EPFL

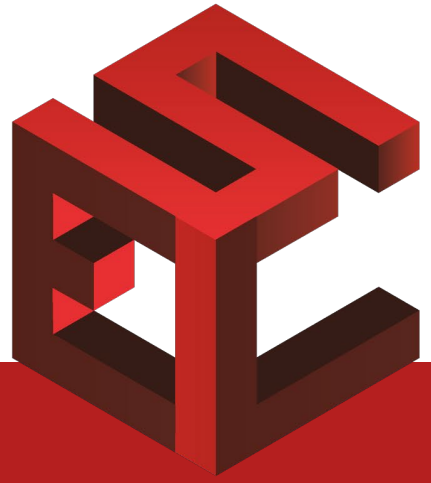

EMBEDDED SYSTEMS LABORATORY

PASC23 Davos Switzerland | 26-28 June 2023

Tradeoffs in Low-Power Accelerators Design for Large-Scale Interferometers

Denisa Constantinescu, Benoît Denkinger, Miguel Peon Quiros, David Atienza

EPFL: Embedded Systems Laboratory & EcoCloud & SKACH

denisa.constantinescu@epfl.ch

Outline

What brings me here?

- EcoCloud, ESL, SKACH

Motivation

- Green computing accelerators for SKAO

Challenges

- High throughput requirements
- Software pipeline not fixed and evolving
- Scalability

Trends & Tradeoffs

- Energy efficient HPC
- Energy efficiency vs performance vs programmability

HW-SW Codesign - Interferometry Domain

Discussion

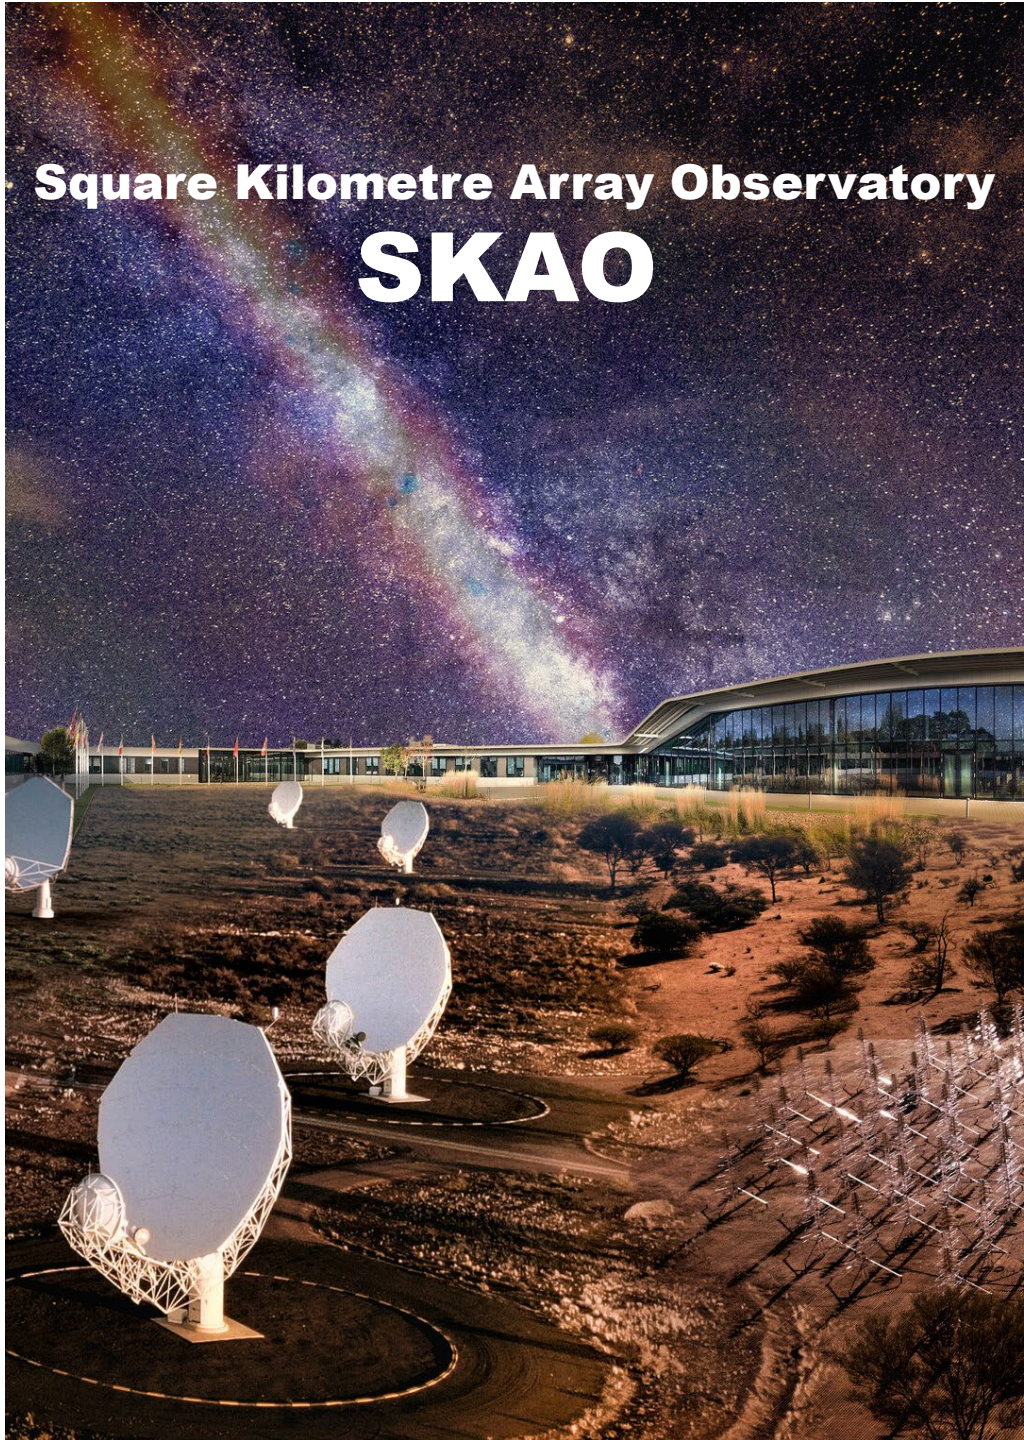

EcoCloud: Center for Sustainable Computing

Enabling environmental sustainability with IT

- Founded at EPFL in 2011
- 12 industrial affiliates, 29 faculty, 4 schools at EPFL
- 150+ researchers involved
- Experimental facility on sustainable computing

Three main themes of research & networking

- **Sustainable IT infrastructure**
- IT as enabler for sustainability (e.g.: UrbanTwin)
- Best practices in IT for a sustainable digital society

