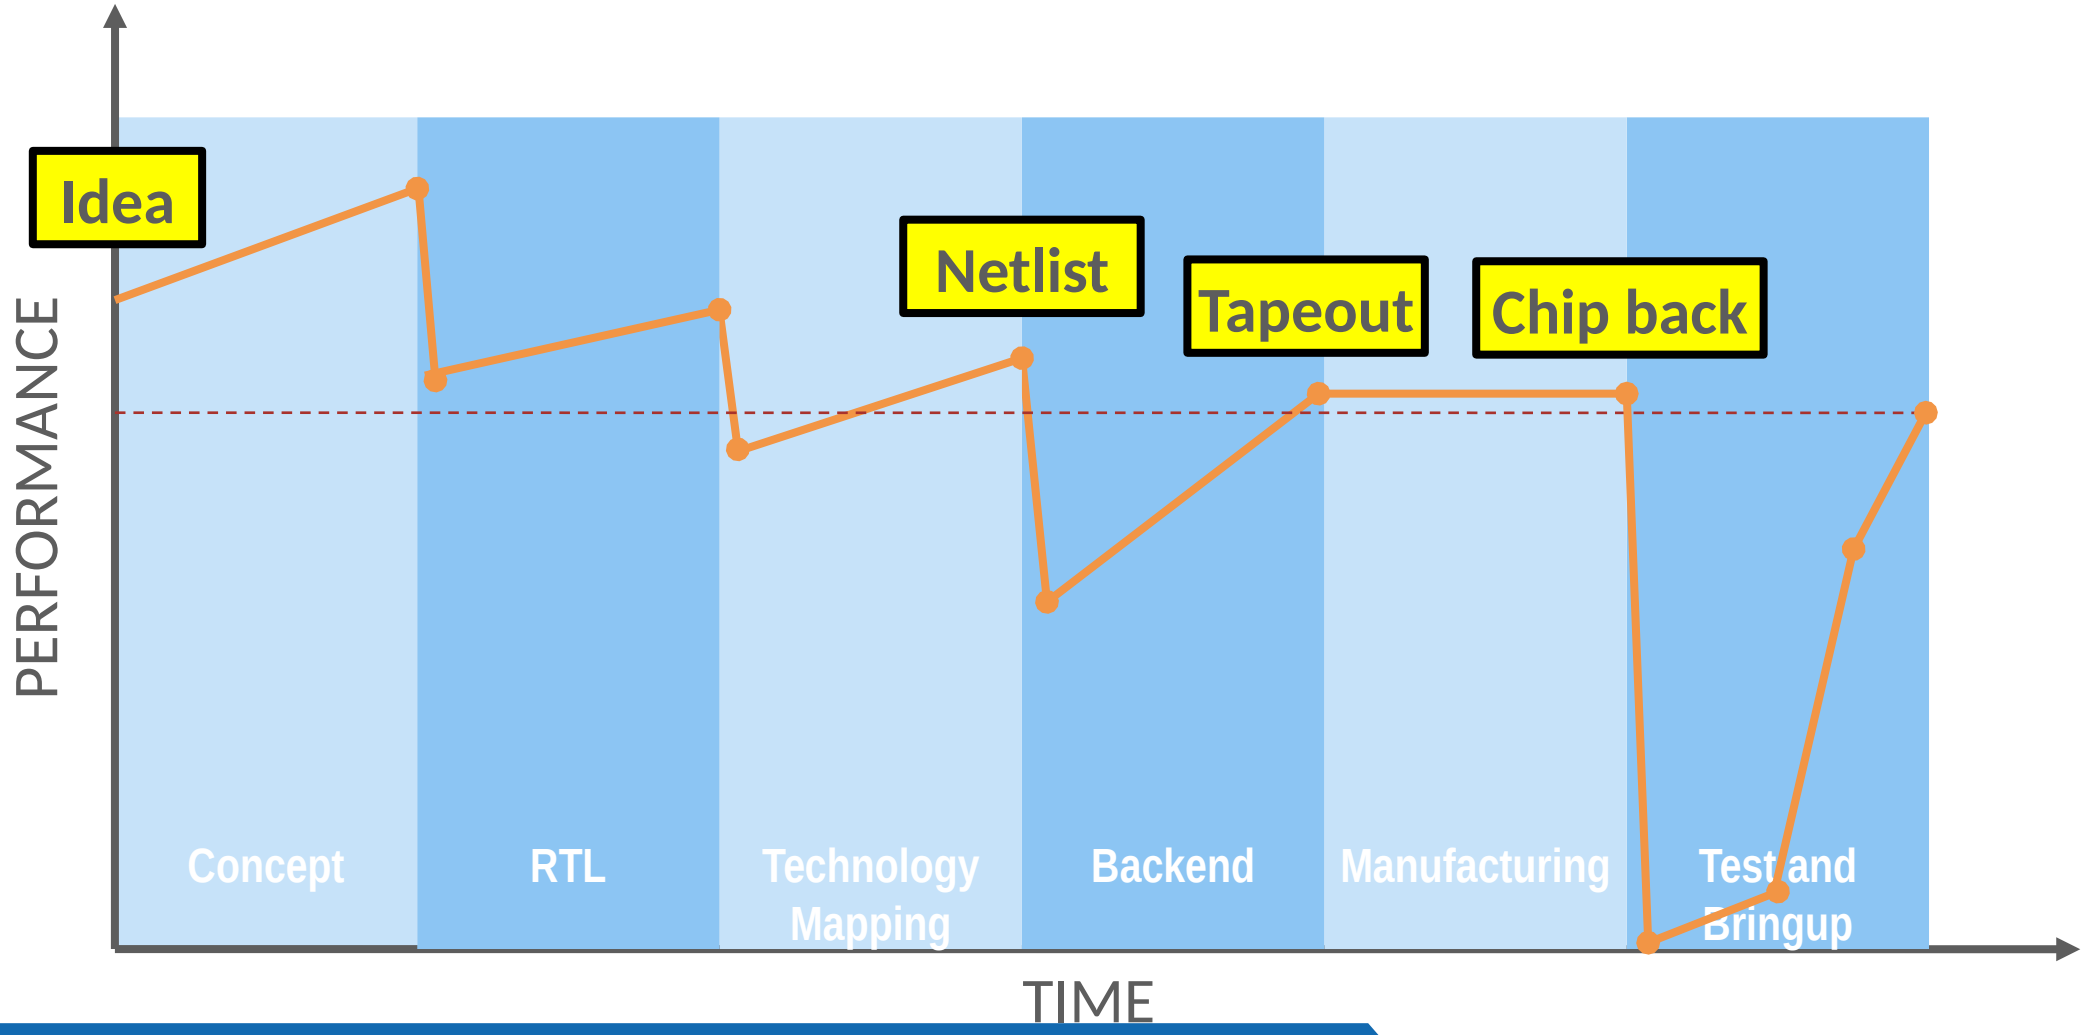

Designs always look better on paper

The tester room

When you figure it out

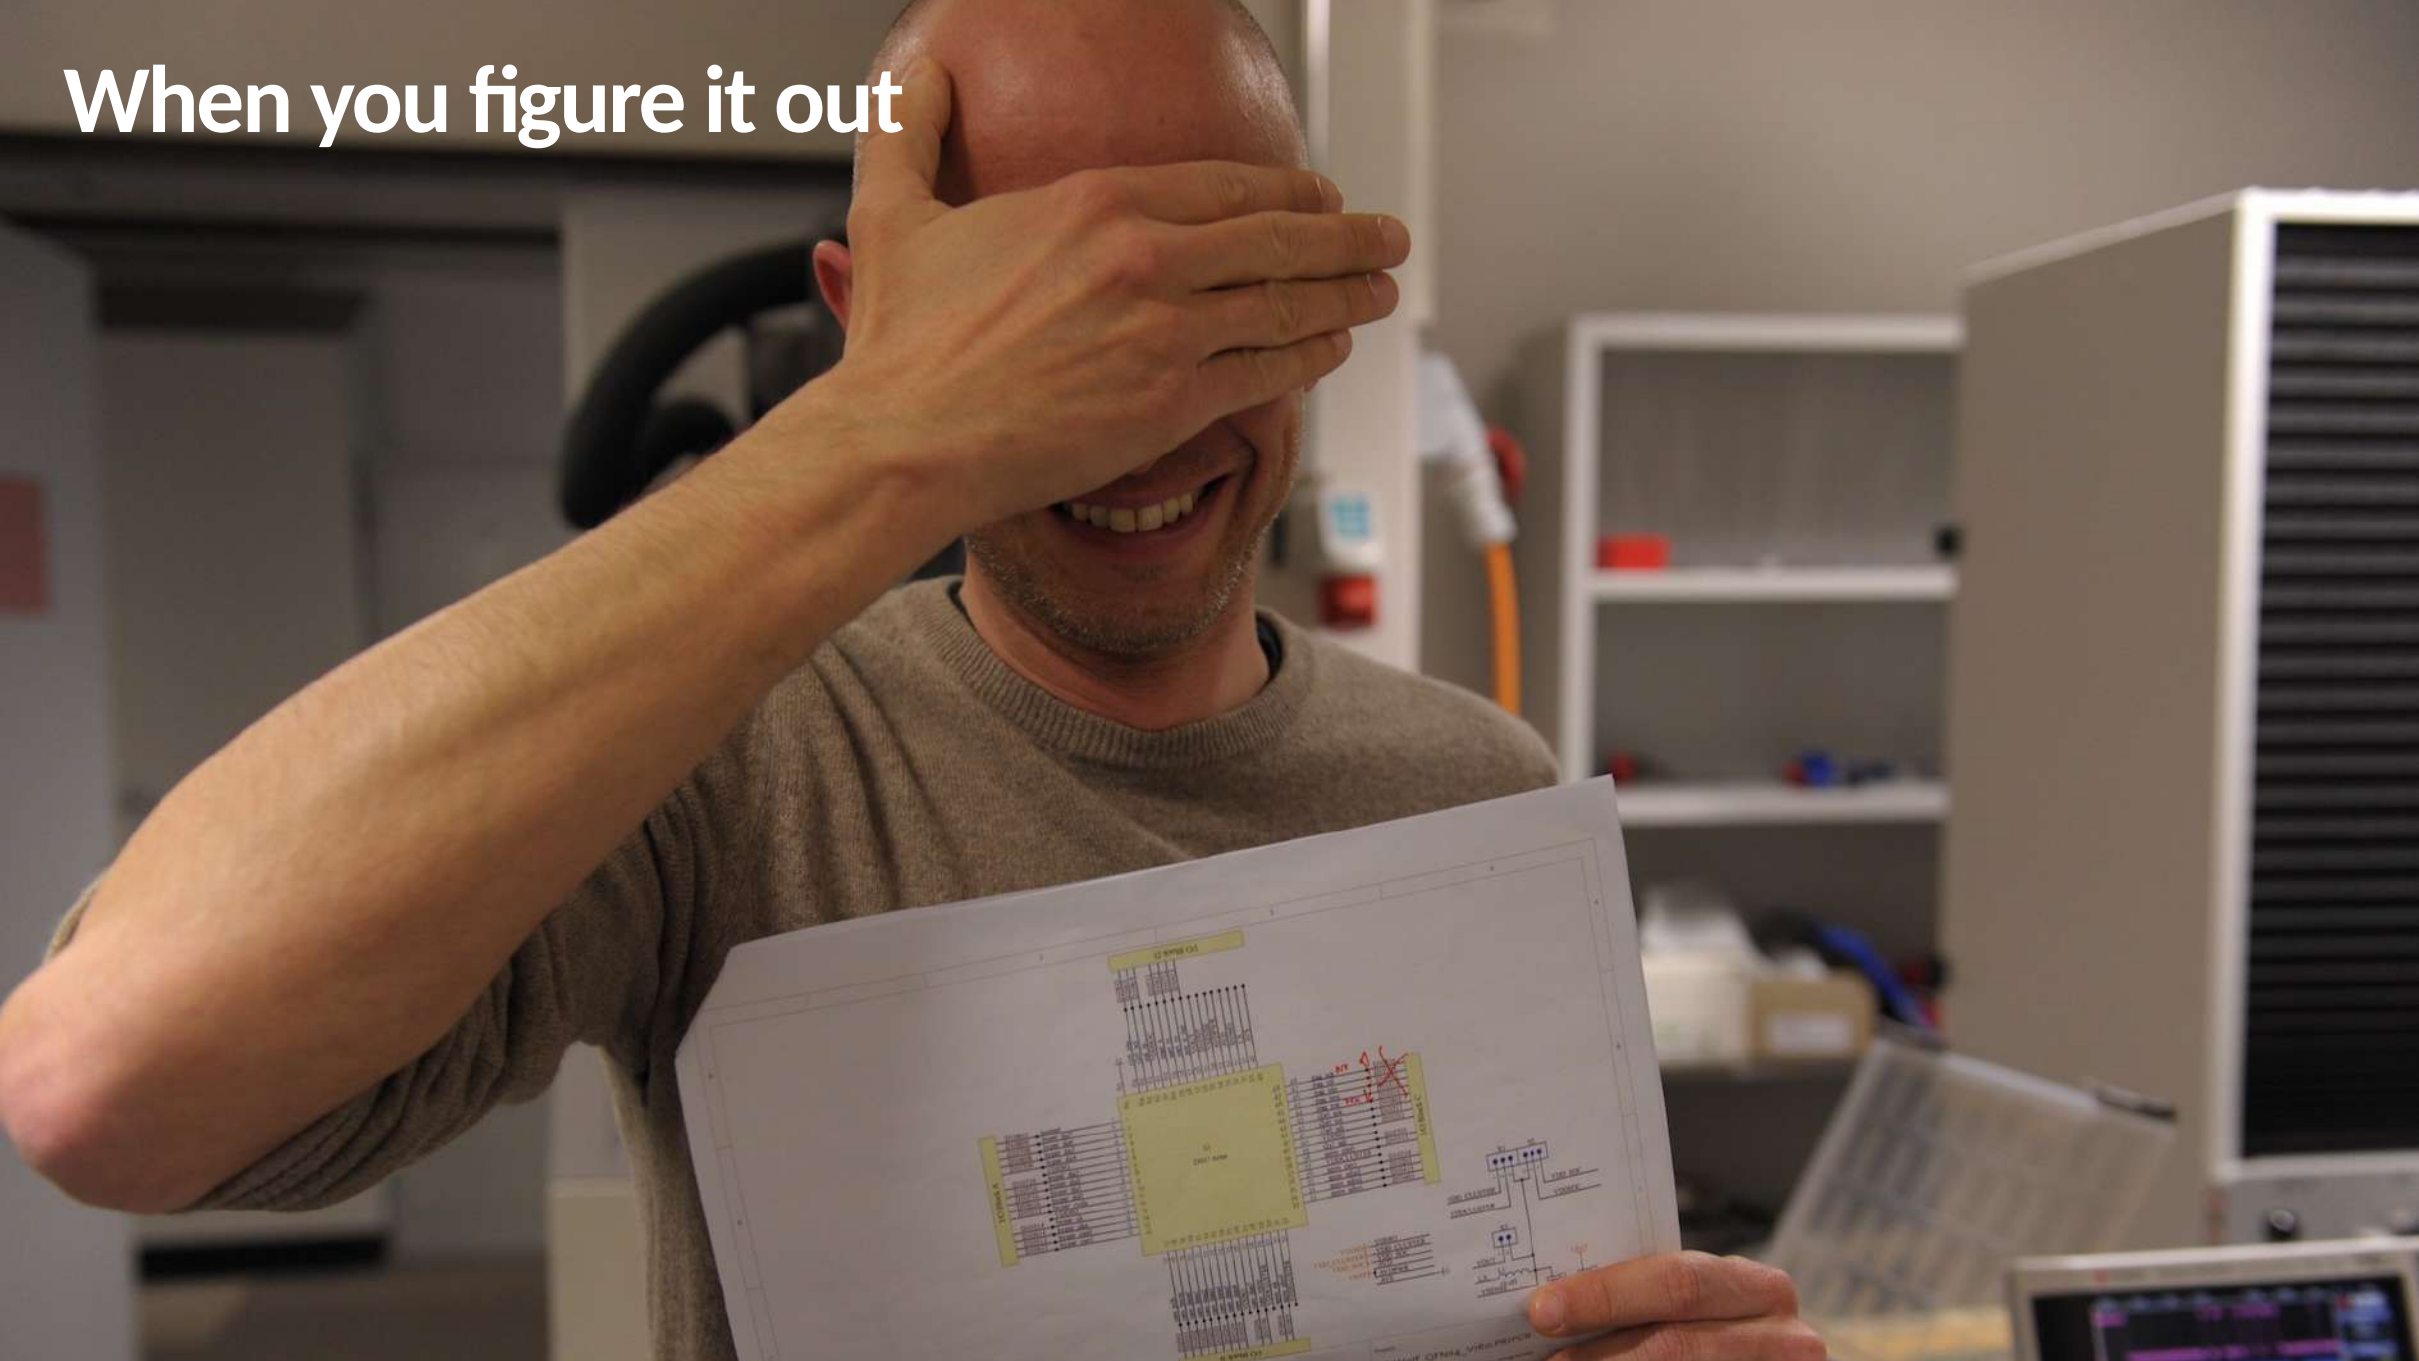

And when things go right

Today the PULP team has grown to more than 70 people

- **Nearly half of the team has had experience with at least 1 tape-out**
- **And some of the rest has just started**

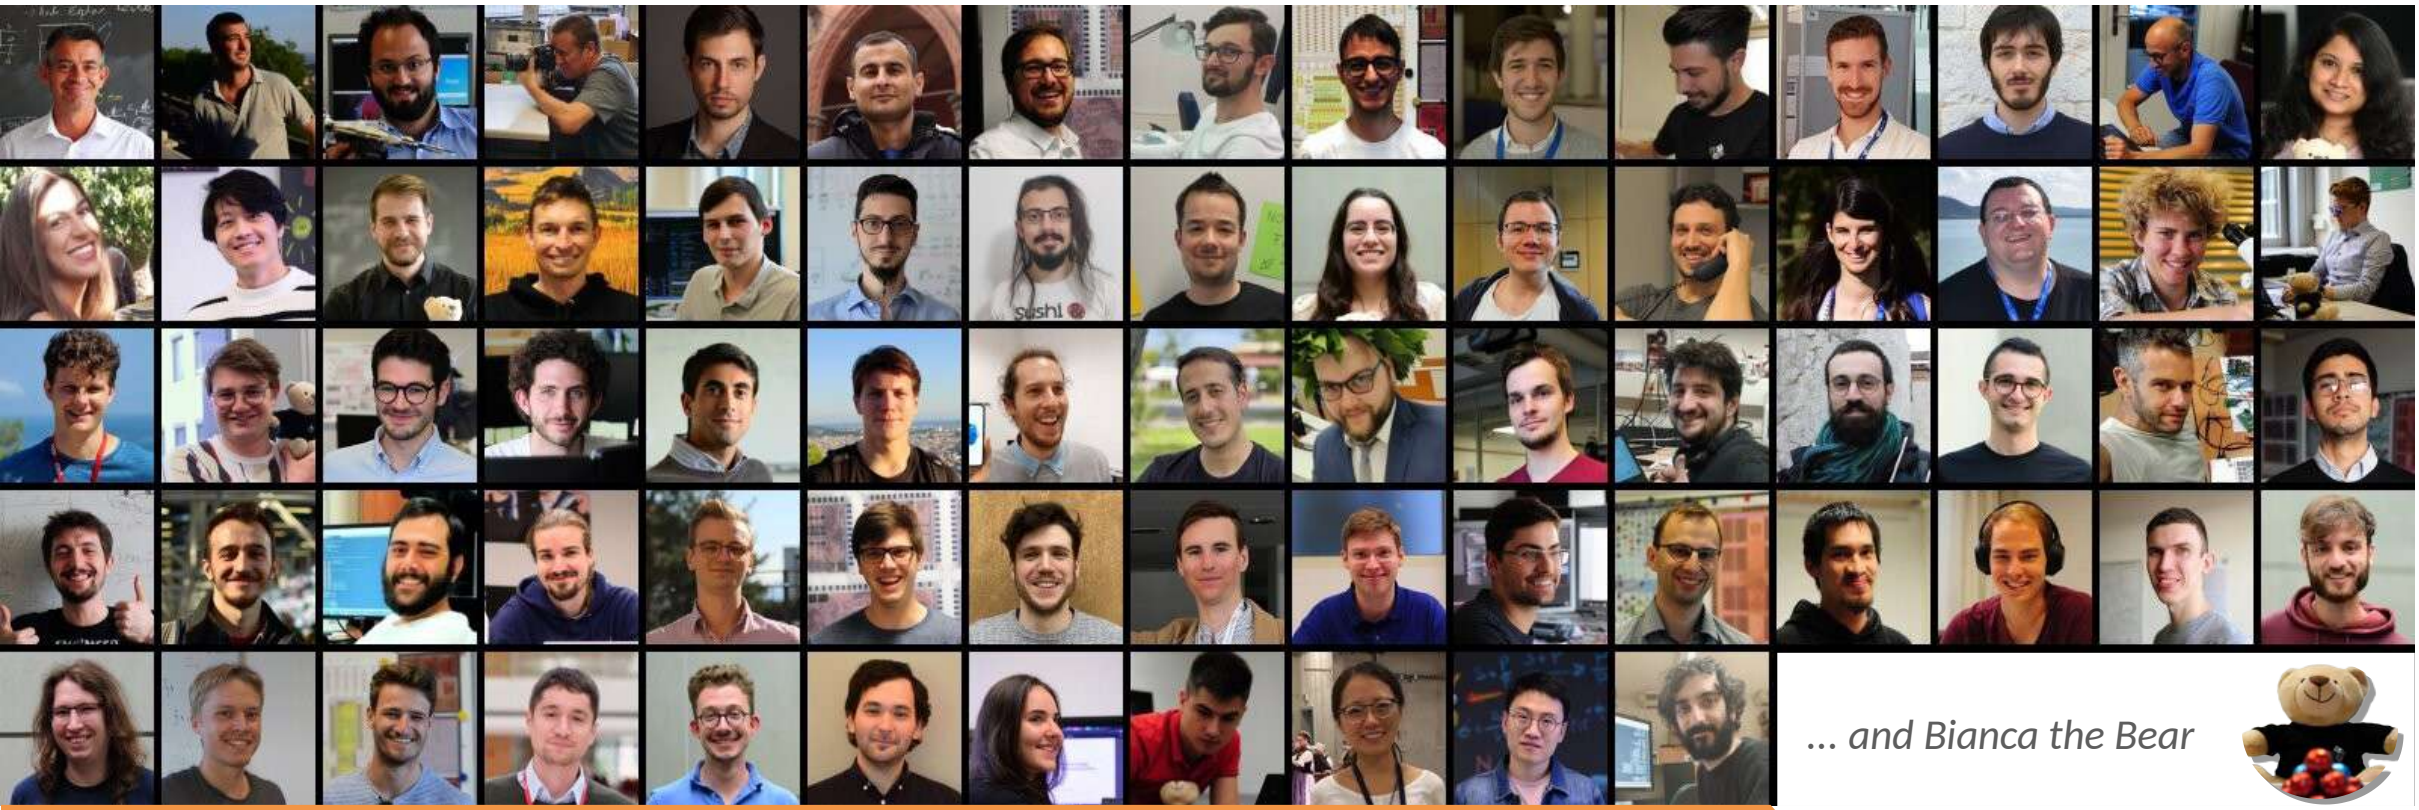

... and Bianca the Bear

If your picture is missing, you did not send it to me

We started almost excatly 10 years ago (April 2013)

- **Investigating new computing architectures**
  - Efficient over a wide range from IoT applications to HPC systems
- **Key points**
  - Parallel processing
  - Near threshold computing
  - Efficient switching between operating modes
  - Making best use of technology
  - Heterogeneous acceleration
- **Parallel Ultra Low-Power (PULP) platform was born**

