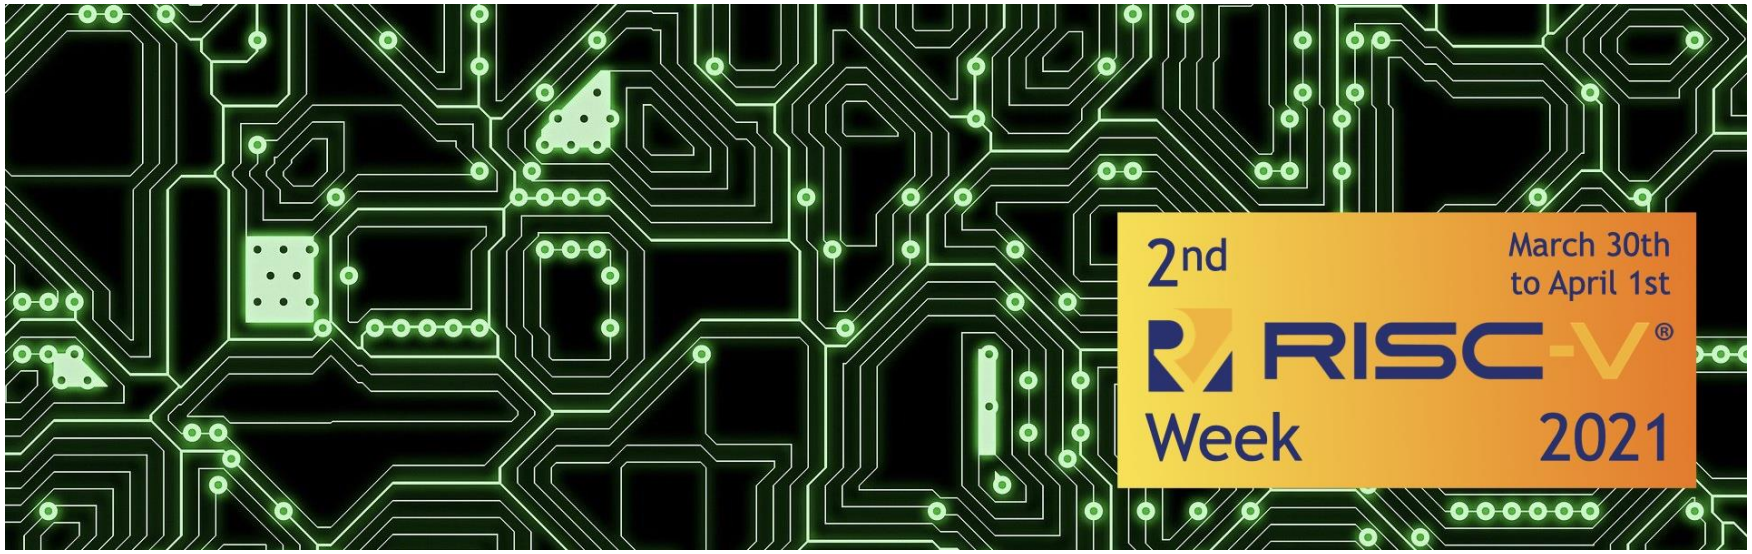

BARE METAL: CACHE COHERENCY INTERCONNECT AND CHIPLET PARTITIONING FOR HETEROGENEOUS MULTI-RISC-V CORES

Nom : Denis Dutoit – CEA LIST
Date: March 30th, 2021

Research focus

Design Platform Facilities:
Embedded Systems
Integrated Circuits

Aiming
to use Open Source
and contribute

High
Performance
SoC

Performance
Frugality
Autonomy
Safety &
Security
Scalability

Cyber
Physical
Systems