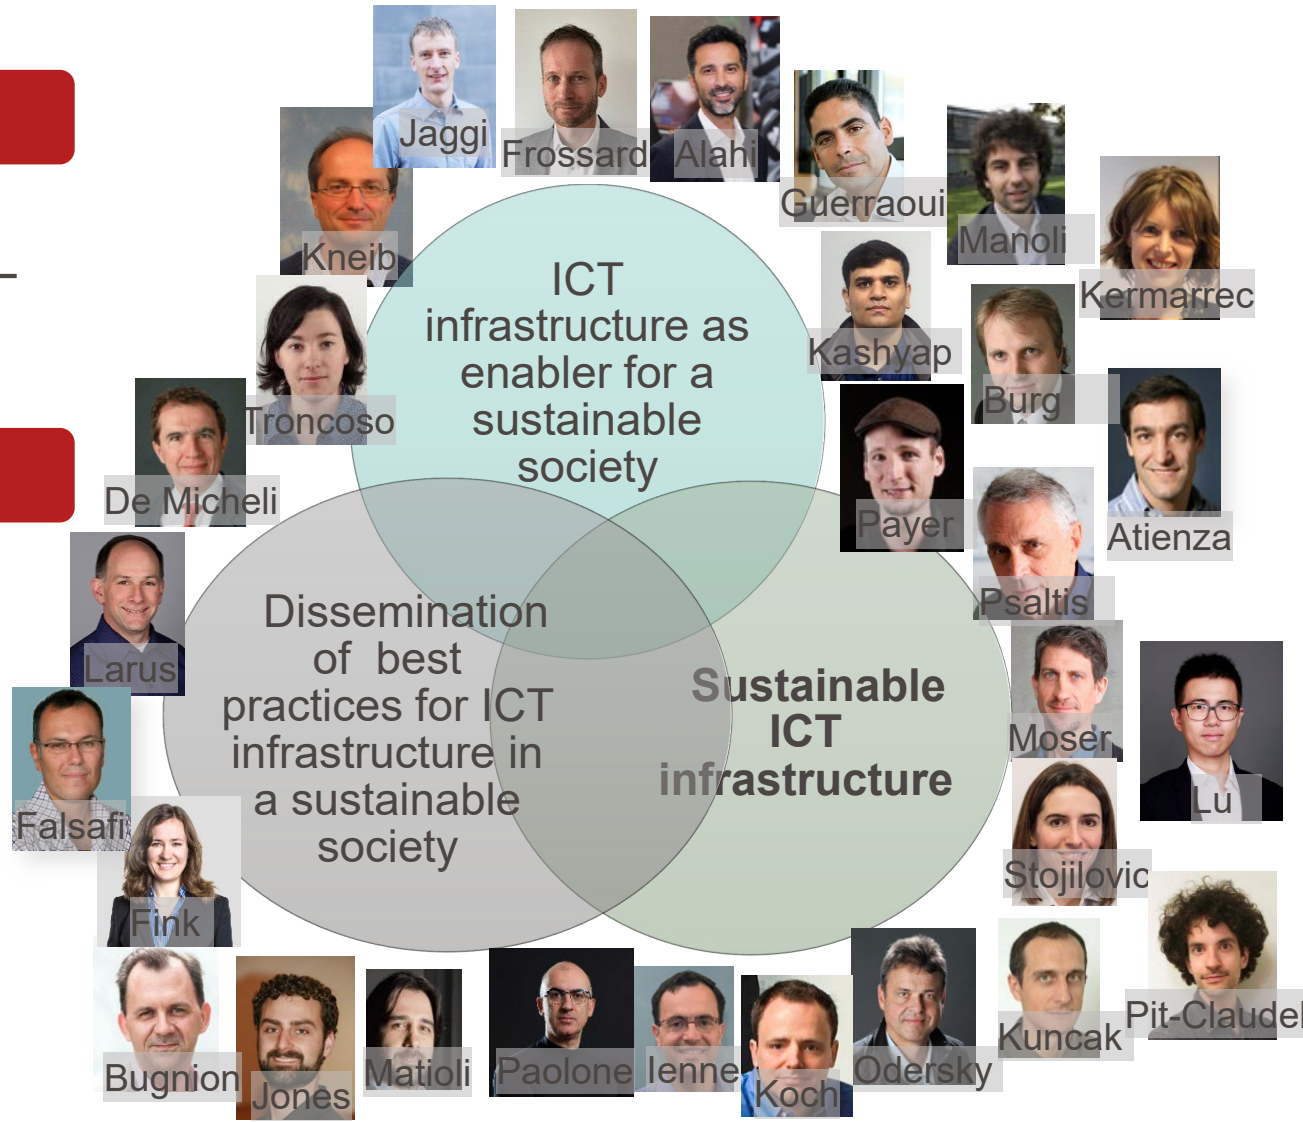

ecocloud.epfl.ch

ESL in a nutshell

Embedded systems and computer engineering

- Wearables and Internet-of-Things (IoT)
- **Low-power hardware-software co-design**
- **System-level design methodologies for HPC**

14 years at EPFL, 48 members today

- 8 post-docs (17 so far, 10 moved to industry)
- 30 PhD students (21 already graduated)

Innovation and tech. transfer: basic and applied research

- 31 companies provided grants and donations

Motivation for green acceleration

- SKAO will operate for 50+ years
- **Sustainability goals: SDP < 1 MWatt**
- **Science goals: high throughput req.**

Challenge

- Plan and design energy efficient data centers
- **Software pipeline not fixed and evolving**

Solution

- HW-SW codesign
- **Heterogeneous** computing
- **Reconfigurable** accelerators

Trends HPC – Heterogeneous & Energy Efficient

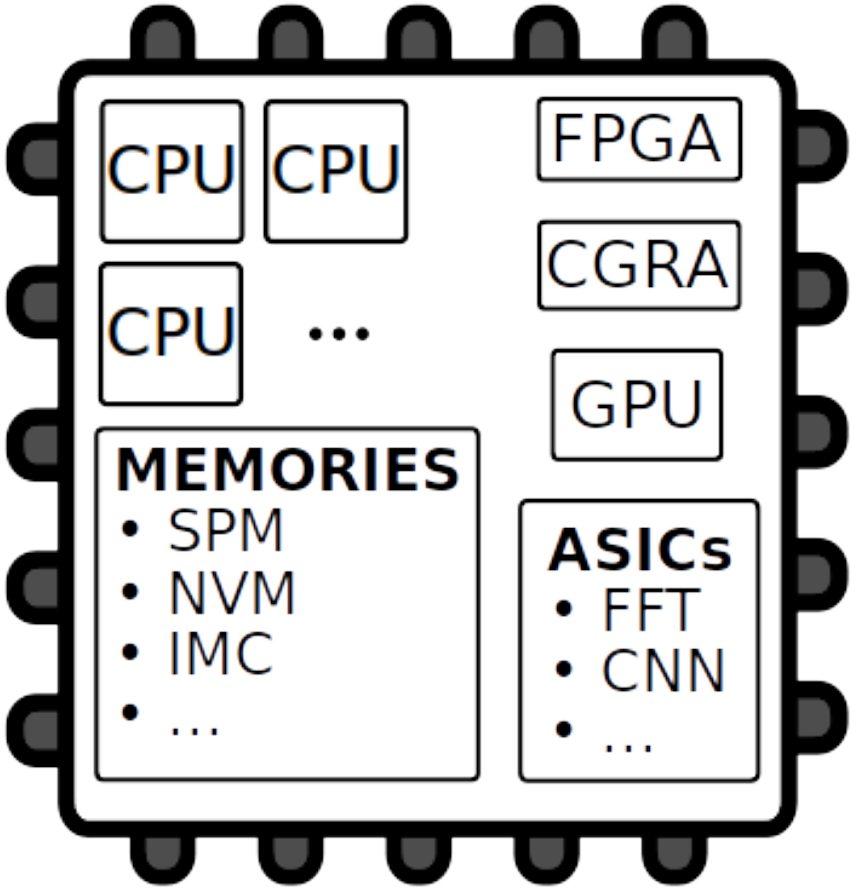

Heterogeneous systems computing

- **Multi-Core**
- **Hardware Accelerators**
  - Field Programmable Gate Array (FPGA)
  - Coarse Grained **Reconfigurable Architectures** (CGRA)
  - Graphic Processing Unit (GPU)
  - Application Specific Integrated Circuit (ASIC)
  - ...
- Memory Hierarchy & Near Memory Computing