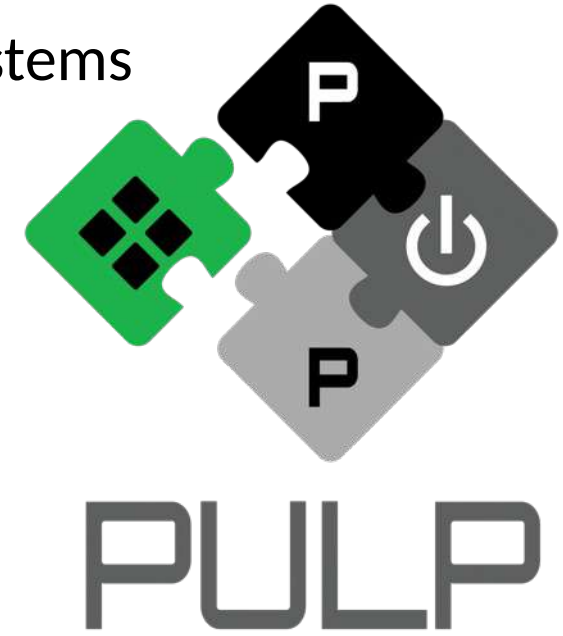

Right out of the gate – our first ASIC was PULPv1 in 28nm

- **Our first complete PULP chip -2013**
  - 4x OpenRISC cores
  - STM 28FDSOI technology (RBB)
  - Explores body-biasing
- **Collaboration with STM (France)**
  - They needed a complete system demo (more than ring oscillators)
  - Demo for technology capabilities
- **Meant for an IC tester**
  - Almost no I/Os

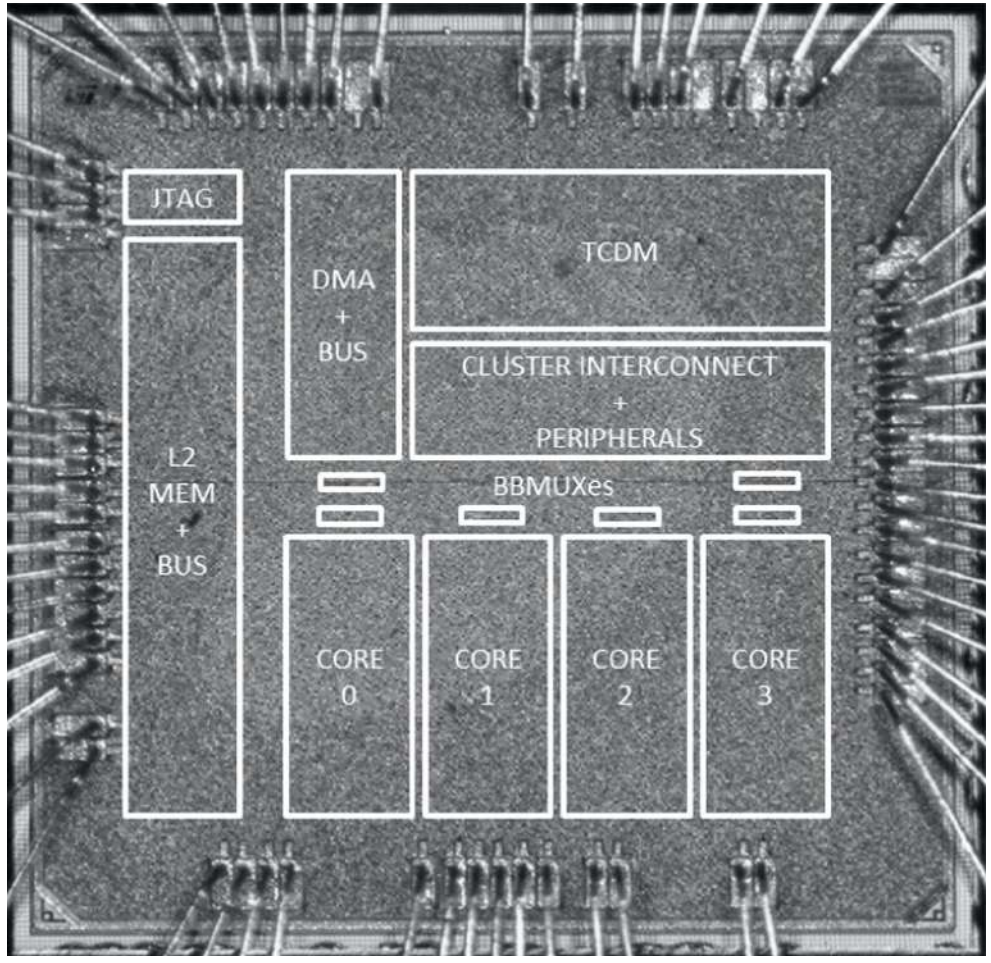

Davide Rossi, Antonio Pullini, Igor Loi, Michael Gautschi, Frank K. Gurkaynak, Andrea Bartolini, Philippe Flatresse, Luca Benini, "A 60 GOPS/W, -1.8 V to 0.9 V body bias ULP cluster in 28 nm UTBB FD-SOI technology", Journal of Solid-State Electronics, Volume 117, March 2016, Pages 170-184, DOI: 10.1016/j.sse.2015.11.015

This graph that we use, is from measurements of PULPv1

[Rossi et al. IEEE Micro 2017]

- **As VDD decreases, operating speed decreases as well.**
- **However efficiency increases more work done per Joule**
  - Until leakage effects start to dominate
- **Put more units in parallel to get performance up and keep them busy with a parallel workload**

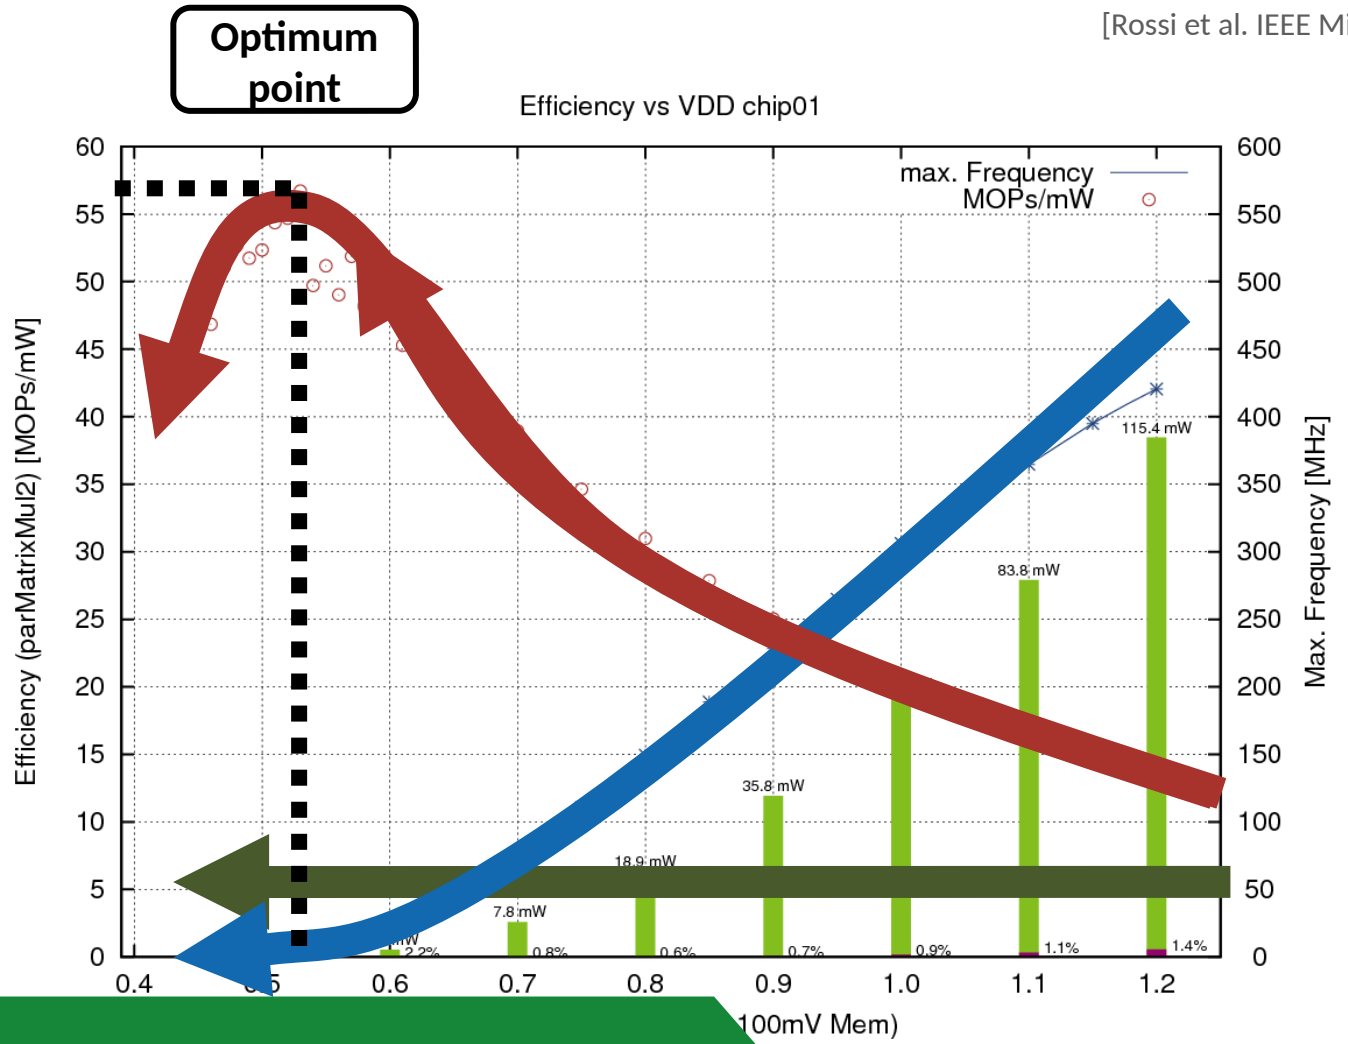

N cores running at lower VDD is more energy efficient

First steps to open source, how to start?

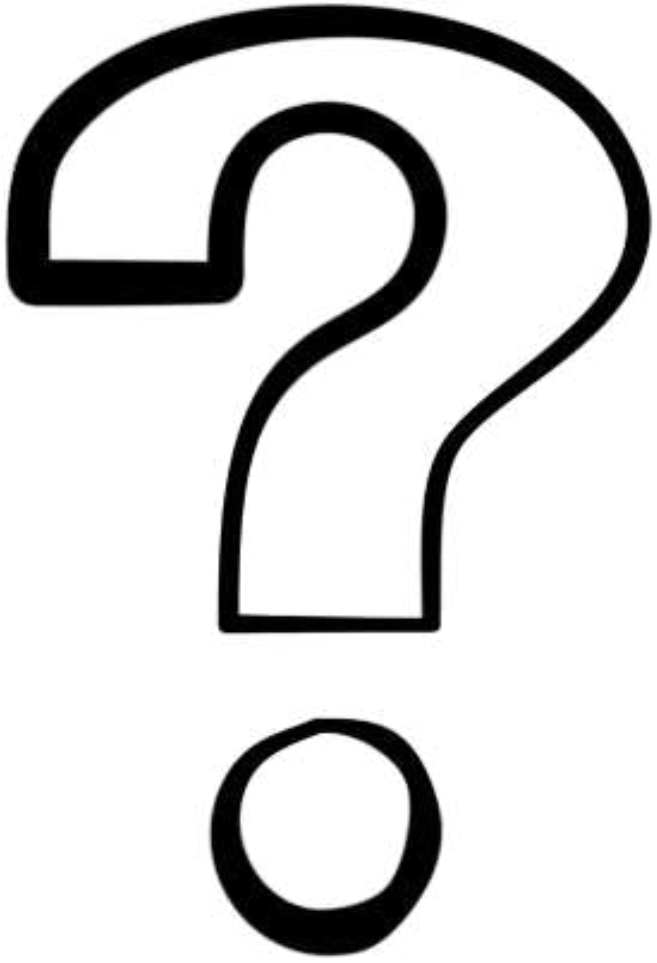

- **At this time nothing was released**
  - We were 100% sure it would become open source
  - But we had no idea how
    - What can we open source, and what not
    - We work for ETH Zurich, we have to ask their permission
  - We also did not have much idea about licensing
- **We need support of industry**
  - This project was supported by ST Microelectronics
    - They would not support a project where they can not use our work ‘freely’
  - Permissive licenses are the only way
    - Even though purists consider it not ‘free’ enough