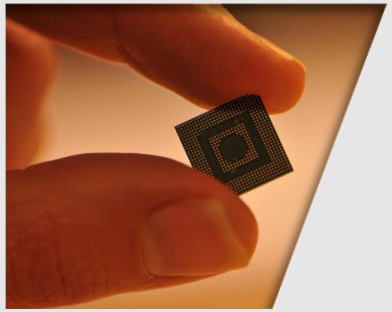

Community member

RISC-V
RISC-V: The Free and Open RISC Instruction Set Architecture

Platinium member

GROUP OPENHW
PROVEN PROCESSOR IP

Strategic member

ECLIPSE FOUNDATION

RISC V innovation partner

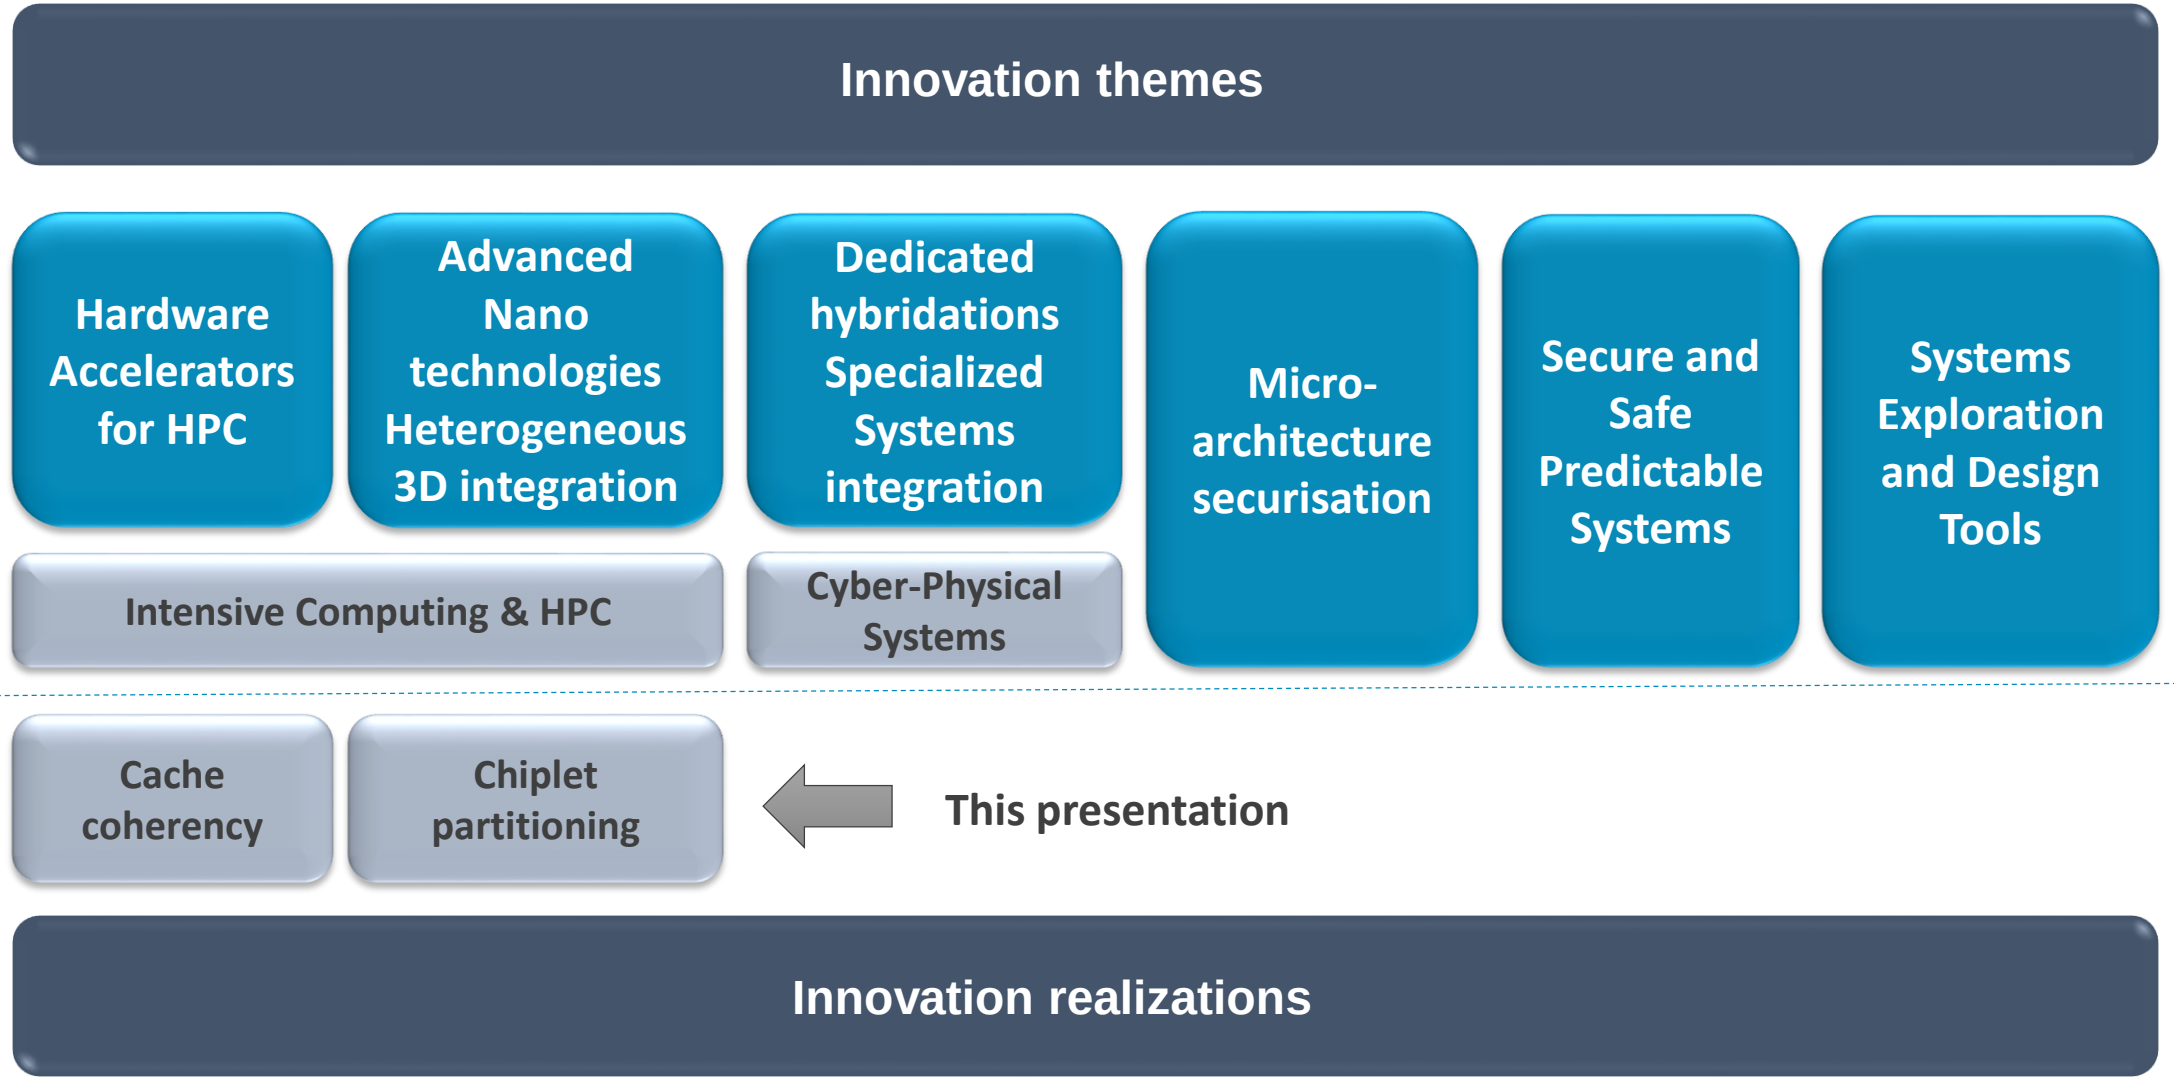

OUTLINE

- How to get « bare metal » performances from Silicon ?
- Cache coherency
- Chiplet partitioning
- Conclusion

HOW TO GET BARE METAL PERFORMANCES FROM SILICON ?