Optimal design for **energy and throughput?**

Trends: Compute vs Power of Reconfigurable Architectures

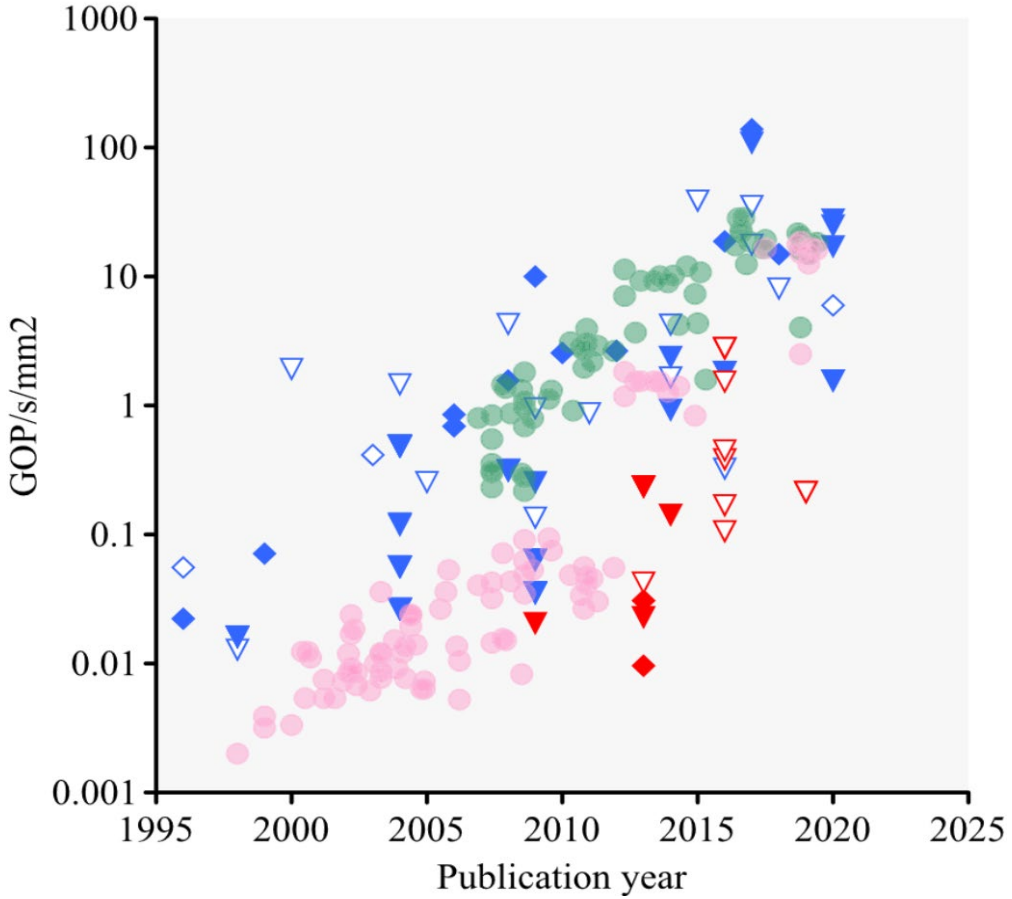

CGRA Compute Density

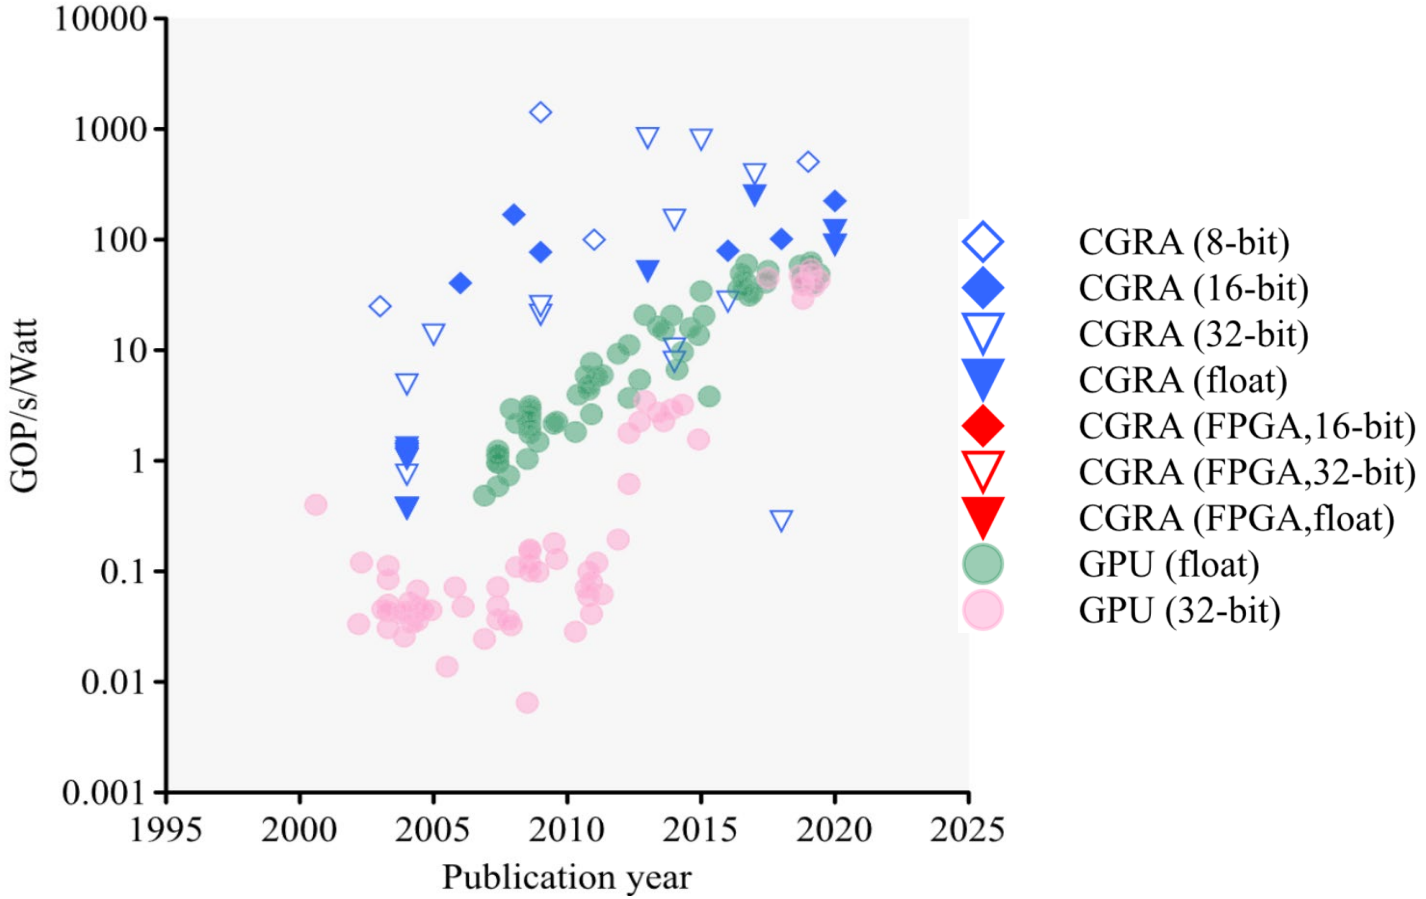

CGRA Compute Power Efficiency

Podobas, et. al. A survey on coarse-grained reconfigurable architectures from a performance perspective. *IEEE Access 8*, 2020.