Fulmine: the award winning one

- **UMC65**
- **Earlier chip (2015)**
  - 4x OpenRISC cores (not yet RISC-V)
  - 192 kBytes L2 + 64 kBytes TCDM
  - 2x HW accelerators
    - HW – Crypt (together with TU-Graz)
    - HW – Convolution Engine
- **Publication from this chip**

Francesco Conti, Robert Schilling, Davide Schiavone, Antonio Pullini, Davide Rossi, Frank K. Gurkaynak, Michael Muehlberghuber, Michael Gautschi, Igor Loi, Germain Haougou, Stefan Mangard, Luca Benini, "An IoT Endpoint System-on-Chip for Secure and Energy-Efficient Near-Sensor Analytics", IEEE Transactions on Circuits and Systems I: Regular Papers, Vol: 64, Issue: 9, Sept. 2017,pp 2481 - 2494, DOI: 10.1109/TCSI.2017.2698019

We have a base to work on and expand

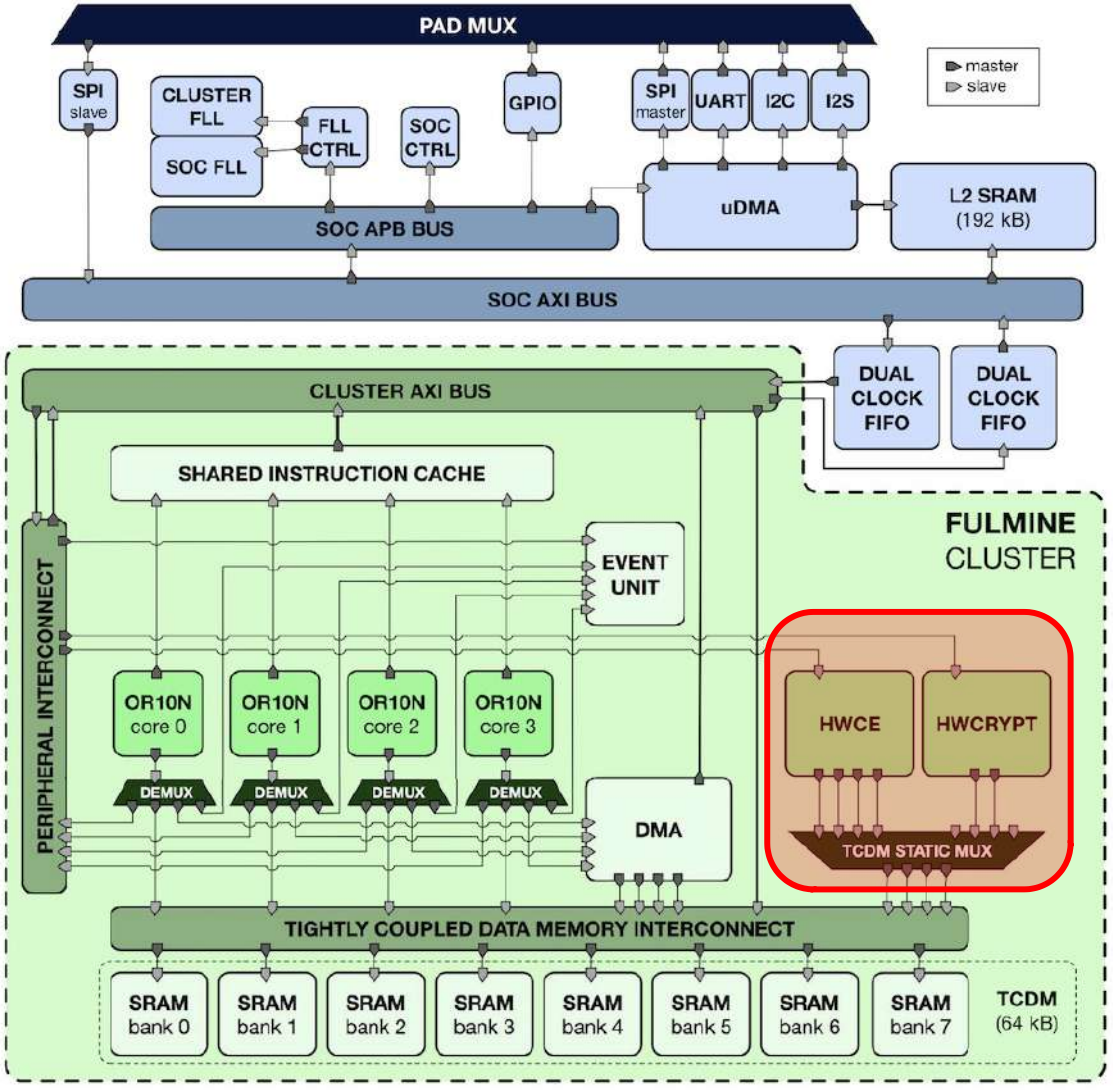

- **Much more than a core**
  - Peripherals (SPI, UART, I2C, I2S)
  - DMA, Busses, event unit
- **First chip with accelerators**
  - zero-copy connection to the memory
  - Allows independent systems (HWCrypt/HWCE) to be added easily.
- **Still not openly released**
  - Using our OpenRISC core (3rd gen)

Mr. Wolf (TSMC 40) 2017: 8+1 core IoT Processor

- **One cluster with**
  - 8 RISC-V cores
  - 2x shared FPU units
  - 64 kByte of TCDM
- **One controller with**
  - 512 kByte L2 RAM
  - Peripherals
- **On chip voltage regulators**
  - By Dolphin Integration

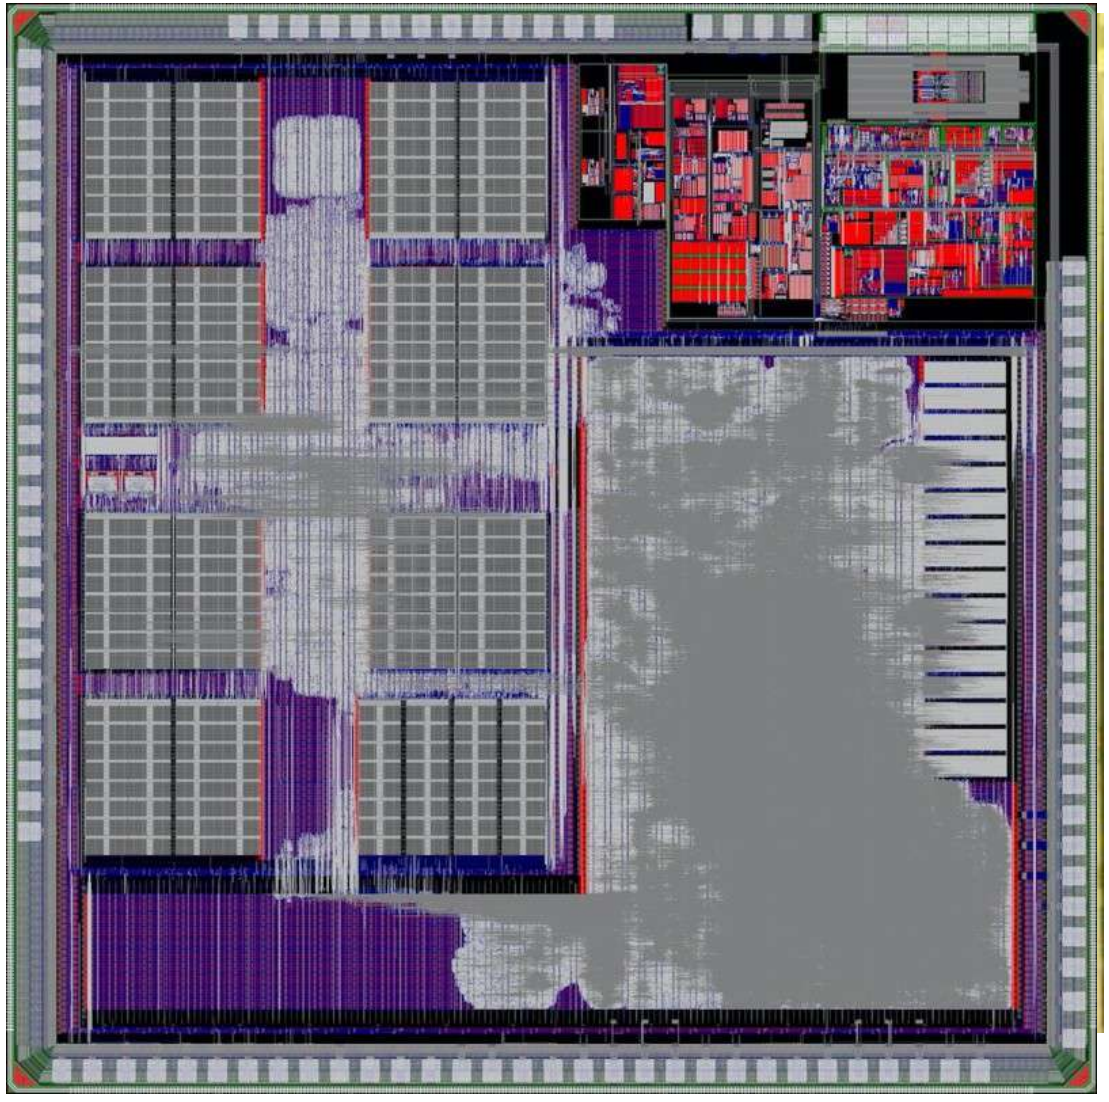

Antonio Pullini, Davide Rossi, Igor Loi, Alfio Di Mauro, Luca Benini, "Mr.Wolf: A 1 GFLOP/s Energy-Proportional Parallel Ultra Low Power SoC for IoT Edge Processing", In Proc. European Solid State Circuits Conference (ESSCIRC) 2018, 3-6 Sep 2018, Dresden, DOI: 10.1109/ESSCIRC.2018.8494247

What a difference two years make (Fulmine to Mr. Wolf)

- **With Mr. Wolf, most of what we have is open sourced**
  - This is a **complex IoT processor**, not like the much simpler PULPino
  - 8 + 1 cores, FPUs, shared accelerators, multiple power down modes.
- **The cores are now RISC-V**
  - Supports RV32IMCF and custom extensions (xPULP)
- **Interesting collaboration with Dolphin Integration (SOITEC)**
  - They have their IP demonstrated on an complex design, they can freely share
  - We get to use industrial IP in our chip
- **Still many parts can still not be open source**
  - FLL, analog macros, I/O cells, memory cuts (affects performance), P&R scripts

Mr. Wolf has been used in multiple systems

- **Designed as an application processor**
  - We still build boards with it
  - Despite only 200 manufactured
- **Widespread industrial use:**
  - Dolphin IP was validated on this chip
  - Greenwaves GAP8 is based on the open source release OpenPULP
  - BitCraze AI Deck is related
- **Phenomenally successful**
  - University chips are not meant to work this well

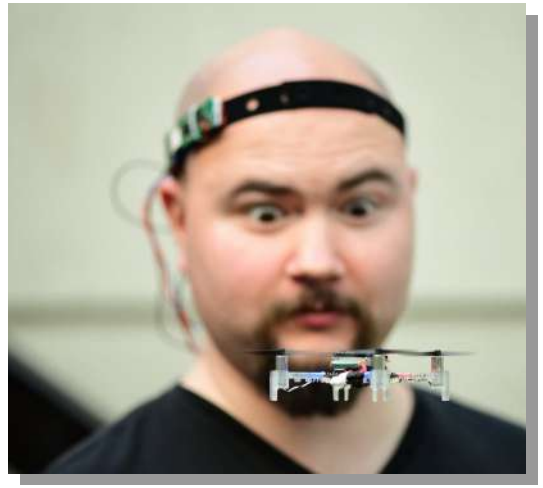
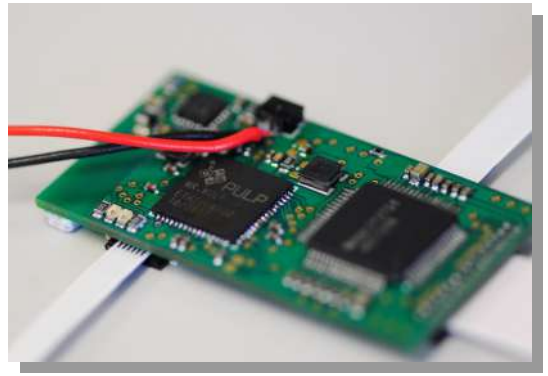