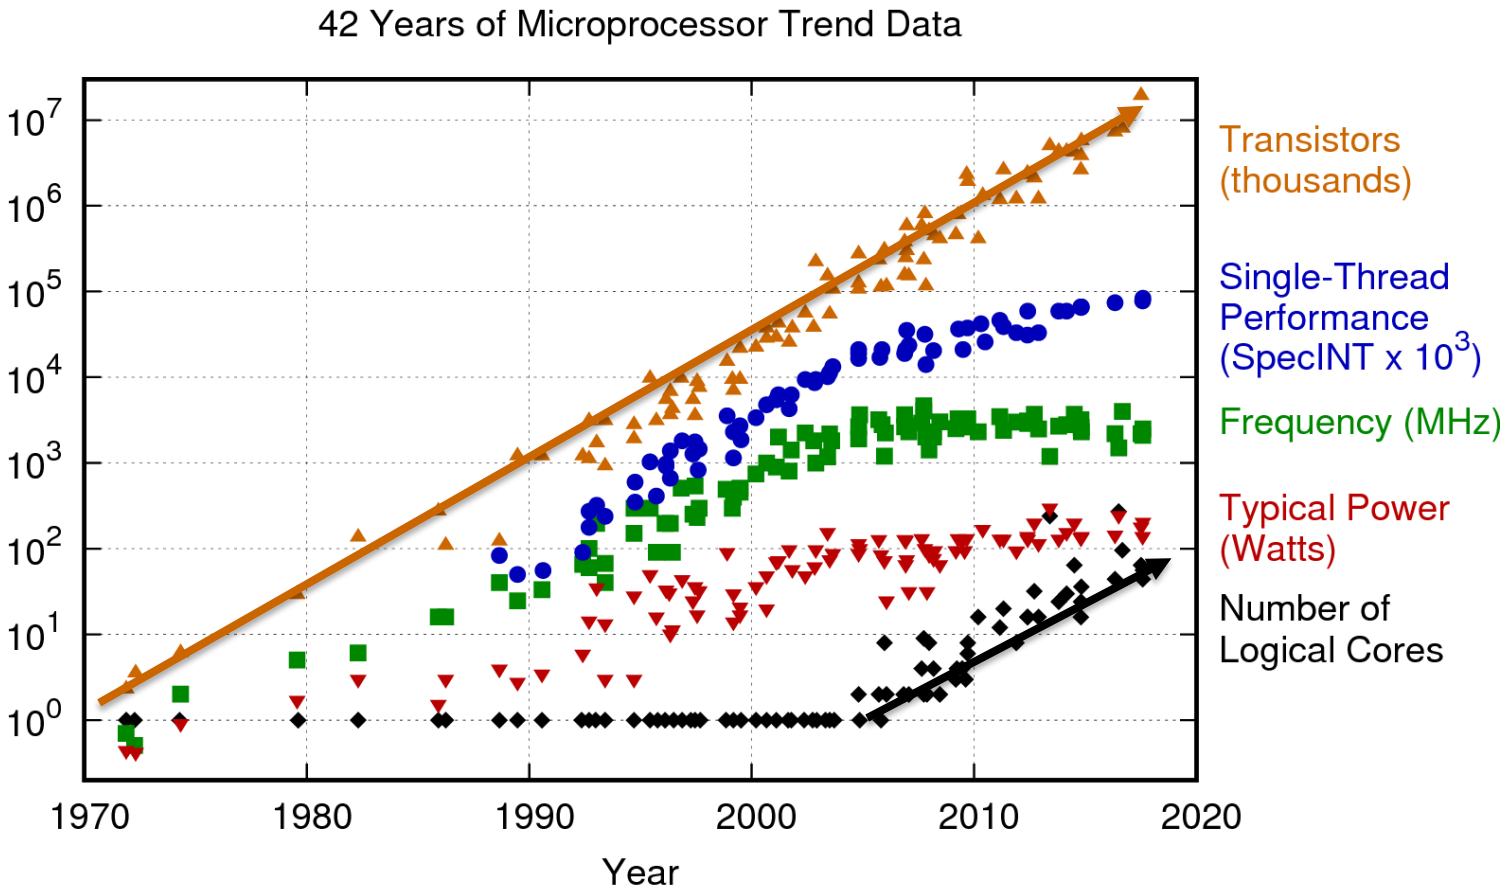

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp https://goo.gl/bb6wZW

Moore's law: transistor density is still increasing

- **Many-core architecture**

HOW TO GET BARE METAL PERFORMANCES FROM SILICON ?

→ At reasonable Total Dissipated Power (TDP)

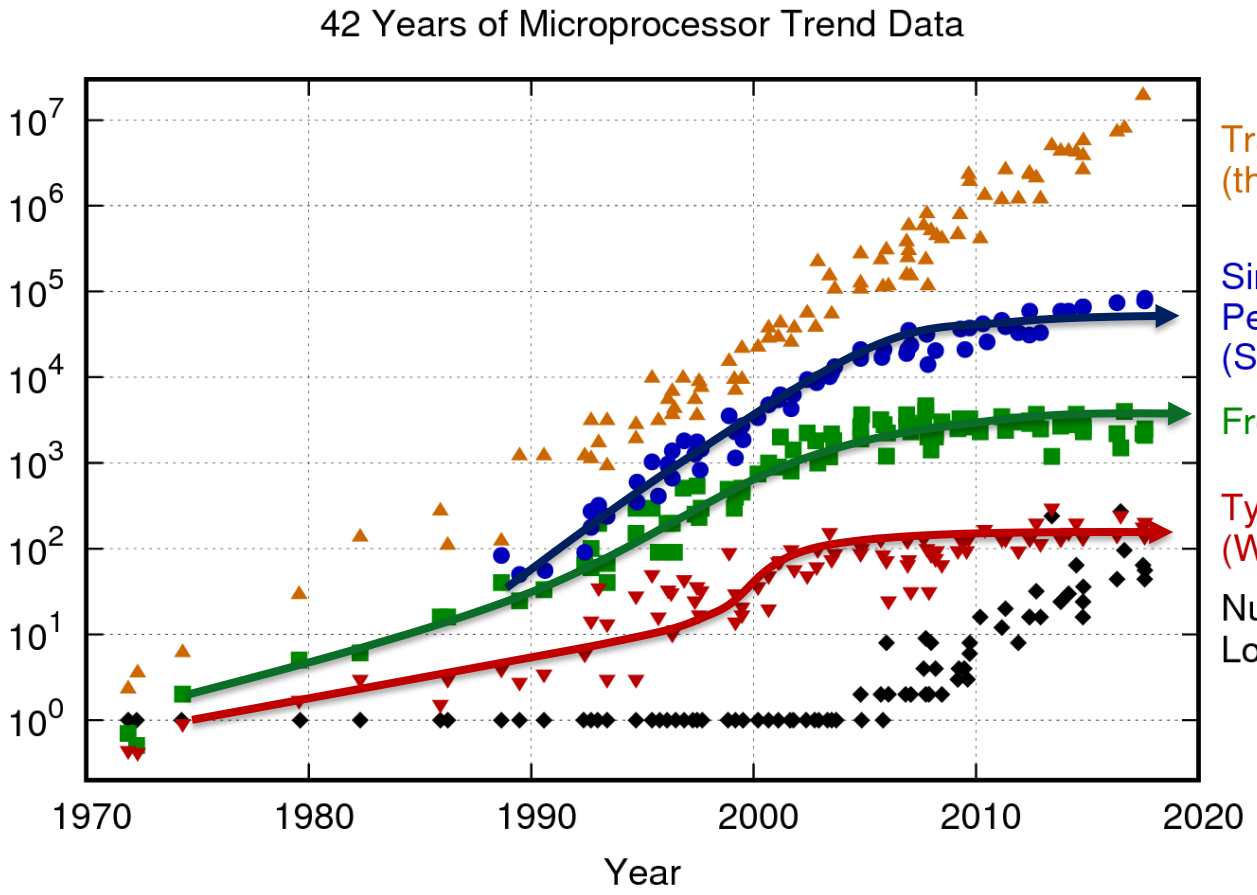

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp https://goo.gl/bb6wZW

Dennard's Law (as transistors shrink, they get faster, use less power) ended.