Tradeoffs: Programmability vs Energy Efficiency

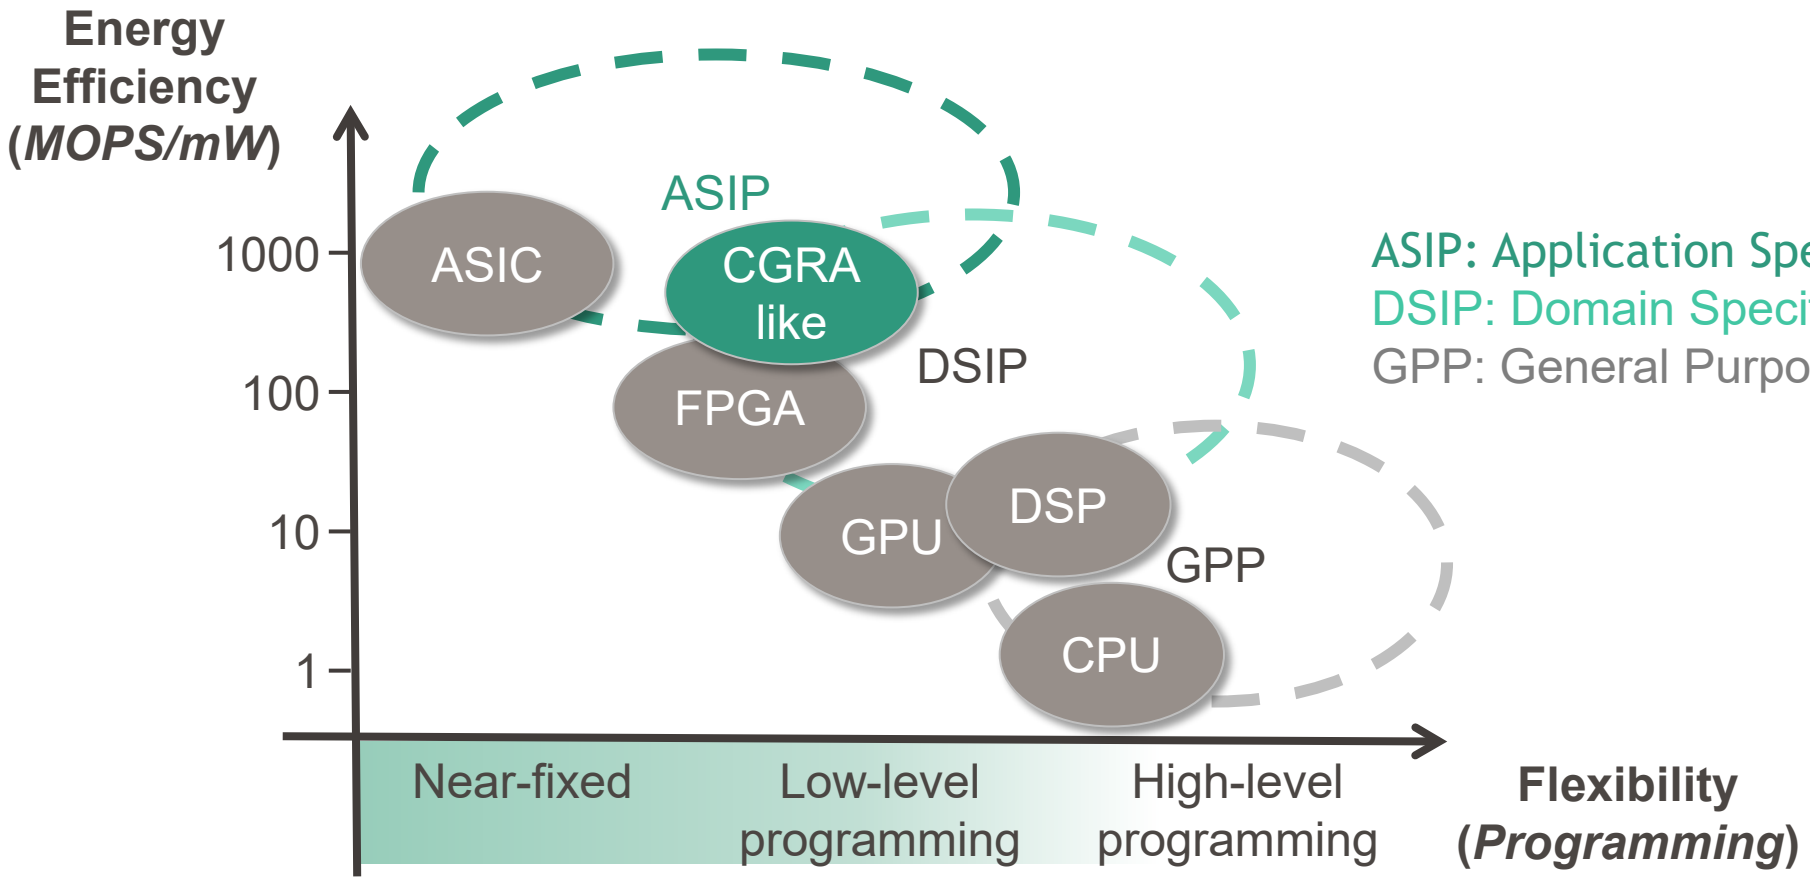

ASIP: Application Specific Instruction Processor
DSIP: Domain Specific Instruction-set Processor
GPP: General Purpose Processor

Adapted from *Kevin J M Martin. IPDPSW 2022*

Quest for green acceleration: HW-SW codesign

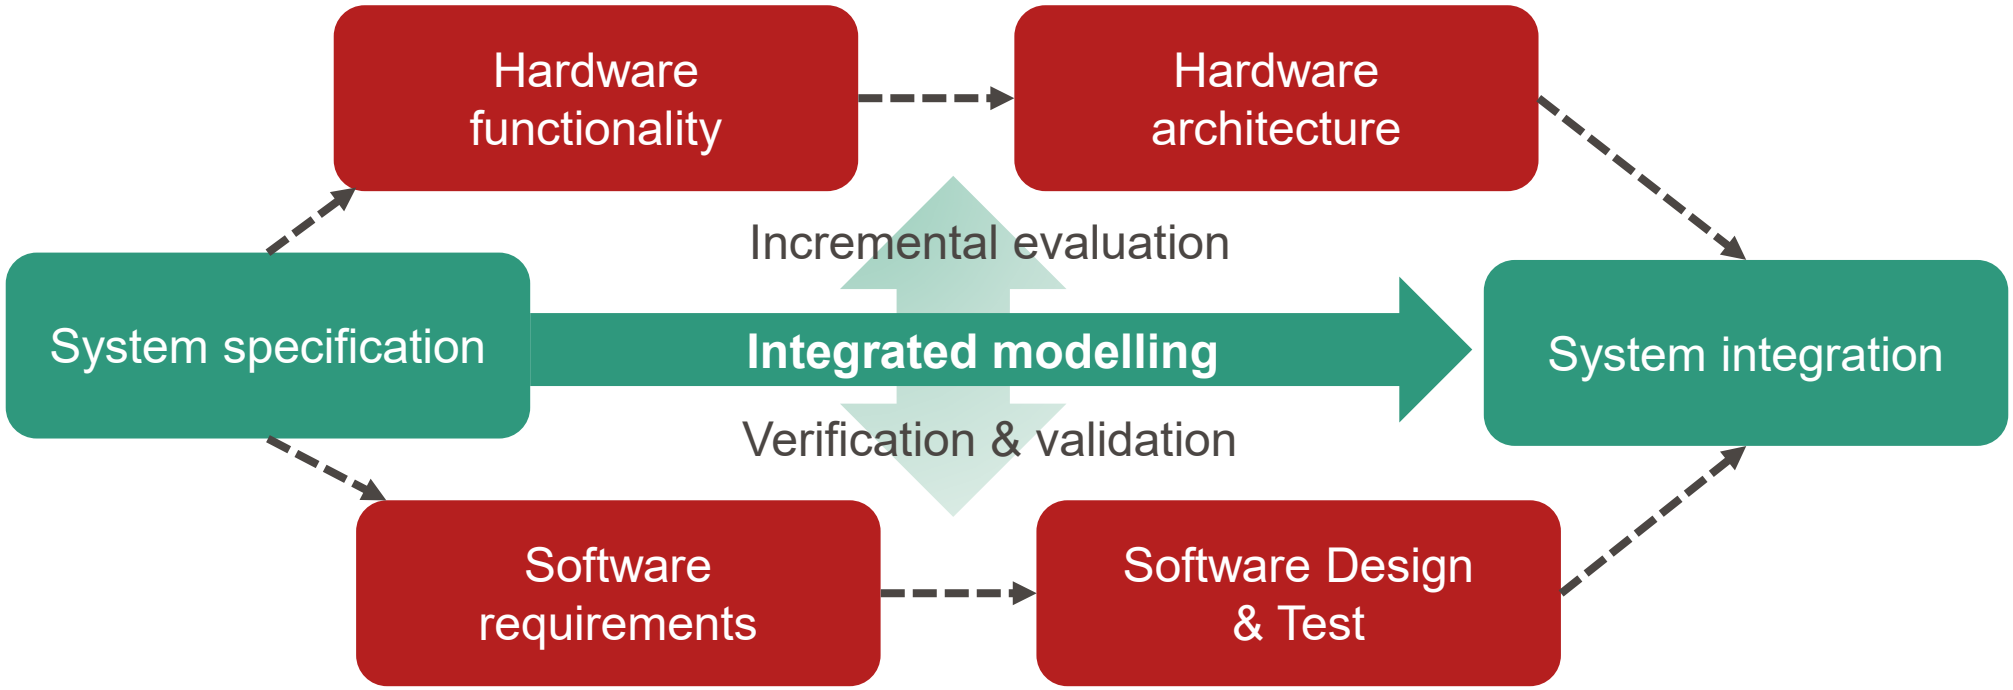

3 paths HW-SW codesign, given starting point

1. Only **specifications** - open choice for a model
2. Existing **software** implementation
3. Existing **hardware** (chip – CPU, GPU)

We only have pieces of the puzzle from each

Specifications SKAO (Partial)

- Science data processor **energy budget of 1 MWatt** (x2 sites)
  - Long term feasibility?
  - Risk of power shortages: need for green computing

- **High throughput requirements**
  - **8.8 Tb/s** [Mid-frequency array]
  - **7.2 Tb/s** [Low-frequency array]
  - **data dependency**: seconds → hours → days worth of data
    - needs to be processed in **real-time**

- **HPC at Scale**
  - Dynamic workload
  - Highly heterogeneous data center infrastructure