GREENWAVES TECHNOLOGIES

bitcraze

IcySoC chips (Manny, Sid, Diego), made in Switzerland

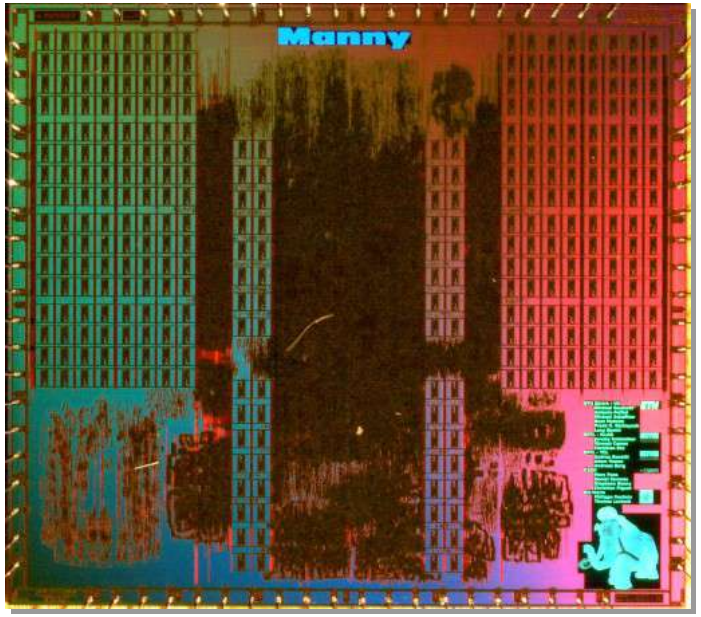

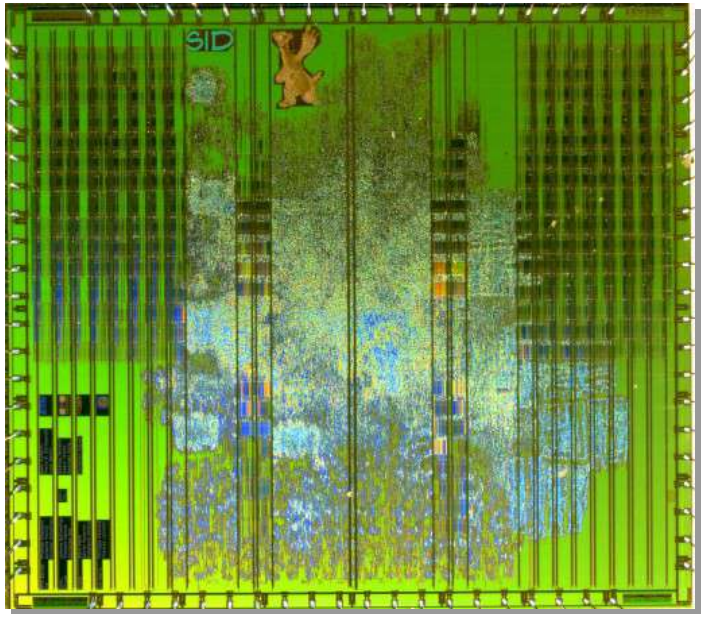

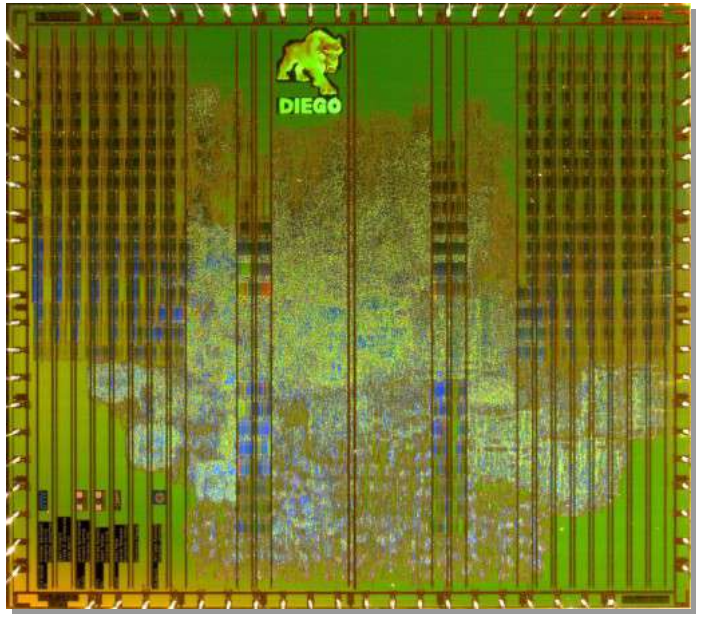

- **ALP 180 nm technology from EM-Marin, near Neuchatel**
  - Part of the IcySoC Nano-Tera project
  - Partners EM-Marin, CSEM, EPFL
  - Each chip uses a different library combination (regular, lowVT, superlowVT
  - Each chip is 7.20mm x 8.16mm (58 mm^2) - for comparison Occamy is only 25% larger (73 mm^2)

The ICYSoC Wafer

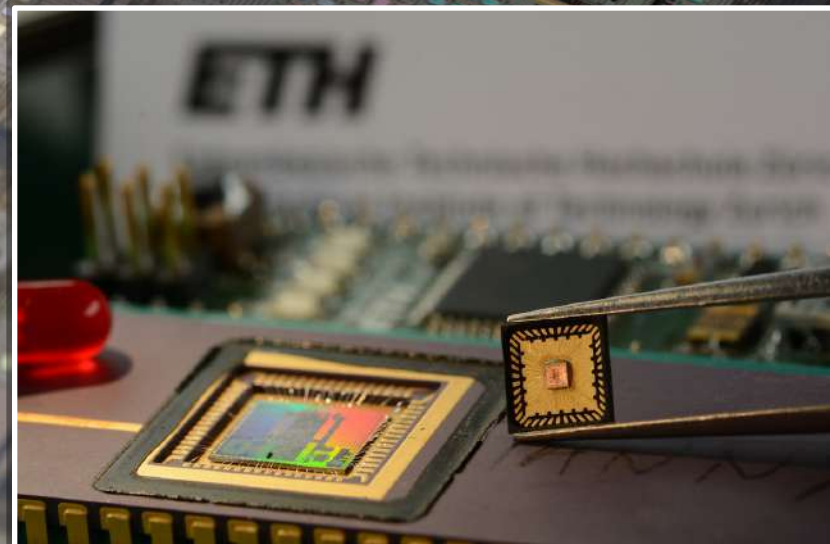

Not all of our ASICs are pure PULP designs – Origami 2014

- **Quite a few standalone chips designed**
  - Accelerators, either meant as a co-processor
  - Or contain just the computation engine
- **Origami is one of our most cited ASICs**
  - UMC 65nm
  - CNN accelerator running at 700 MHz
  - Designed as part of a semester thesis
- **Fun fact:**
  - One of the hardest logos to get it right
  - It does not look that way

L. Cavigelli and L. Benini, "Origami: A 803-GOp/s/W Convolutional Network Accelerator," in IEEE Transactions on Circuits and Systems for Video Technology, vol. 27, no. 11, pp. 2461-2475, Nov. 2017, doi: 10.1109/TCSVT.2016.2592330.

The story of Ariane

- **All started in 2017 with Florian Zaruba**
  - He was looking for a topic for his M.Sc. Thesis
- **At the time..**
  - We already had designed several 32bit RISC-V cores
  - And Luca had famously said:
    "We will never build a 64bit core"
- **Only 6 months later**
  - His thesis was complete
  - PULP had a brand new 64-bit Linux capable core
  - And a new chapter started

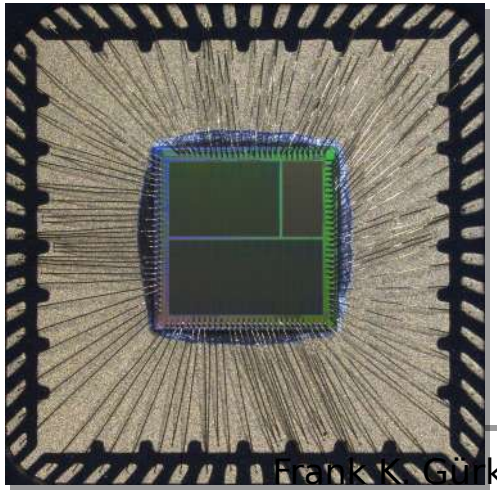

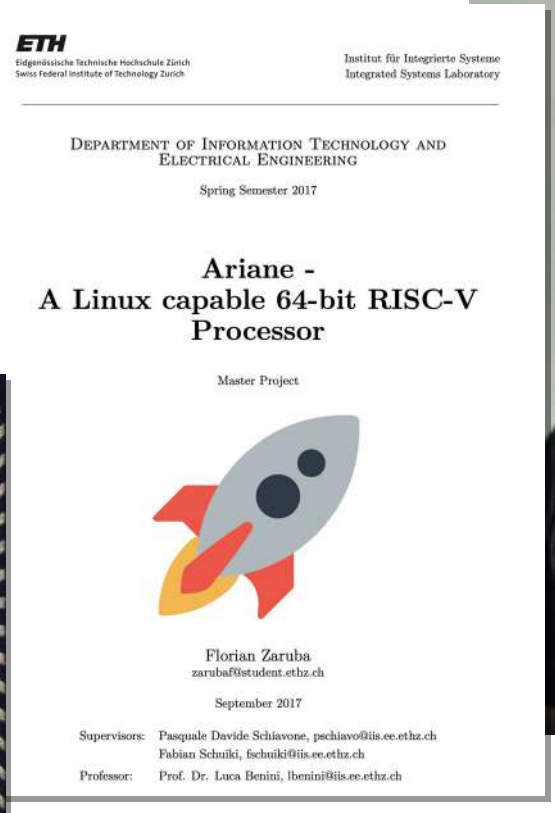

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Institut für Integrierte Systeme
Integrated Systems Laboratory

Department of Information Technology and Electrical Engineering

Spring Semester 2017

**Ariane -
A Linux capable 64-bit RISC-V Processor**

Master Project

Florian Zaruba
zarubaf@student.ethz.ch

September 2017

Supervisors: Pasquale Davide Schiavone, pschiavo@iis.ee.ethz.ch
Fabian Schuiki, fschuiki@iis.ee.ethz.ch
Professor: Prof. Dr. Luca Benini, lbenini@iis.ee.ethz.ch