➢ **Specialized hardware (accelerators)**

HOW TO GET BARE METAL PERFORMANCES FROM SILICON ?

→ At reasonable development cost/time & manufacturing cost

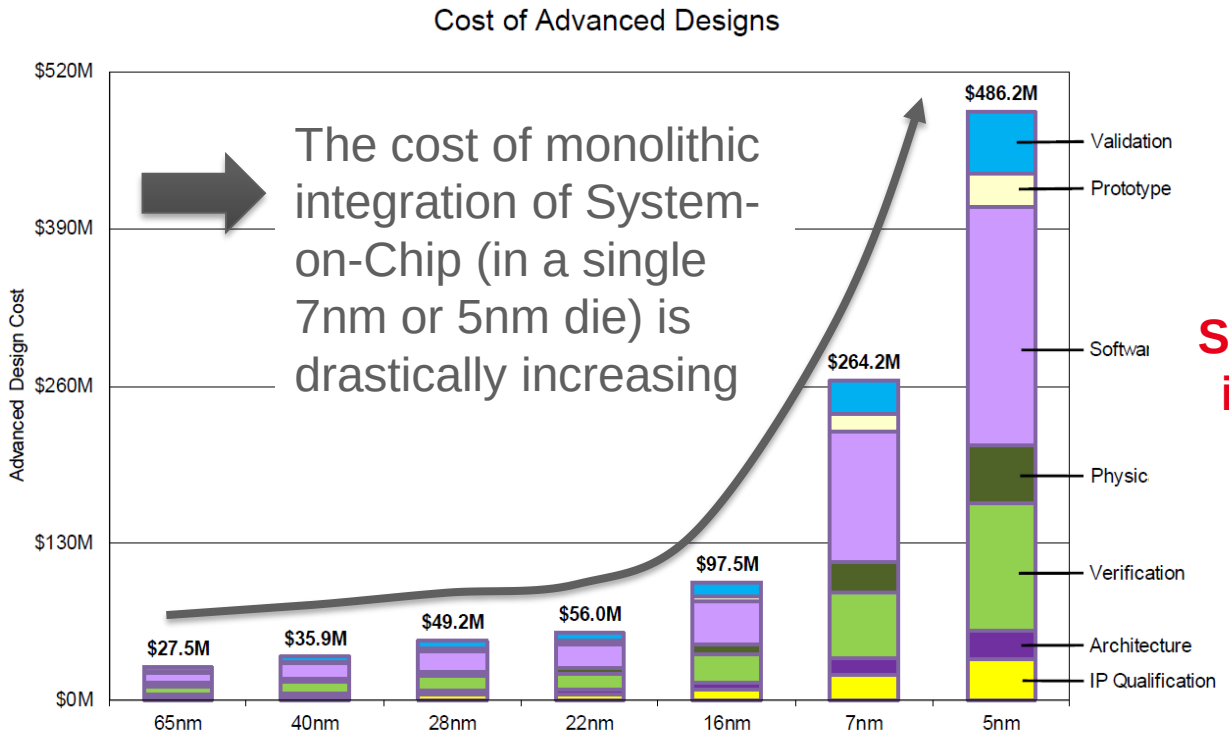

Source IBS 2020

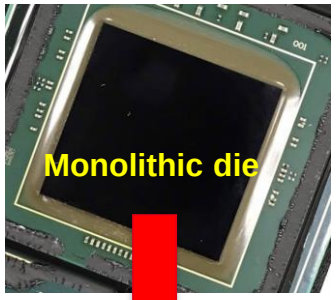

SoC disaggregation into 3D Integrated Circuit

- **Chiplets**: heterogeneous compute is mapped on multiple dies to reduce manufacturing costs (higher yield) and development cost/time (reusability, scalability)
- **Active Silicon Interposer**: integrating I/O functions that are more difficult to scale

HOW DOES IT FIT IN A COMPUTING CHIP ?