HW will not change much in the next years

but SW pipeline is not fixed & evolving

→ **Reconfigurable** accelerators

Tradeoffs: Programmability vs performance

Accuracy – problem dependent

Software: Imaging Pipeline Bottleneck

Visibility domain

Image domain

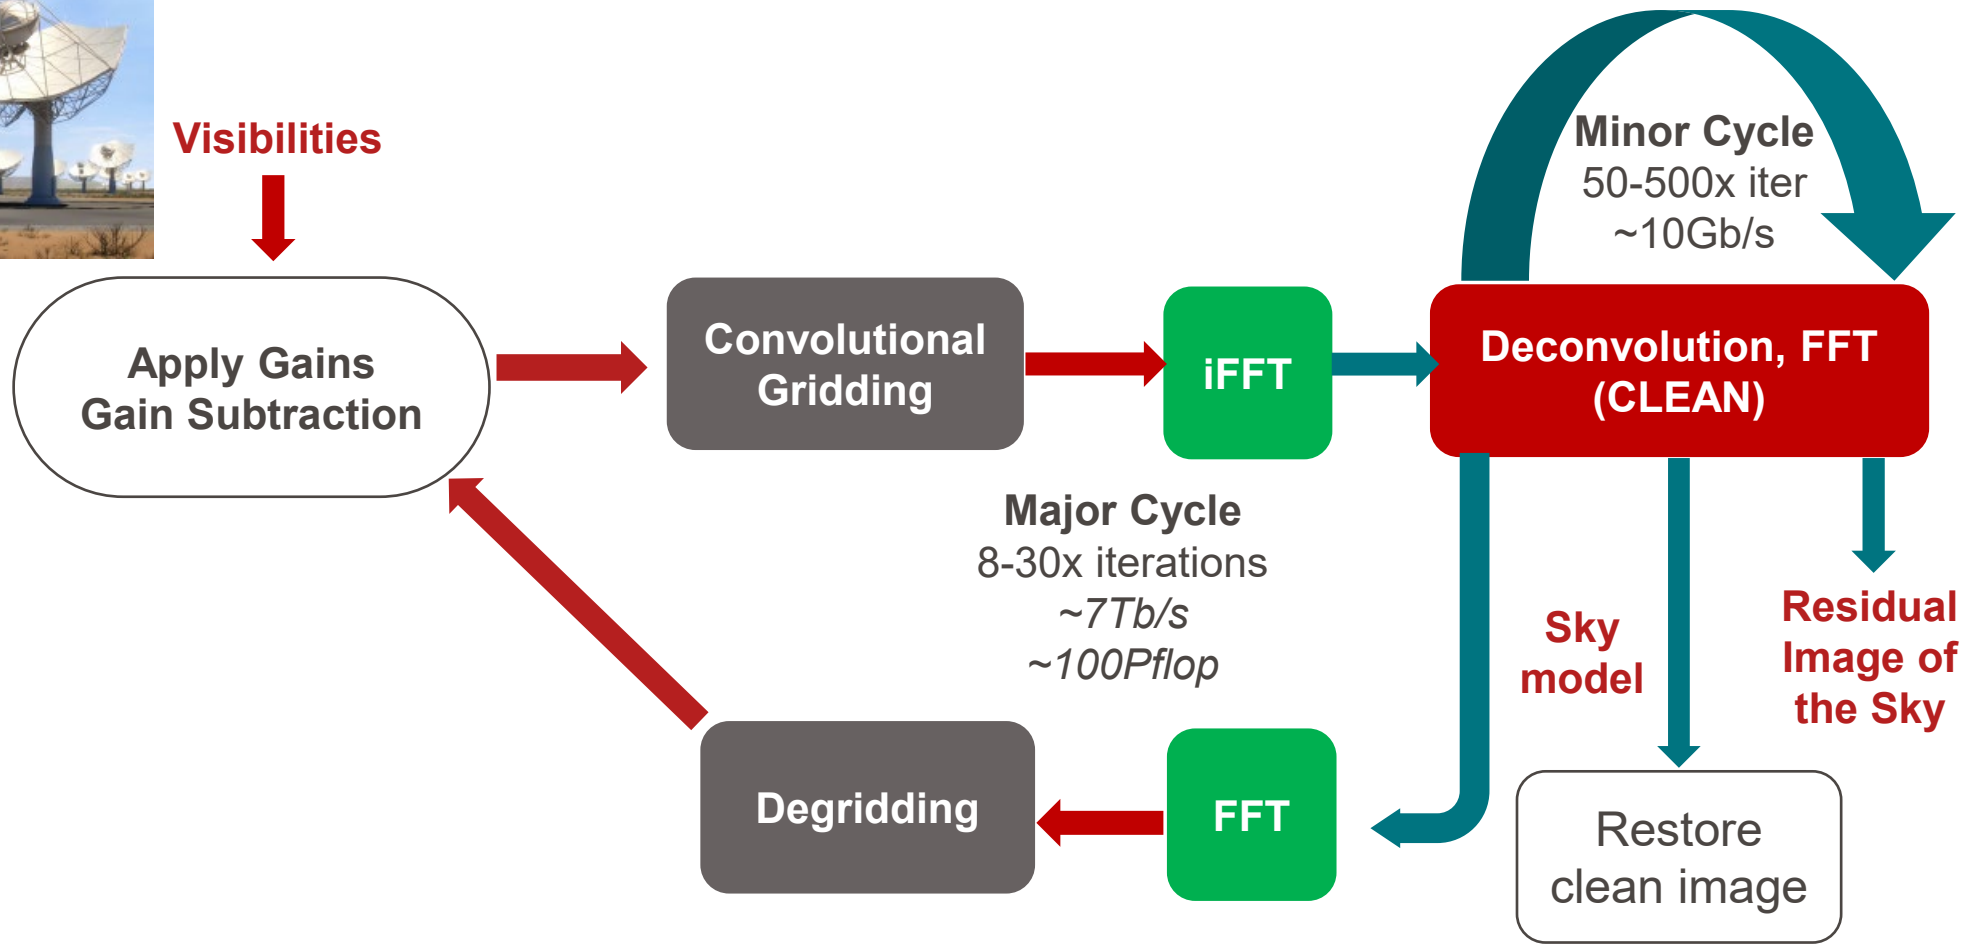

- Numbers based on MeerKAT
- – precursor of SKA Mid, Karoo desert

Image domain: typical interferometry pipeline steps

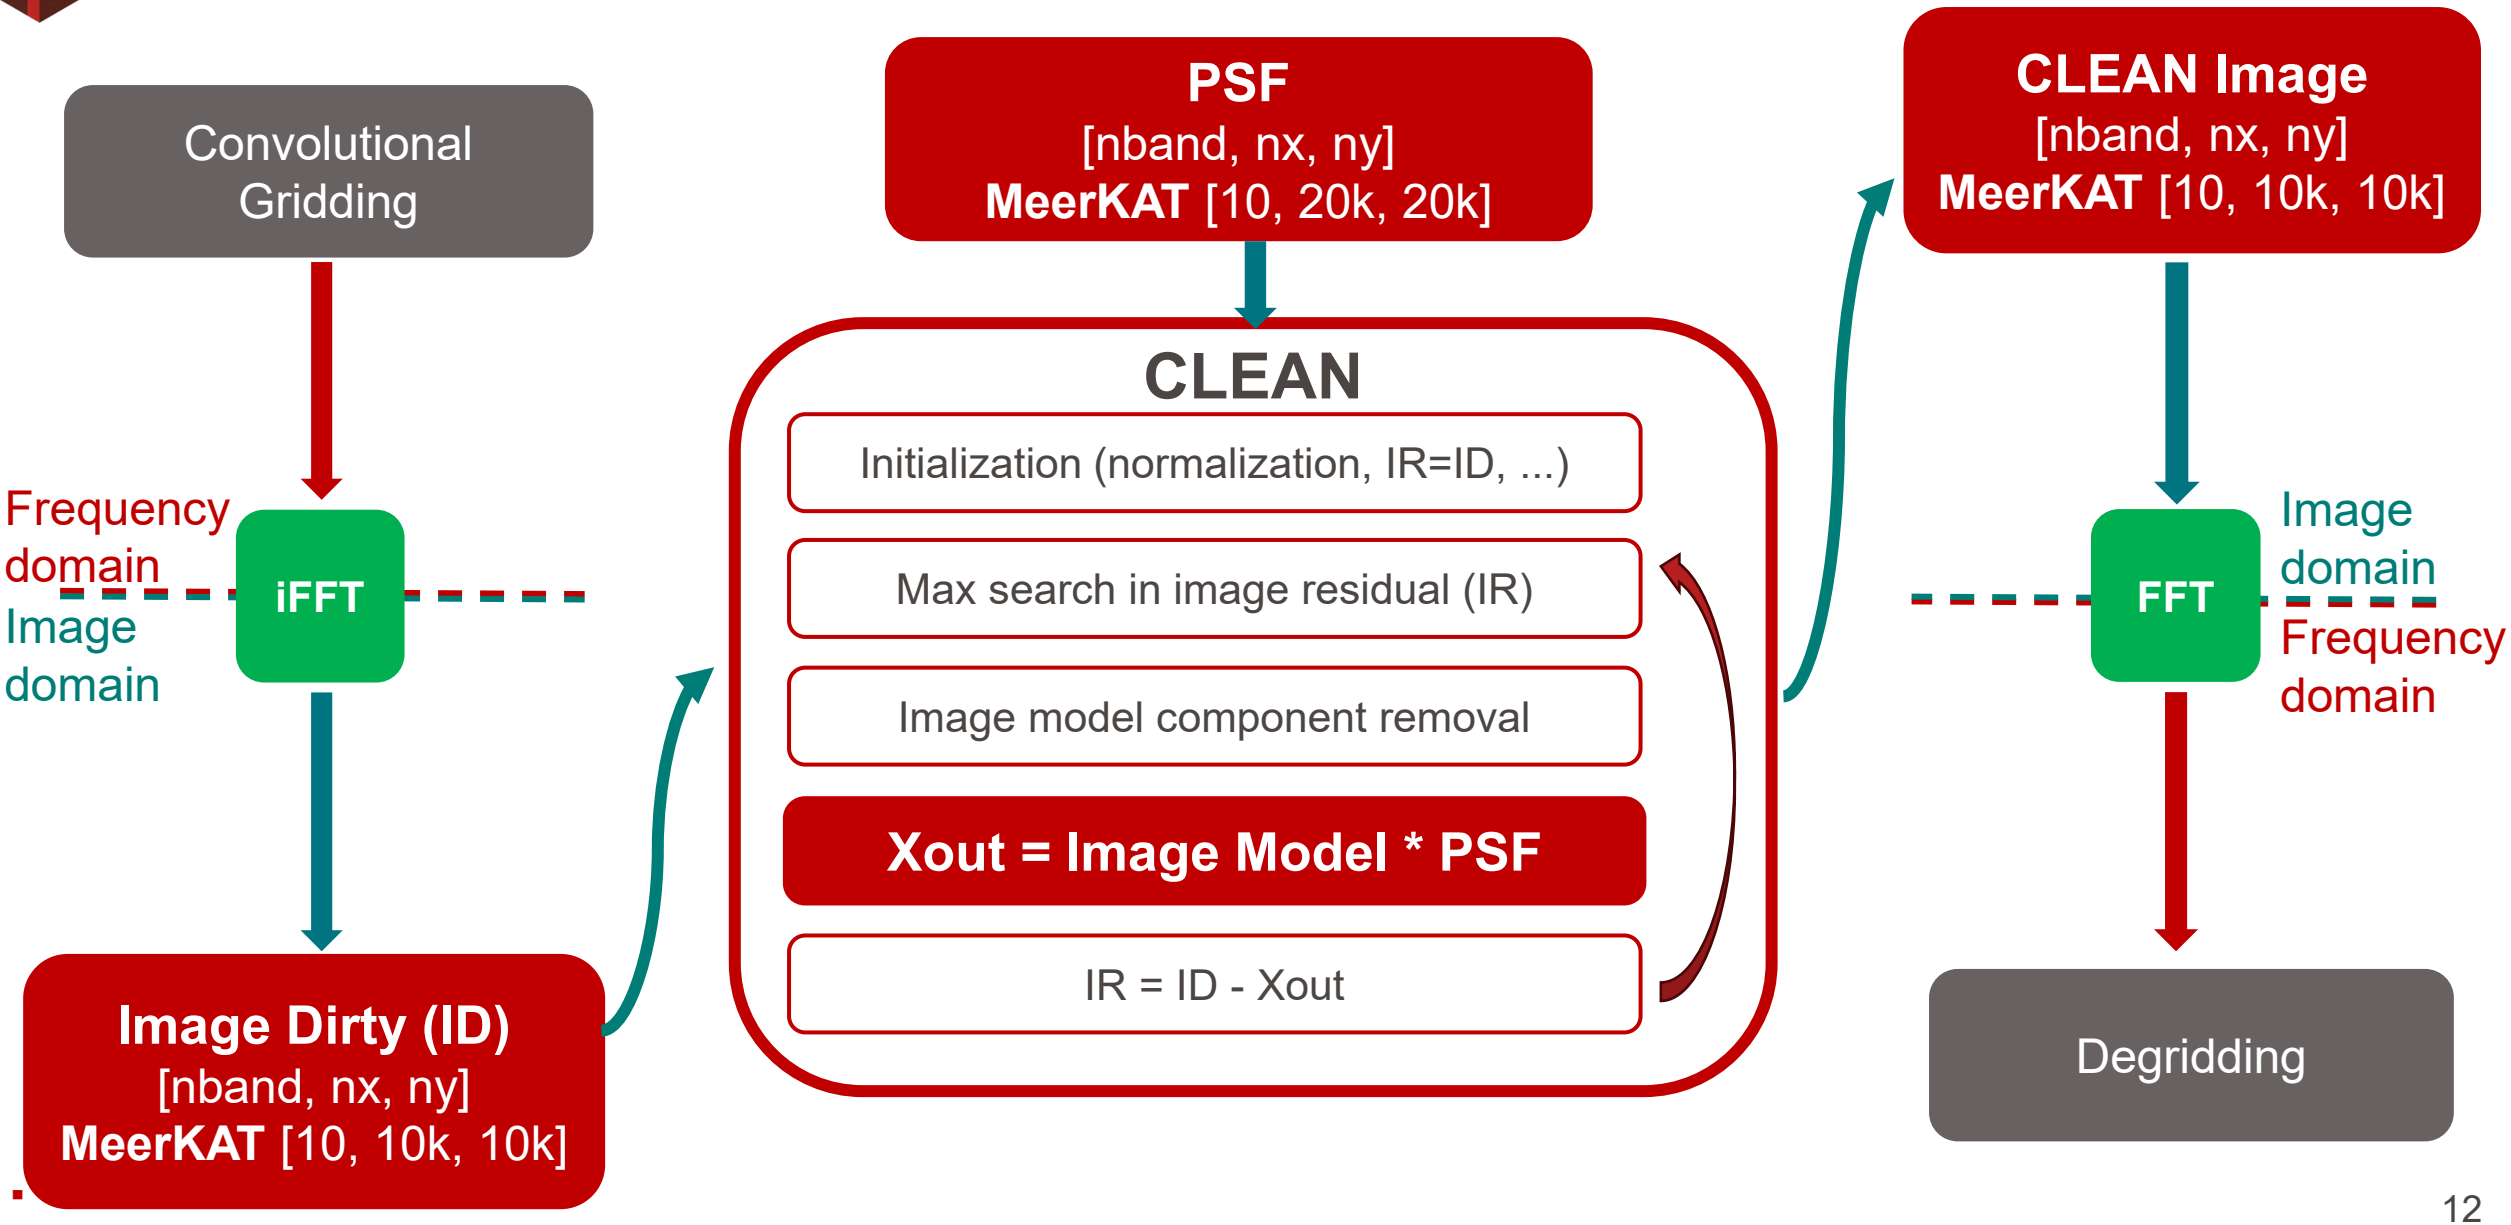

Hardware tradeoffs

- **PPA: Performance** vs **Power** vs **Area** [~cost chip]

- **E.g: communication between Processing Elements (PE)**
  - Discrete routers use 6x more resources for comms [more area!]
  - BUT have up to 78% better utilization of PE [better performance]
  - Energy?

- **Memory architecture considerations**
  - Mem hierarchy?
  - Networked mem?
  - Unified shared memory?
  - Streaming data?
  - Memory coalescence?

to be studied!