Poseidon was the first Ariane chip

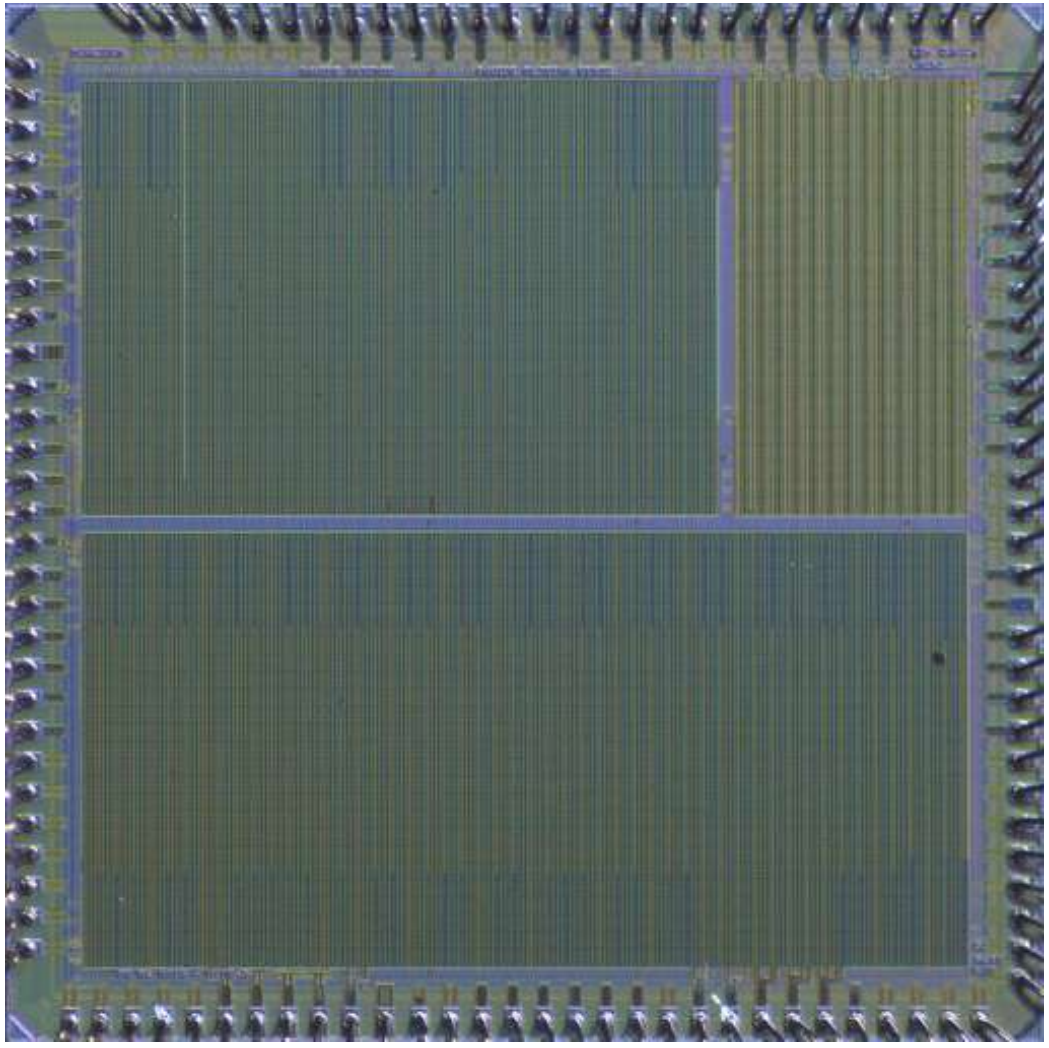

- **First GF22nm chip**
  - Used Europractice IC service
  - Cost 150k CHF for 50 samples
- **Has three parts (trident..)**
  - PULPissimo system
  - Ariane core
  - Independent ML accelerator
- **30 of 50 chips were packaged**
  - We provide a bonding diagram
  - Mostly simple manual work

And we nearly had some serious issues

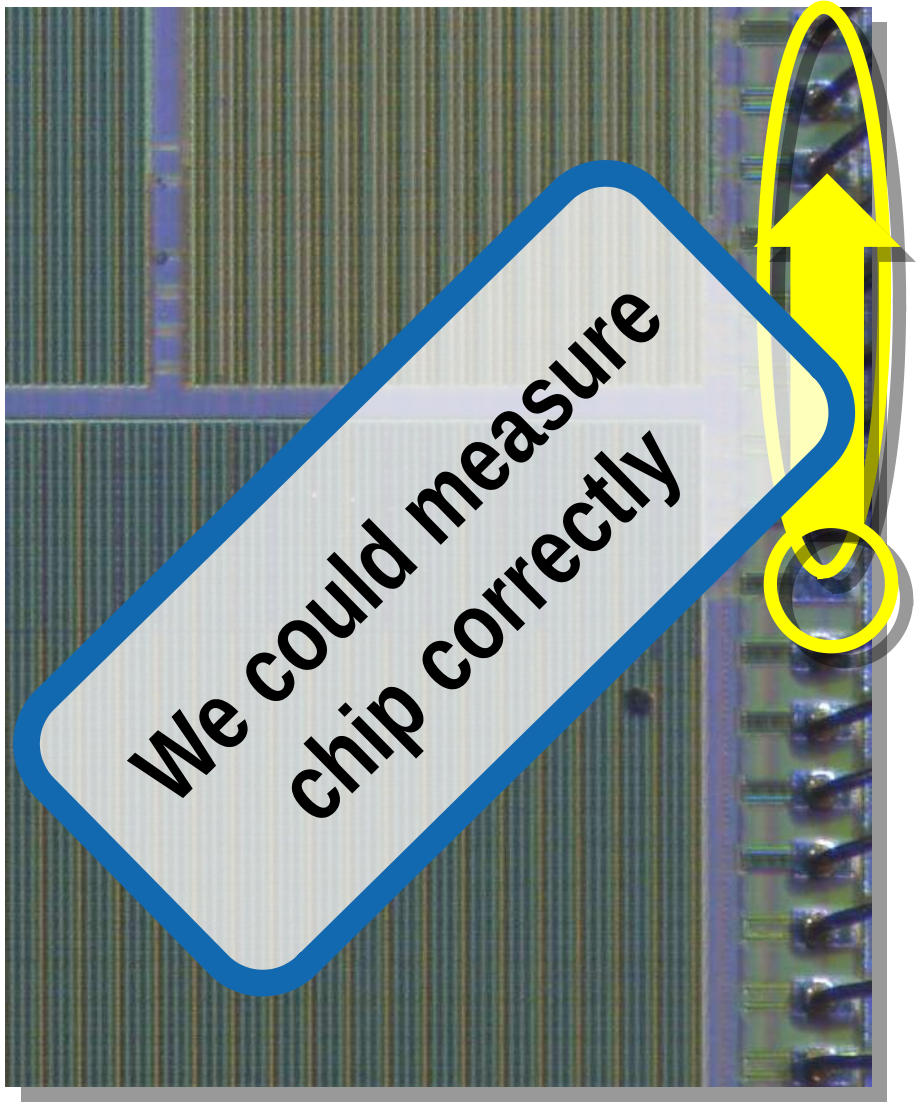

- **Look closer on the right side**
  - There is a pad that is not bonded
- **We skipped one pad**
  - All connections are shifted by one
- **VDD and GND are one after other**
  - Bonding causes shorts between VDD and GND
  - Pretty much catastrophic
- **Fortunately: unpackaged dies**
  - There were 20 unpackaged dies
  - We could bond those correctly
- **Very trivial errors, can be very costly**
  - We learn from every mistake

VivoSoC chips – Analog and Digital

- **Actually 4+ VivoSoCs since 2015**
- **SMIC 130/110 technology**
  - Many Analog IPs
    - ExG interfaces, A/D converters
    - Pulse Oximetry
    - Neuro stimulators
- **PULP cluster for post processing**
  - 4x RISC-V cores
  - Digital interfaces
  - DMA transfer from analog block to digital

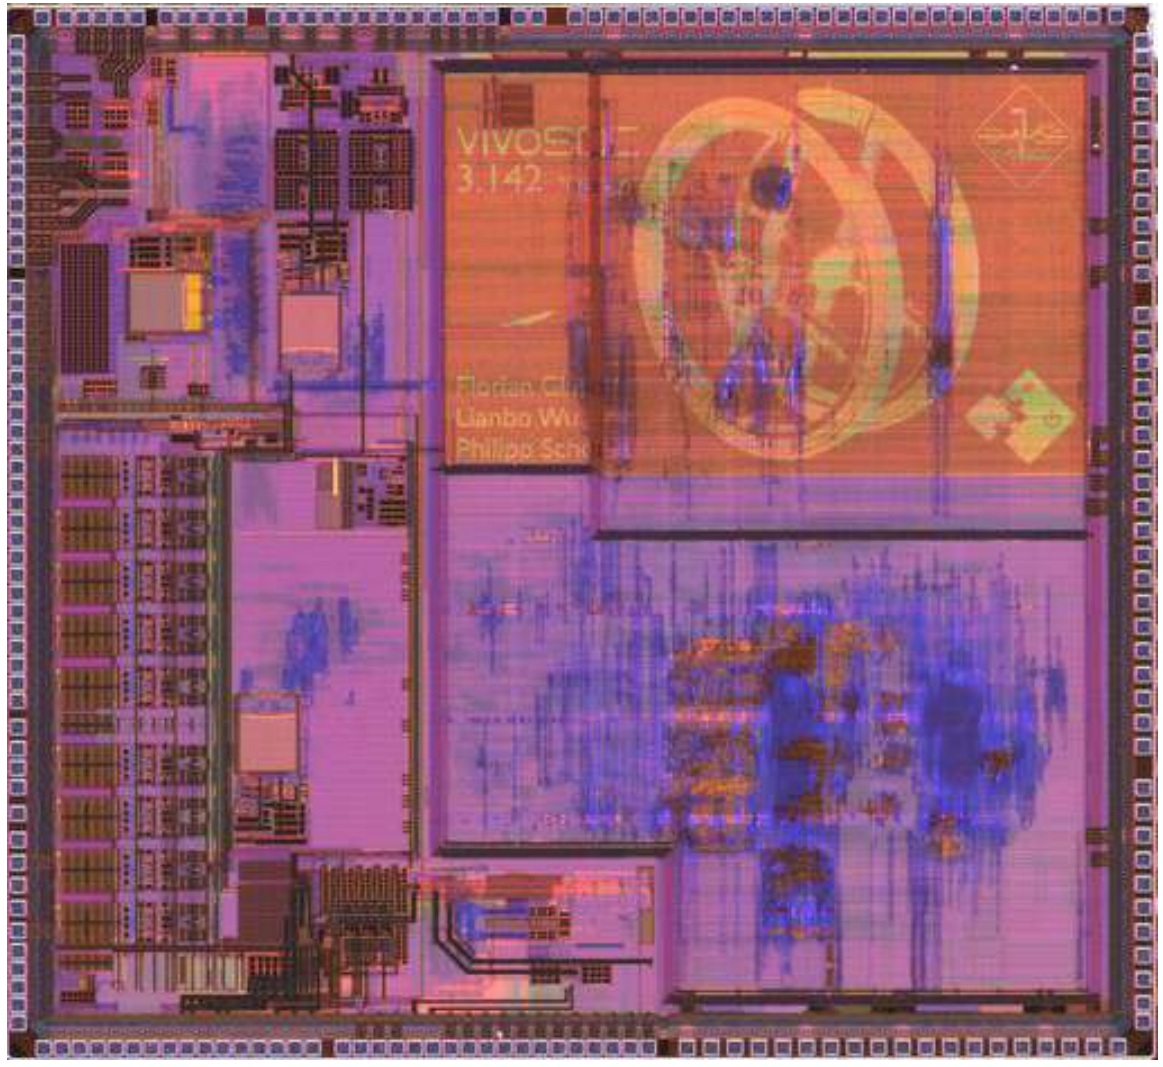

Philipp Schoenle, Florian Glaser, Thomas Burger, Giovanni Rovere, Luca Benini, Qiuting Huang, "A Multi-Sensor and Parallel Processing SoC for Miniaturized Medical Instrumentation", IEEE Journal of Solid-State Circuits PP issue:99, pp 1-12, DOI: 10.1109/JSSC.2018.2815653