- **Many-core architecture**

RISC-V is a unique opportunity to leverage heterogeneous many-core architectures

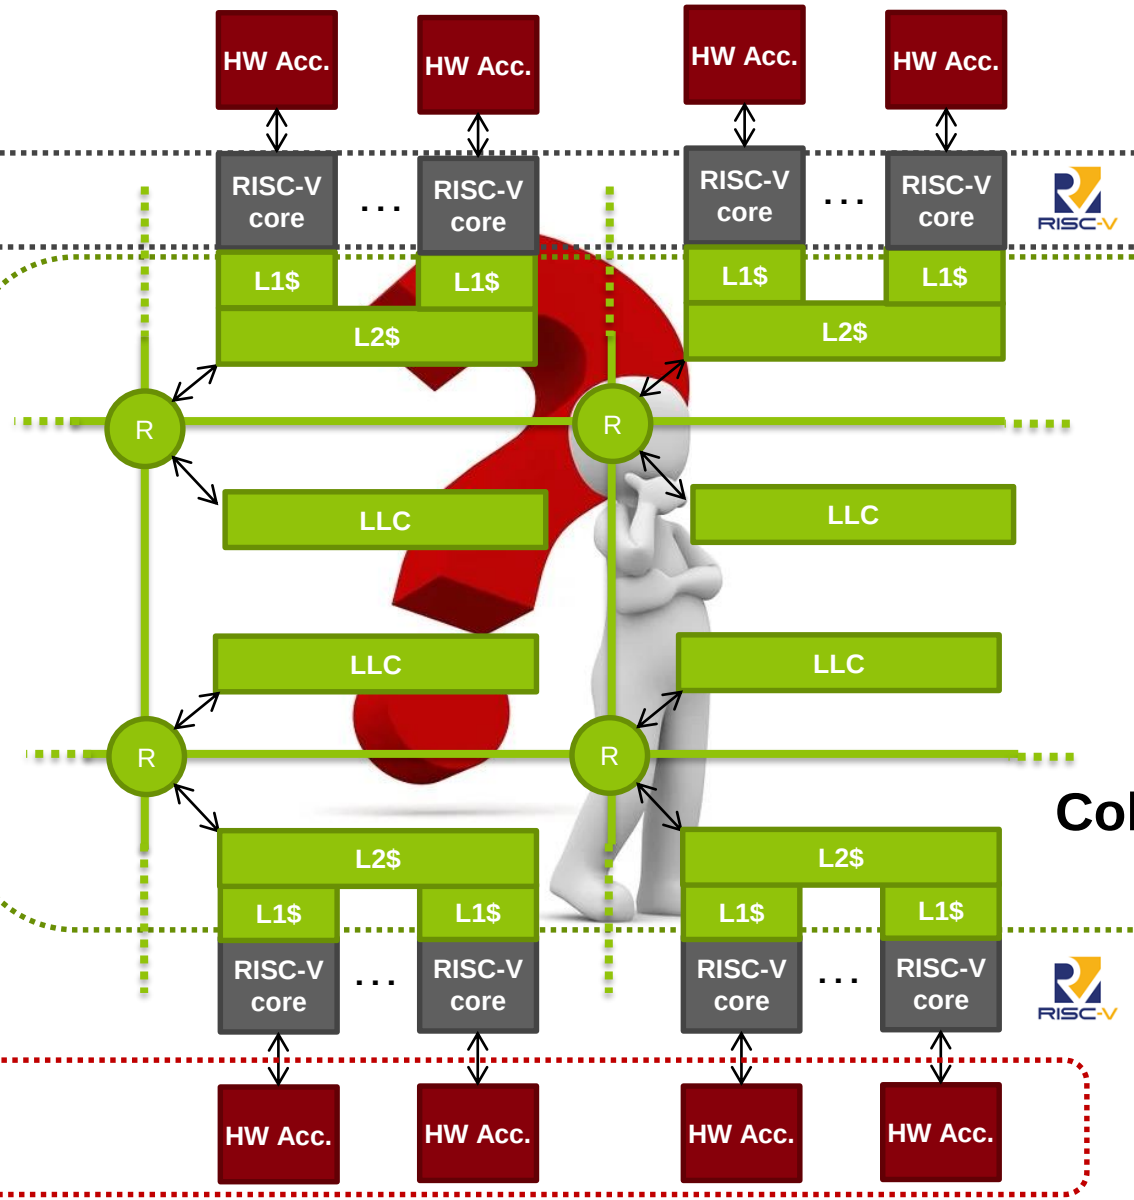

- **Specialized hardware (accelerators)**

CACHE COHERENCY : KEY FEATURES

- Up to 1024 cache-coherent cores with private L1 caches
- Physically distributed and shared L2 caches (NUMA)
- Adaptive fault tolerant L3 caches (NUCA)

- **Efficient, low-latency, hardware cache-coherency protocol**
  - Relaxed-Write-Through (RWT) protocol
  - Virtual Memory: TLB coherency ensured by RWT protocol
  - Write-back for private data, Write-through for shared data
  - Prevention of False sharing issue

- **Low-cost, scalable, coherent Network-on-Chip**
  - 4-channel NoC **with 2 channels for coherency data**
  - Directory-based coherency with linked-list directory
  - Broadcast & atomic operation support

O(N×log(N)) hardware cost

- **Adaptive L3 caches**
  - **Adaptive remapping** for application optimization
  - **Fault tolerance** support

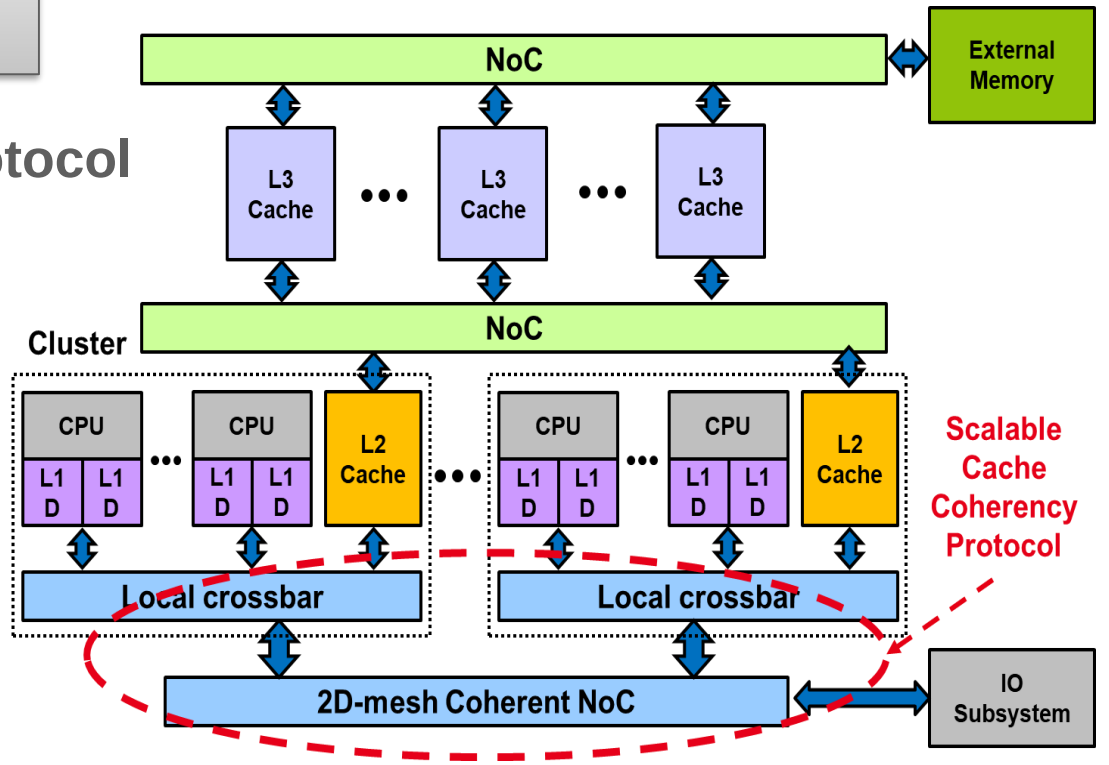

➔ *Support of similar cores or heterogeneous cores*

[E. Guthmuller, VLSISoC' 2013] [E. Guthmuller, ESSCIRC' 2018]

CACHE COHERENCY : BENEFITS & RESULTS

- **Within INTACT, implemented 96 cores in 6 chiplets**
  - For system and many-core scalability, each chiplet contains cores + distributed L1/L2/L3 caches
  - Support up to 56 chiplets (1024 cache-coherent cores)

- **Area-efficient Cache Coherency solution**
  - L1 I-caches + D-caches (16 kB / core)
  - Distributed Shared L2-caches (256 kB / cluster)
  - Adaptive & fault tolerant L3-caches (4 tiles of 1 MB)