Multi-scale co-design process

- ✓ Science objectives
- ✓ Sustainability goals

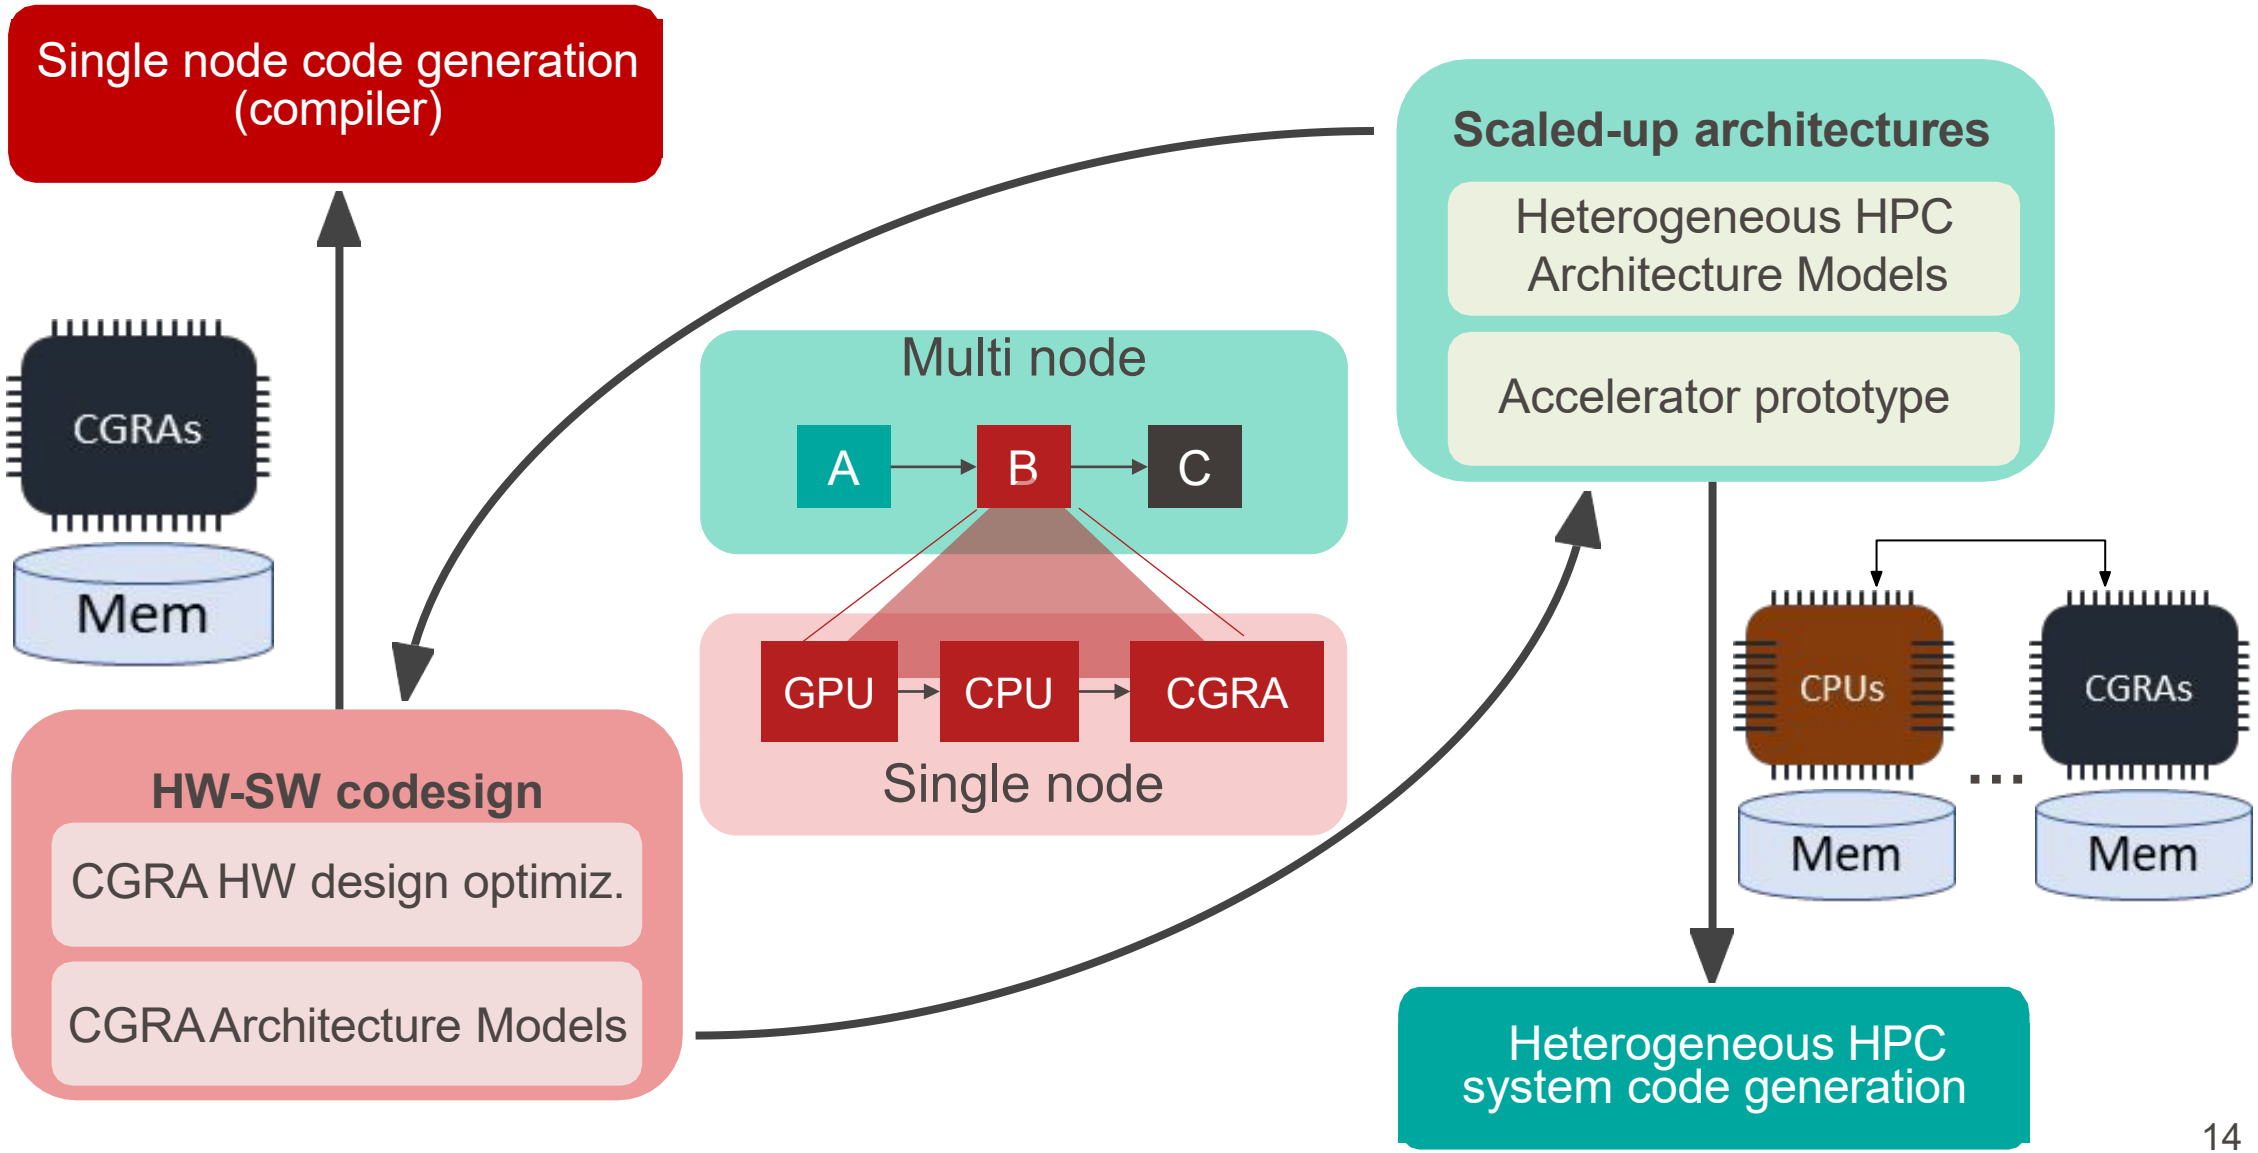

Next steps: open-source design frameworks

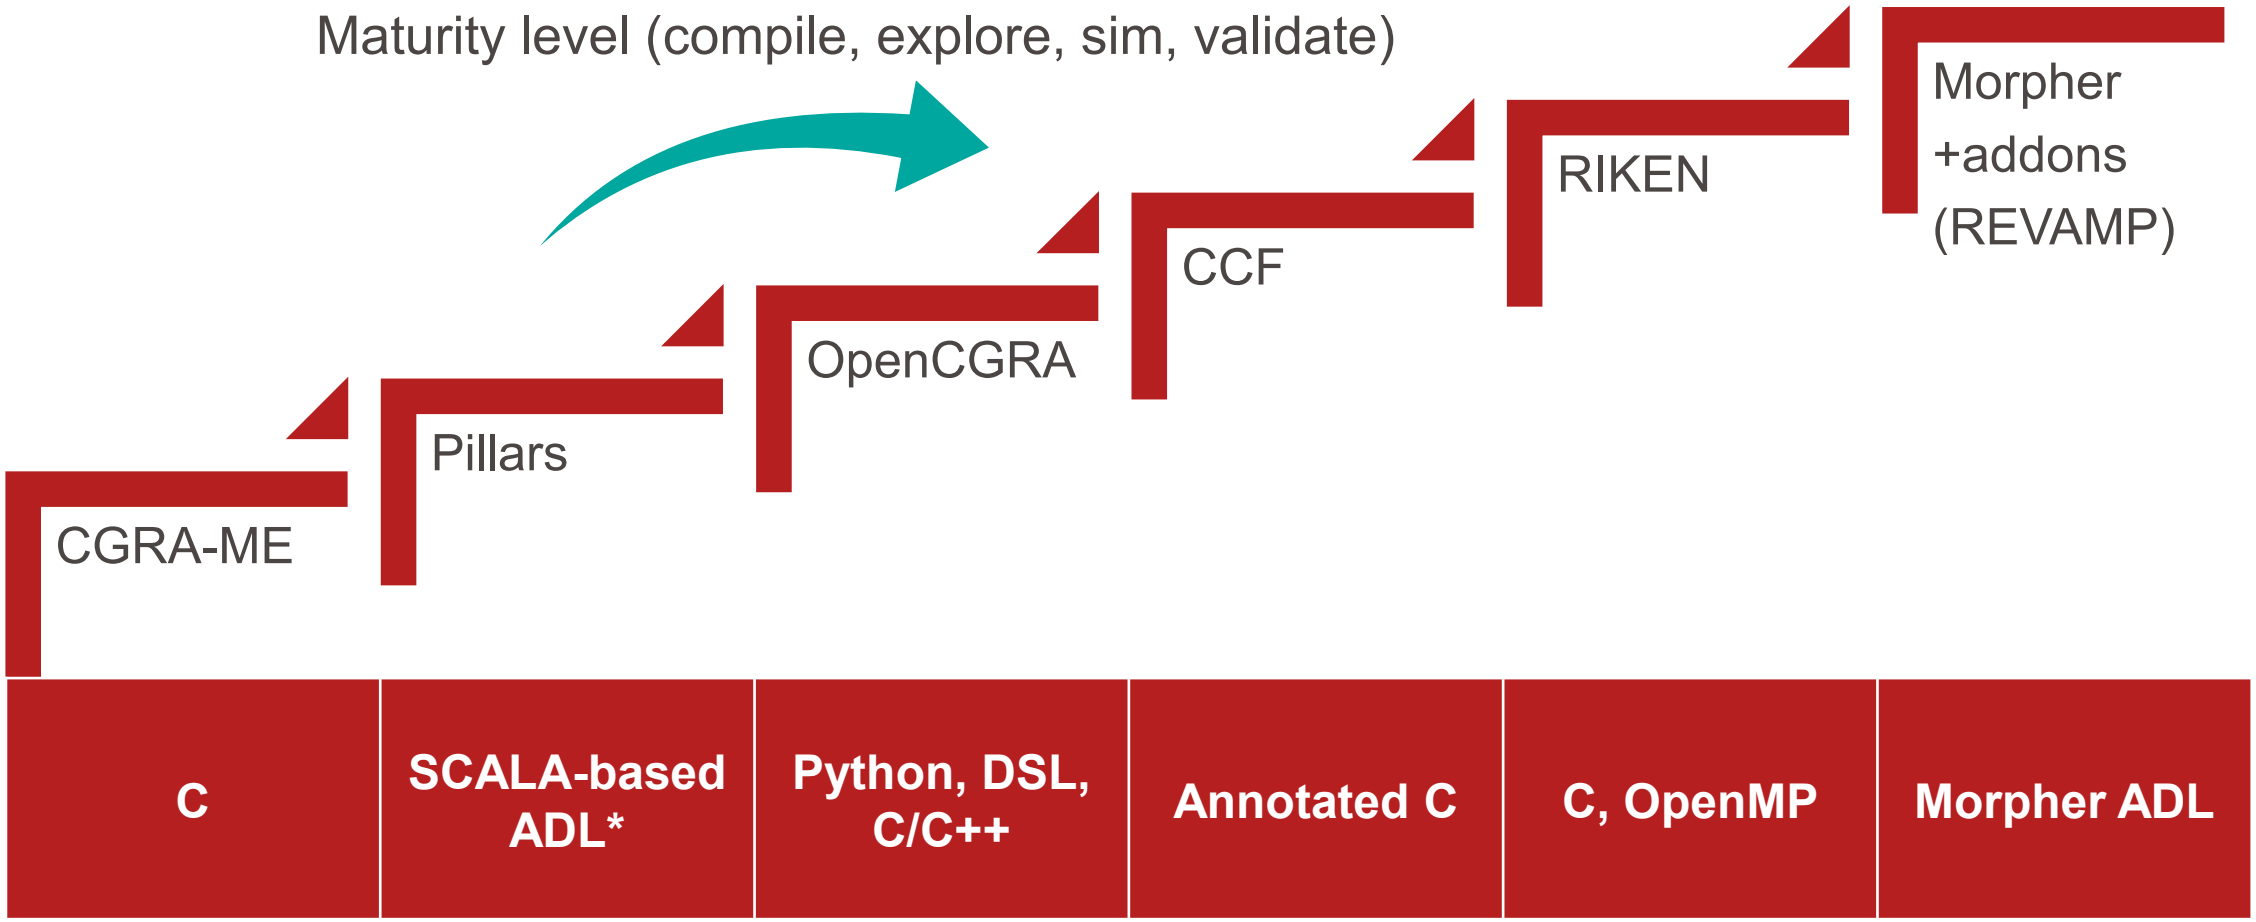

*Architecture Description Language (ADL)

Discussion

- **Sustainable computing** + **high throughput requirements**
  - Risk of power shortages: **need for green computing**
  - 8.8 Tb/s (Mid-frequency array) + 7.2 Tb/s (Low-frequency array)
  - Data dependency: seconds to days-worth of data needs to be processed in **real-time**

- **Software pipeline not fixed and evolving**
  - => need reconfigurable/programmable accelerators