PULP allows us to co-operate with everyone

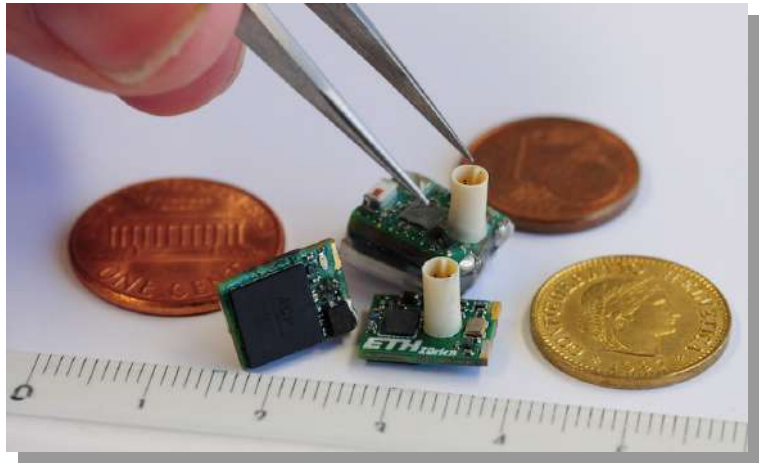

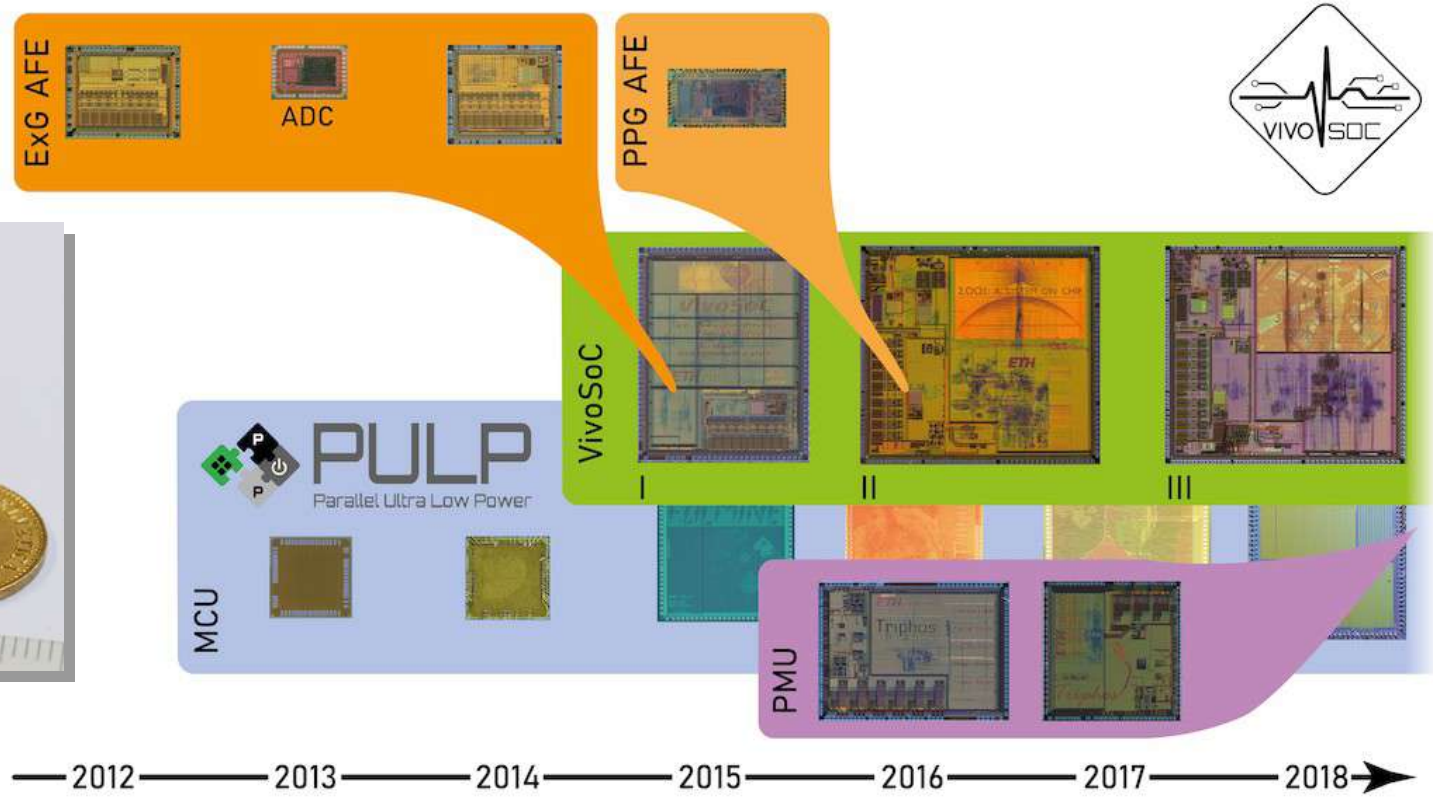

- **Collaboration between Prof. Benini and Prof. Huang**
  - Permissive licensing allows collaboration even if the result is **not** open source

Not everything goes right: reset problem of Urania

- **2 PULP clusters, each with**
  - 4x RV32 RI5CY cores
  - 4x transprecision FPUs
  - 1x PULPO accelerator
  - 64 kB TCDM in 8 banks
- **Ariane RV64 host processor**
  - 128 KiB Shared LLC
  - software-managed IOMMU
- **DDR3 DRAM Controller + PHY by TU-Kaiserslautern**

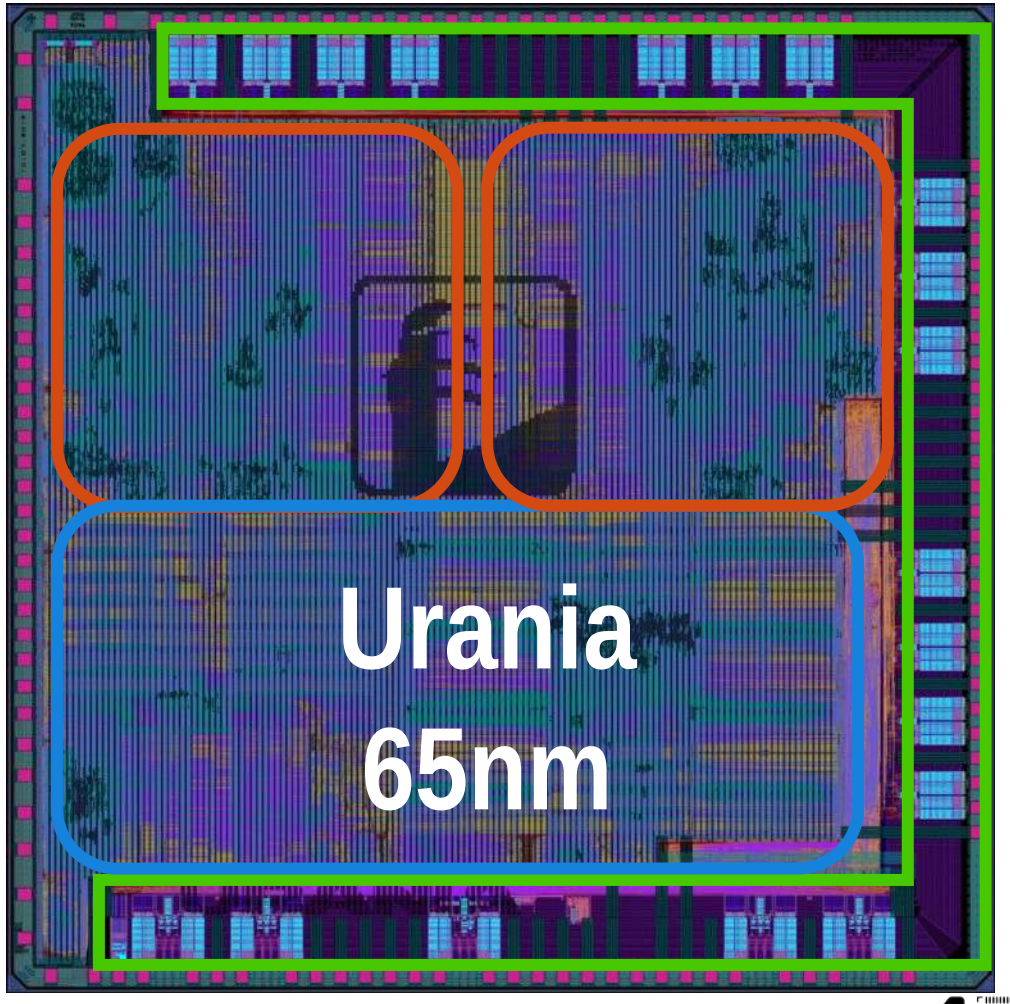

TECHNISCHE UNIVERSITÄT KAISERSLAUTERN

The reset can not be released for clusters

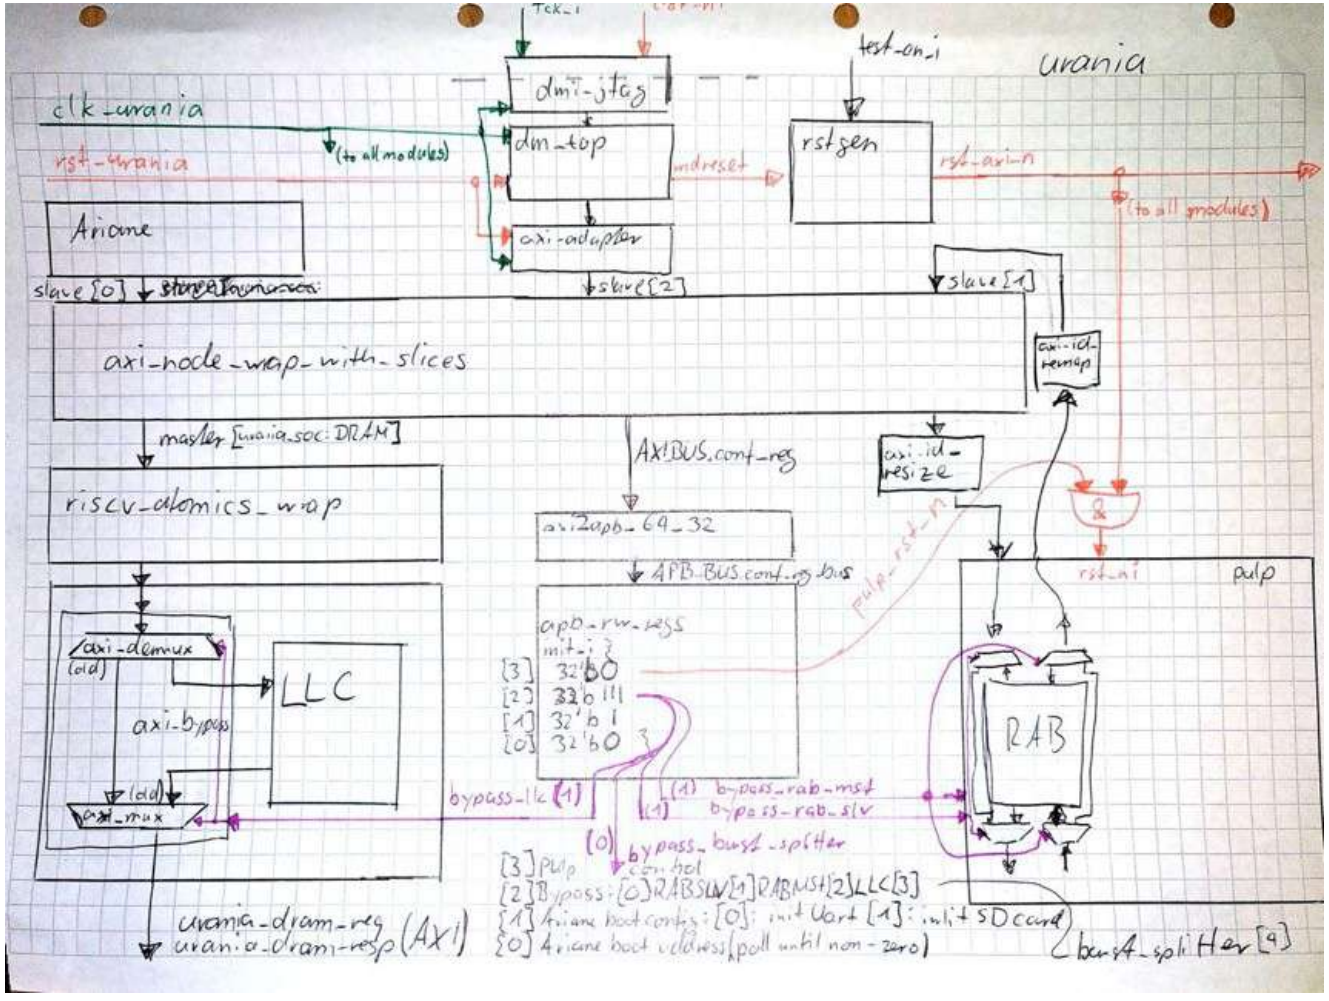

- **Chip has many modules**
  - 1x Ariane core
  - 1x DDR interface
  - 2x Clusters
- **Reset to clusters is stuck 0**
  - Design flow mistake
  - Some other control signals are stuck as well affecting Ariane performance
- **DDR interface is functional**
  - Not everything is lost