  34 MByte L1+L2+L3 caches for 96 cores

  Coherency area impact : less than 2%

- **Energy efficiency**

  Coherency traffic : < 1% power budget

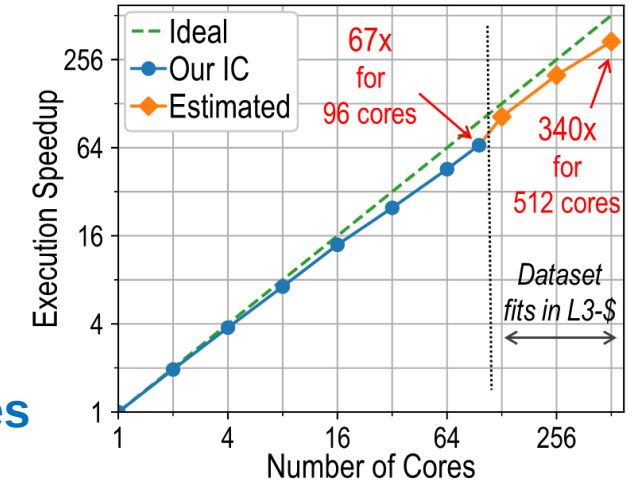

Quasi-linear speedup
with respect to the number of cores

4Mpixel filtering applications
(convolution, transposition, synchronization with barriers)

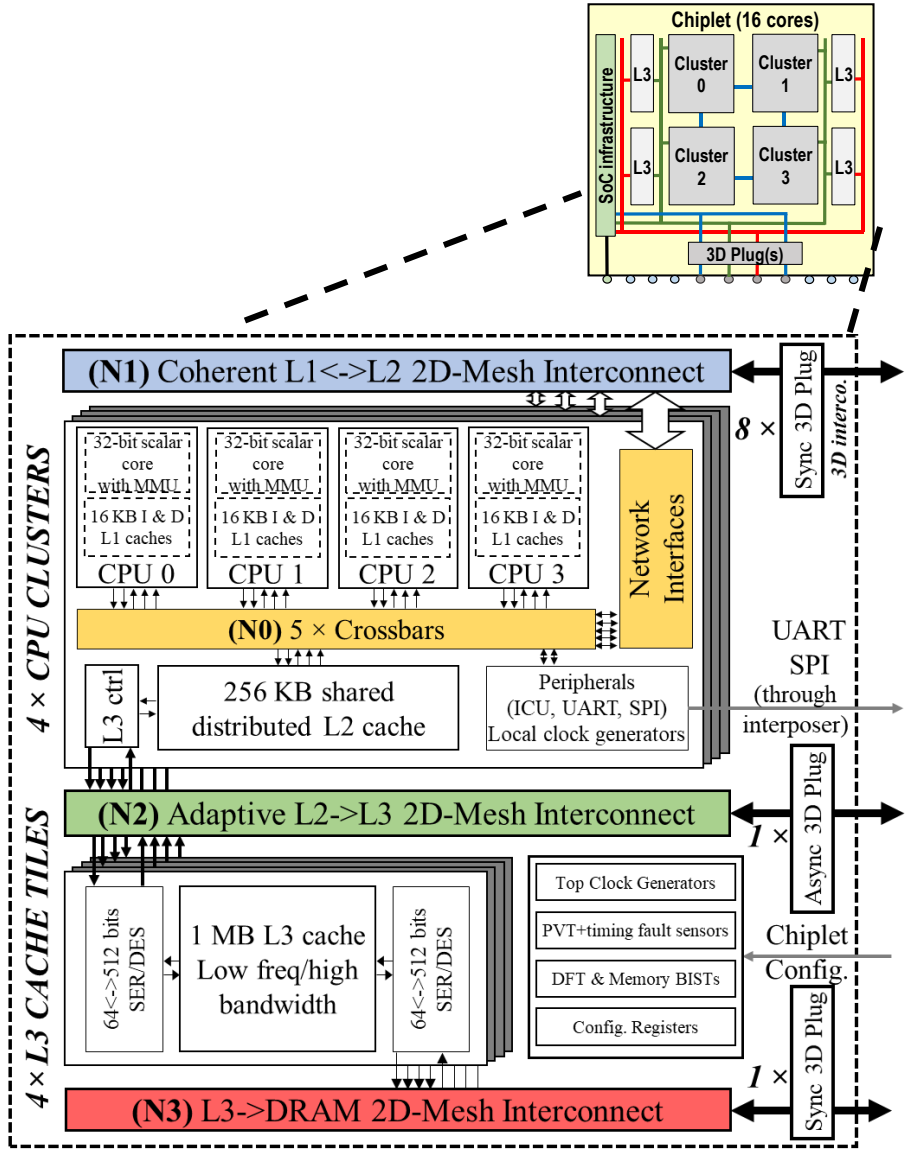

Chiplet Interconnects extended from/to active interposer

HOW DOES IT FIT IN A COMPUTING CHIP ?

- Chiplets
- Active Silicon Interposer

Challenges are multiple: 3D technology, 3D design enablement, 3D design flow, architecture partitioning, test, manufacturing…

INTACT demonstrator from CEA

INTACT : 6 CHIPLETS 3D-STACKED ON AN ACTIVE INTERPOSER

- **Chiplet Overview**
  - 4 cluster of 4 cores
  - Distributed L1$ + L2$ + L3$
  - Scalable Cache Coherency
  - Power Management

- **Active Interposer**
  - Distributed flexible interconnects *(low latency chiplet-to-chiplet traffic)*
  - Integrated Voltage Regulators *(1 per chiplet for local DVFS)*
  - Memory Controller & System IO's *(off chip communication)*
  - SOC Infrastructure, Design-for-Test