  - **Tradeoffs**: flexibility vs energy efficiency vs high throughput vs portability

- **Single node – design decisions**
  - **Extendable architecture:** likelihood of needing new types of computation?
  - **Parametric architecture**: IMG size, accuracy?

- **Multi node: scalability challenge**
  - Dynamic workload – **many different pipelines** – 1 solution for all?
  - Highly heterogeneous data center infrastructure
  - Heterogeneous execution framework – work in progress: single vs multi-node solution

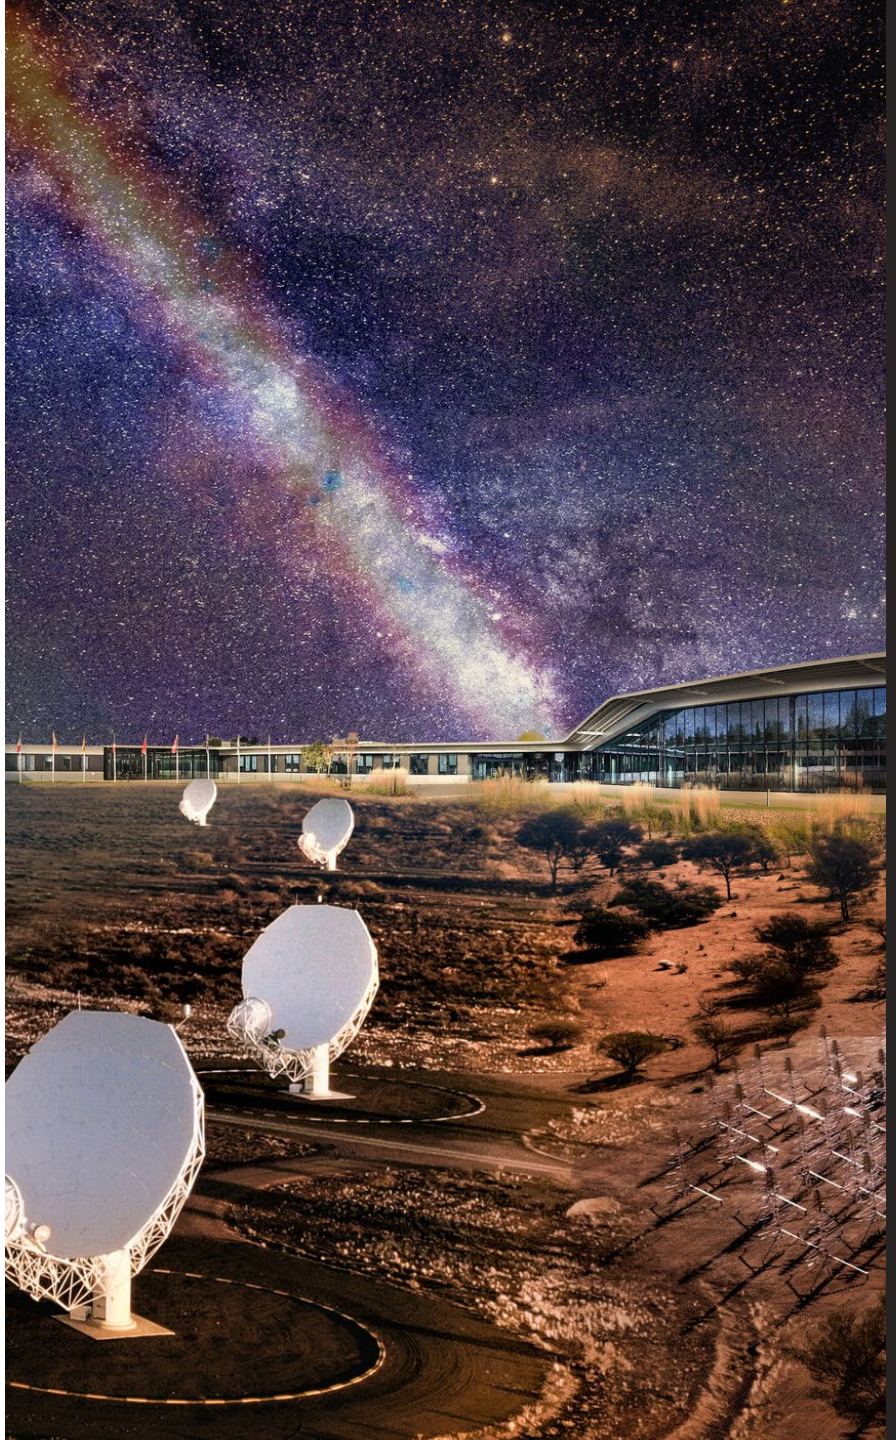

Thank you!

Denisa Constantinescu
EPFL: Embedded Systems Laboratory

denisa.constantinescu@epfl.ch