IC Design is tricky and demands attention

- **Even the simplest things can derail a complex chip**
  - A copy paste error in a bonding diagram, a mistake in reset
- **Academic research chips are not industrial products**
  - Designed to test and verify ideas, not mass production
  - Much more effort needed in DfT and verification to make a successful product
- **Experience is key in IC Design**
  - All the mistakes we make, add to our future success
  - Some lessons you learn the hard way
  - But these stay with you and help you for your future designs

Industrial Collaborations

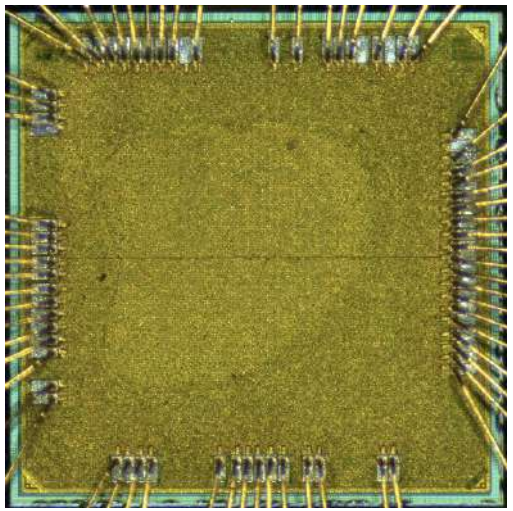

life.augmented

PULPv1,2,3 (ST28 FD)
Demonstrators of 28nm FD-SOI capabilities

Various publications ‘15 – ‘18

QuickLogic®

Arnold (GF22)
IoT SoC combining eFPGA with RISC-V core

Schiavone et al TVLSI ‘19

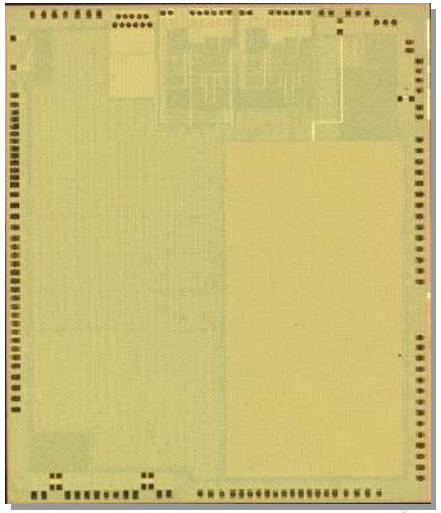

GREENWAVES TECHNOLOGIES

Vega (GF22)
IoT Processor with ML acceleration

Rossi et al ISSCC ‘21
Rossi et al JSSC ‘22

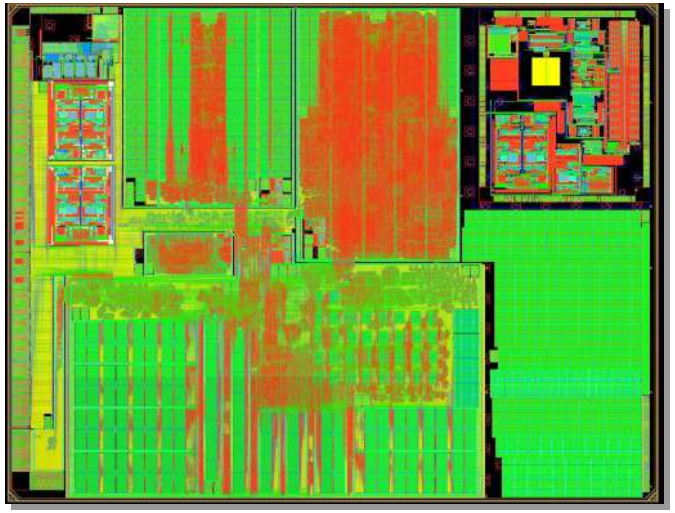

DOLPHIN INTEGRATION
The enabler of low-power Systems-on-Chip

Marsellus (GF22)
IoT Processor with low power modes and AI Accelerators

Conti et al ISSCC ‘23

Currently working with Meta, Intel, GF, IHP, PragmatIC, IIT

So proud to have supported others in their research

RISC-V week Barcelona 2018

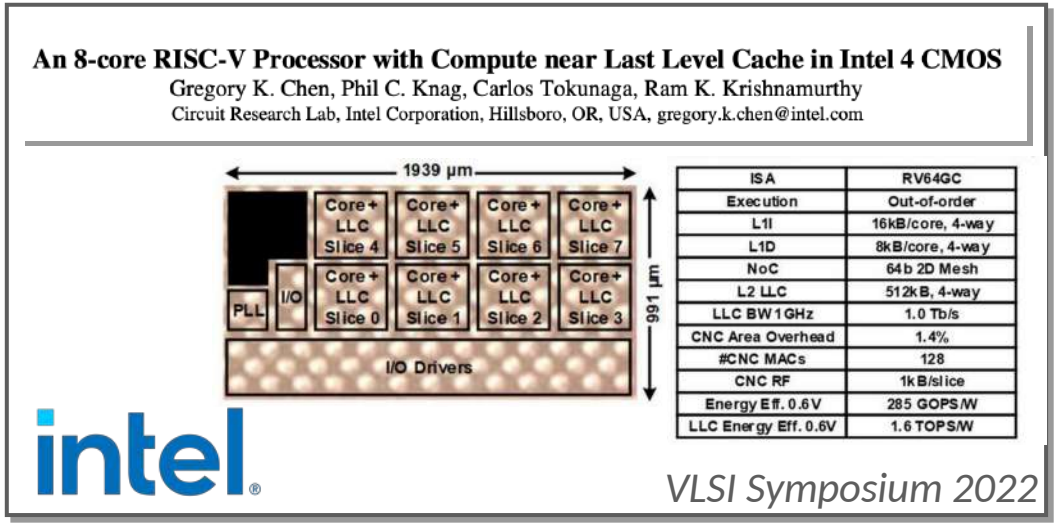

VLSI Symposium 2022

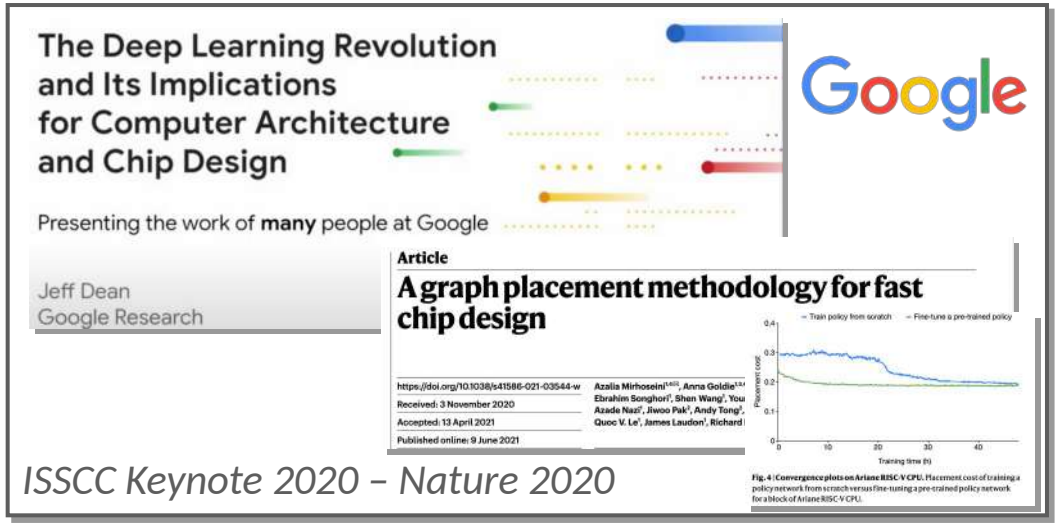

ISSCC Keynote 2020 – Nature 2020

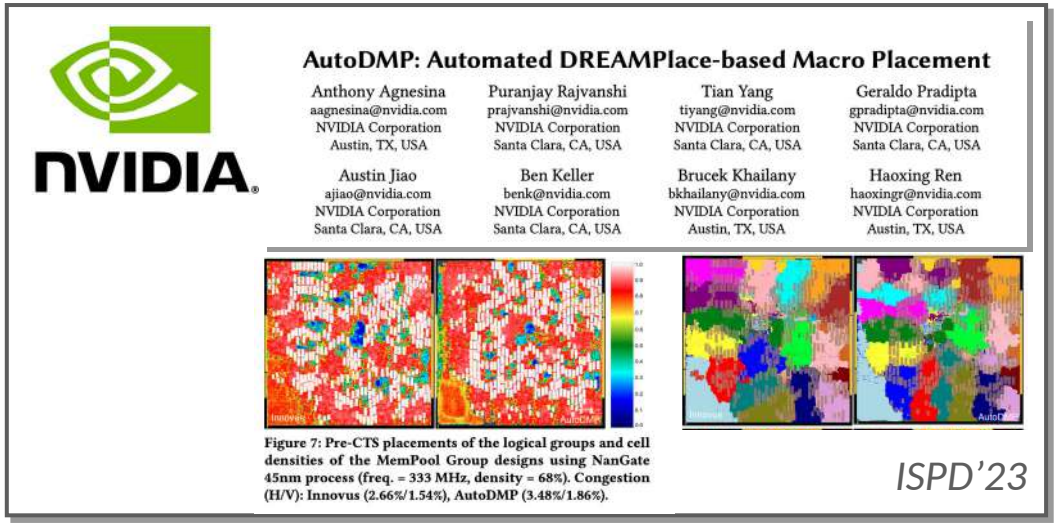

ISPD'23

Coming soon from the PULP team

FlexIBEX

IoT processor with a twist a kHz range design

Pragmatic

Iguana

Going all the way in open source

ihp

Carfield

Cars (and cats) can also use a bit of PULP

intel

Occamy

Bringing up 432 core chiplet in 12nm

55 PULP chips manufactured until now – more on the way

2013 (3)

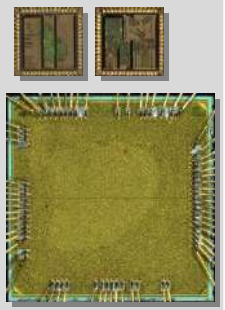

PULPv1
STM 28FDSOI
Multi-core processor

2014 (5)

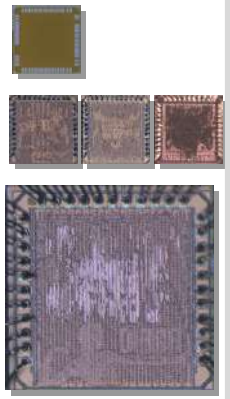

Diana
UMC 65
4-core system with approximate FPUs

2015 (10)

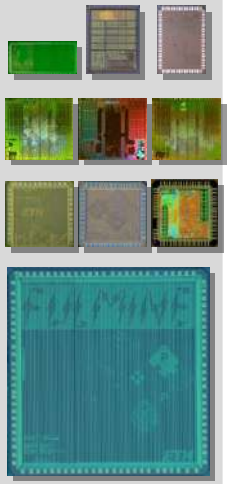

Fulmine
UMC 65
4-core system with ML and Crypto accelerators

2016 (3)

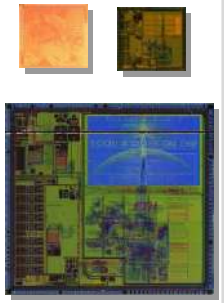

VivoSoC 2.001
SMIC 130
Mixed signal system for biosignal acquisiton

2017 (2)

Mr. Wolf
TSMC 40
8+1 core IoT processor

2018 (6)

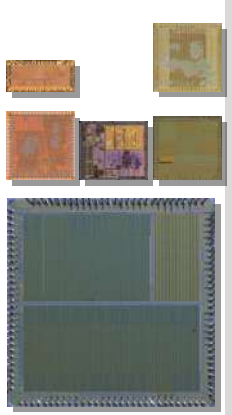

Poseidon
GF 22FDX
64bit RISC-V core, 32bit Microcontroller system, ML ...rator

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 (7)

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

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


http://pulp-platform.org
@pulp_platform
Frank K. Gürkaynak - 50+ ASICs in 10 years: a visual history