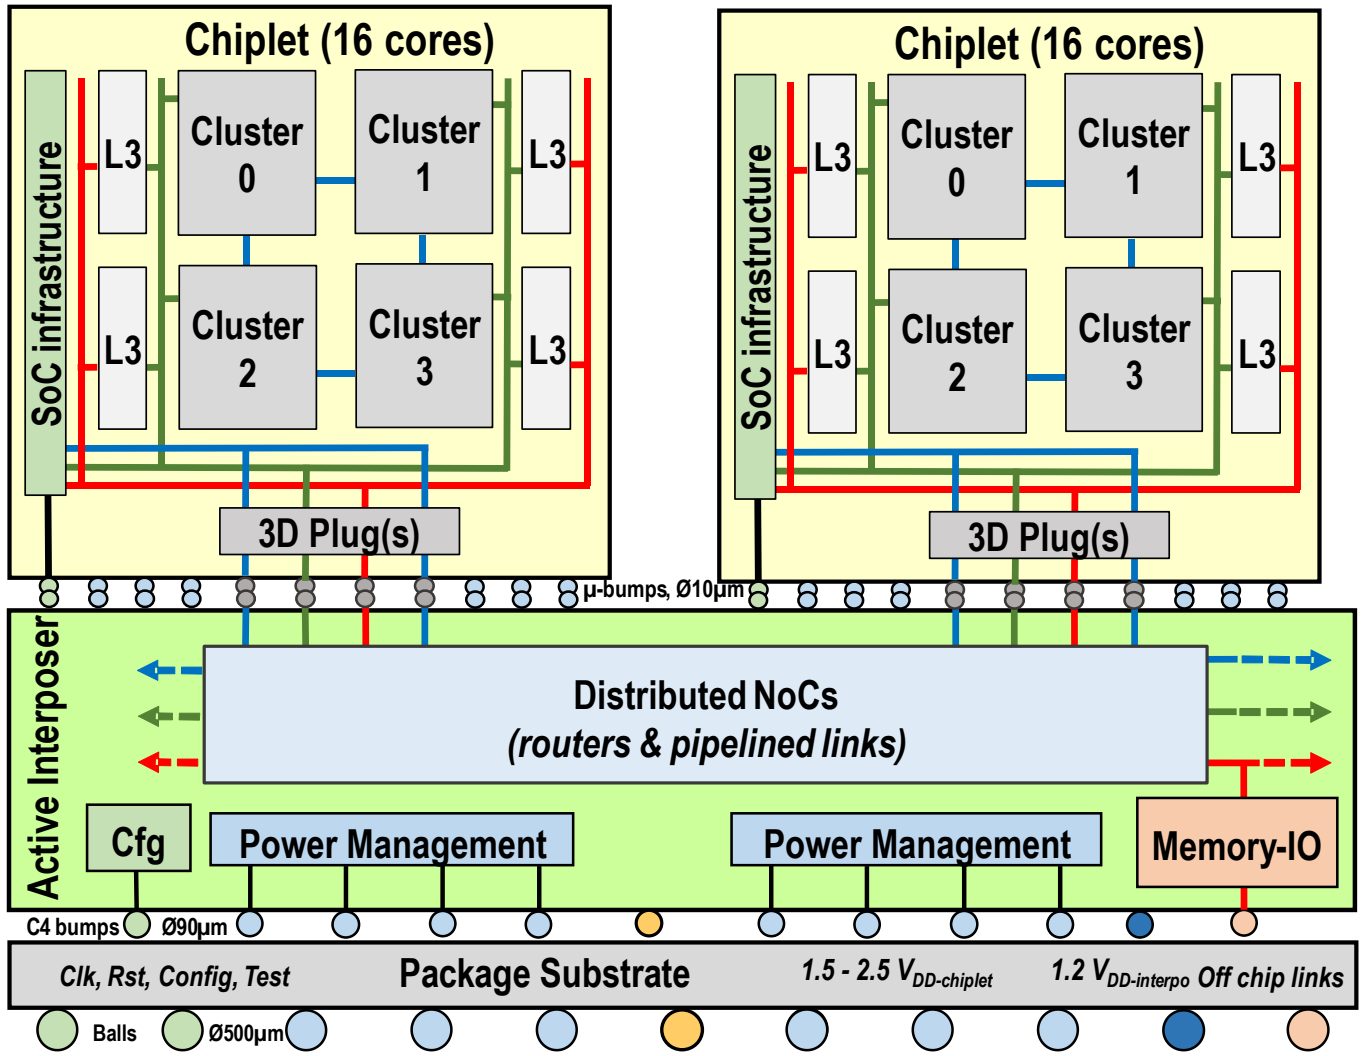

➔ *96 cores cache coherent architecture in 6 chiplets*

INTACT : TECHNOLOGY & CIRCUIT OVERVIEW

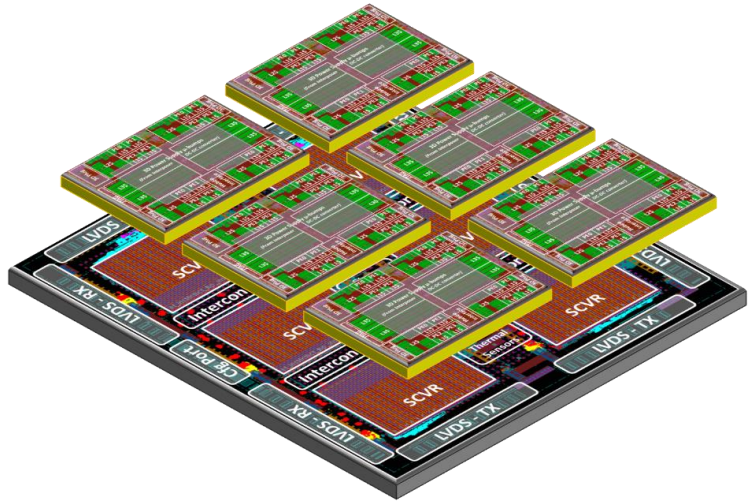

- **Die technologies**
  - *Chiplet*: FDSOI 28nm, Ultra Low Voltage, Body-Biasing, 22mm²
  - *Active Interposer*: CMOS 65nm, MIM option, 200mm²

- **3D technology integration**
  - µ-bumps, 20µm pitch
  - TSV middle, 40 µm pitch
  - Face2Face assembly on package substrate
  - 6 chiplets

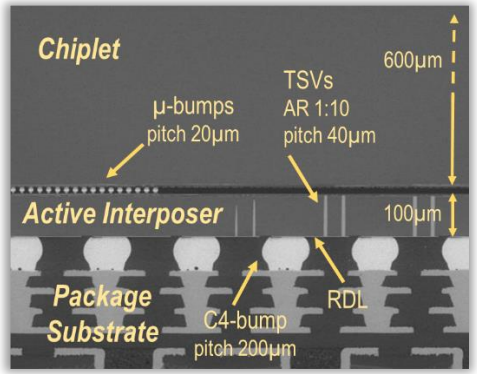

3D cross-section

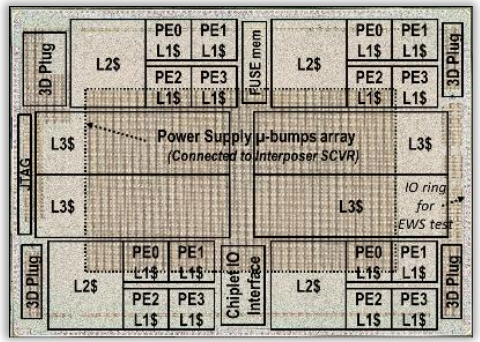

Chiplet front-face

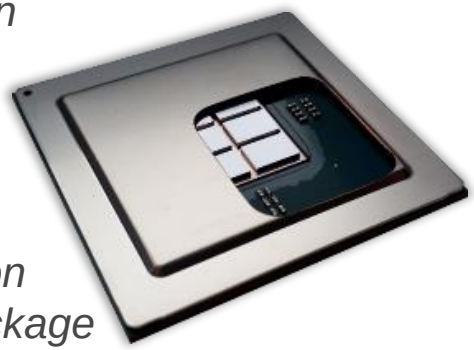

3D integration and final package

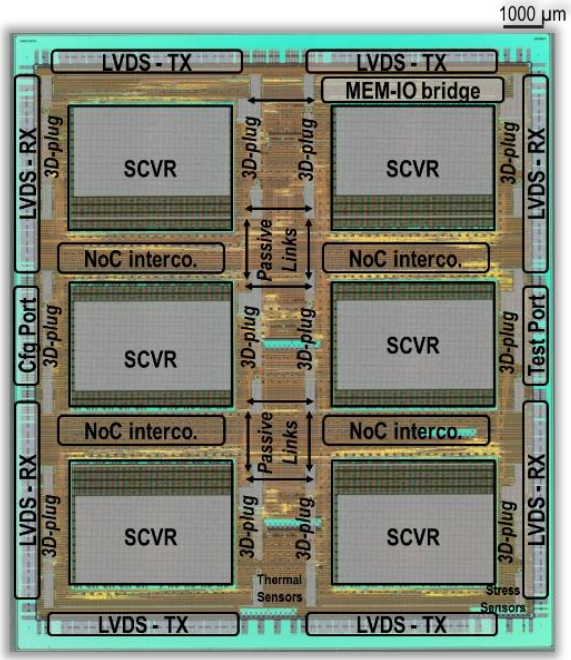

Active Interposer front-face

[P. Vivet, ISSCC'2020]

[P. Coudrain, ECTC'2020]

CONCLUSION

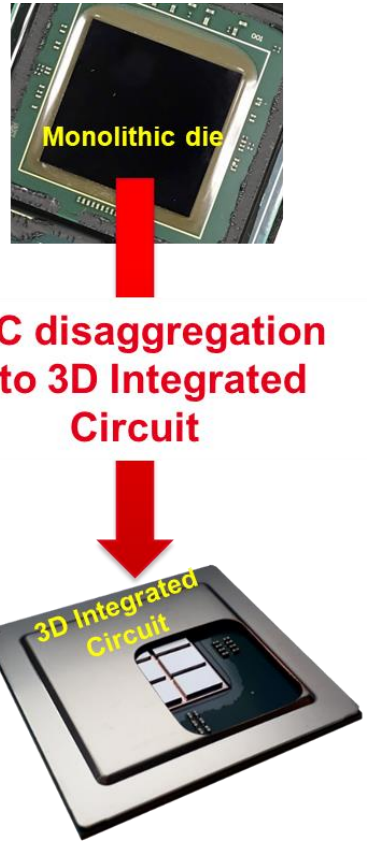

- **How to move from bare metal silicon to power efficient and cost efficient processors ?**
  - Many-core heterogeneous architectures
  - SoC disaggregation into chiplets
  - ➢ RISC-V is a unique opportunity to leverage heterogeneous many-core architectures

- **Challenges are:**
  - Scalable and coherent cache
  - 3D architecture and design

- **Key assets from CEA, thanks to demonstrator development and characterization:**
  - Scalable cache coherency architecture
  - INTACT: Six chiplets 3D-stacked on an active interposer

[P. Vivet, JSSC'2021]

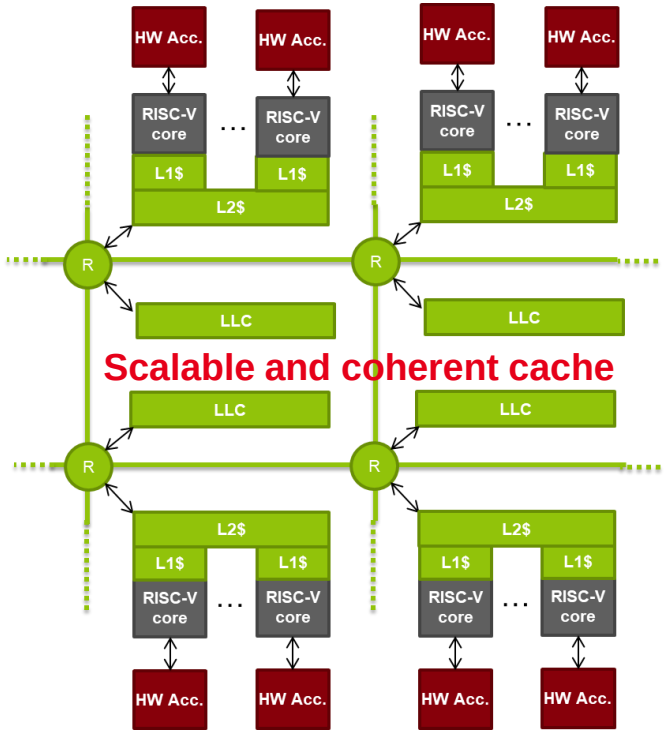

ACKNOWLEDGMENTS & REFERENCES

- **Acknowledgements**
  - IRT-3D project

- **References**

- P. Vivet et al., “A 220GOPS 96-Core Processor with 6 Chiplets 3D-Stacked on an Active Interposer Offering 0.6ns/mm Latency, 3Tb/s/mm2 Inter-Chiplet Interconnects and 156mW/mm2@ 82%-Peak-Efficiency DC-DC Converters”, ISSCC’2020.
- E. Guthmuller et al., “A 29 Gops/Watt 3D-Ready 16-Core Computing Fabric with Scalable Cache Coherent Architecture Using Distributed L2 and Adaptive L3 Caches”, ESSCIRC’2018.
- E. Guthmuller et al., "Architectural exploration of a fine-grained 3D cache for high performance in a manycore context," VLSI-SoC’2013.
- Y. Thonnart, et al., “A Fully Asynchronous Low-Power Framework for GALS NoC Integration”, DATE’2010.
- Y. Thonnart et al., “Latency Improvement of an Industrial SoC System Interconnect using an Asynchronous NoC Backbone”, ASYNC’2019.
- I. Miro-Panades et al., "In-situ Fmax/Vmin tracking for energy efficiency and reliability optimization," IOLT’2017.

Commissariat à l'énergie atomique et aux énergies alternatives
Institut List | CEA SACLAY NANO-INNOV | BAT. 861 – PC142
91191 Gif-sur-Yvette Cedex - FRANCE
www-list.cea.fr

Établissement public à caractère industriel et commercial | RCS Paris B 775 685 019