SIPEARL

http://www.teratec.eu/forum/atelier_7_06.html

Indicateurs de performance clés pour mesurer la consommation énergétique des microprocesseurs généraliste dans les supercalculateurs Exascale

Rhea, qui implémente le cœur ARM Neoverse V1 conçu pour le HPC, repose sur une architecture basse consommation. Il s'appuie également sur des choix architecturaux innovants destinés à maximiser sa performance énergétique. Dans cette présentation, nous détaillerons quels indicateurs de consommation sont pertinents à l'âge de l'exascale et en quoi Rhea innove pour la gestion de l'énergie consommée et l'optimisation de la performance par watt

At the heart of Rhea

With its high-performance, low-power Arm Neoverse V1 architecture, Rhea will meet the needs of all supercomputing workloads.

Key features

Core	- Arm architecture - Neoverse V1 cores - SVE 256 per core supporting 64/32/BF16 and Int8 - ArmVirtualization extensions
SoC	- Arm mesh fabric - Advanced RAS support including Arm RAS extensions - Link protection for NoC & high-speed IO - ECC support for selected memory
Cache	- Large L3 (Shared Level Cache) - RAS supported for all cache levels
Memory	- HBM2e - And DDR5 - ECC for memory and link protection for controllers
High Speed I/O	- PCIe, CCIX & CXL - Root and endpoint support
Other I/O	- USB, GPIO, SPI, I^2C
Power Management	- Power management block to optimize perf/watt accross use cases and workloads.
Security Block Support	- Secure boot and secure upgrade - Crypto - True random number generation - Made in Europe

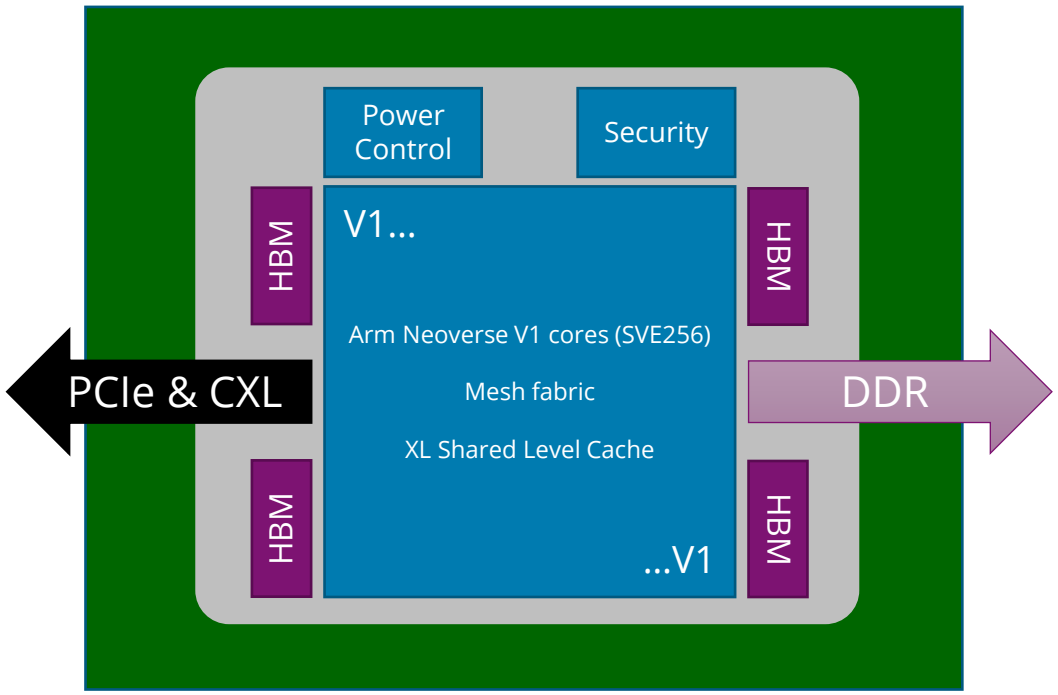

Rhea will deliver extraordinary real compute performance and efficiency with an unmatched Bytes/Flops ratio.

Computation is (almost) free, data movement is not

Memory is diverse

- Caches L1,L2,L3
- DRAM
- GDRAM
- NUMA domains
- HBM/MCDRAM
- NVDIMM
- Node-local SSD
- …
- Object Storage
- GFS

Most used HPC metrics like R_Peak, R_Max or Green500 don't address this problem.

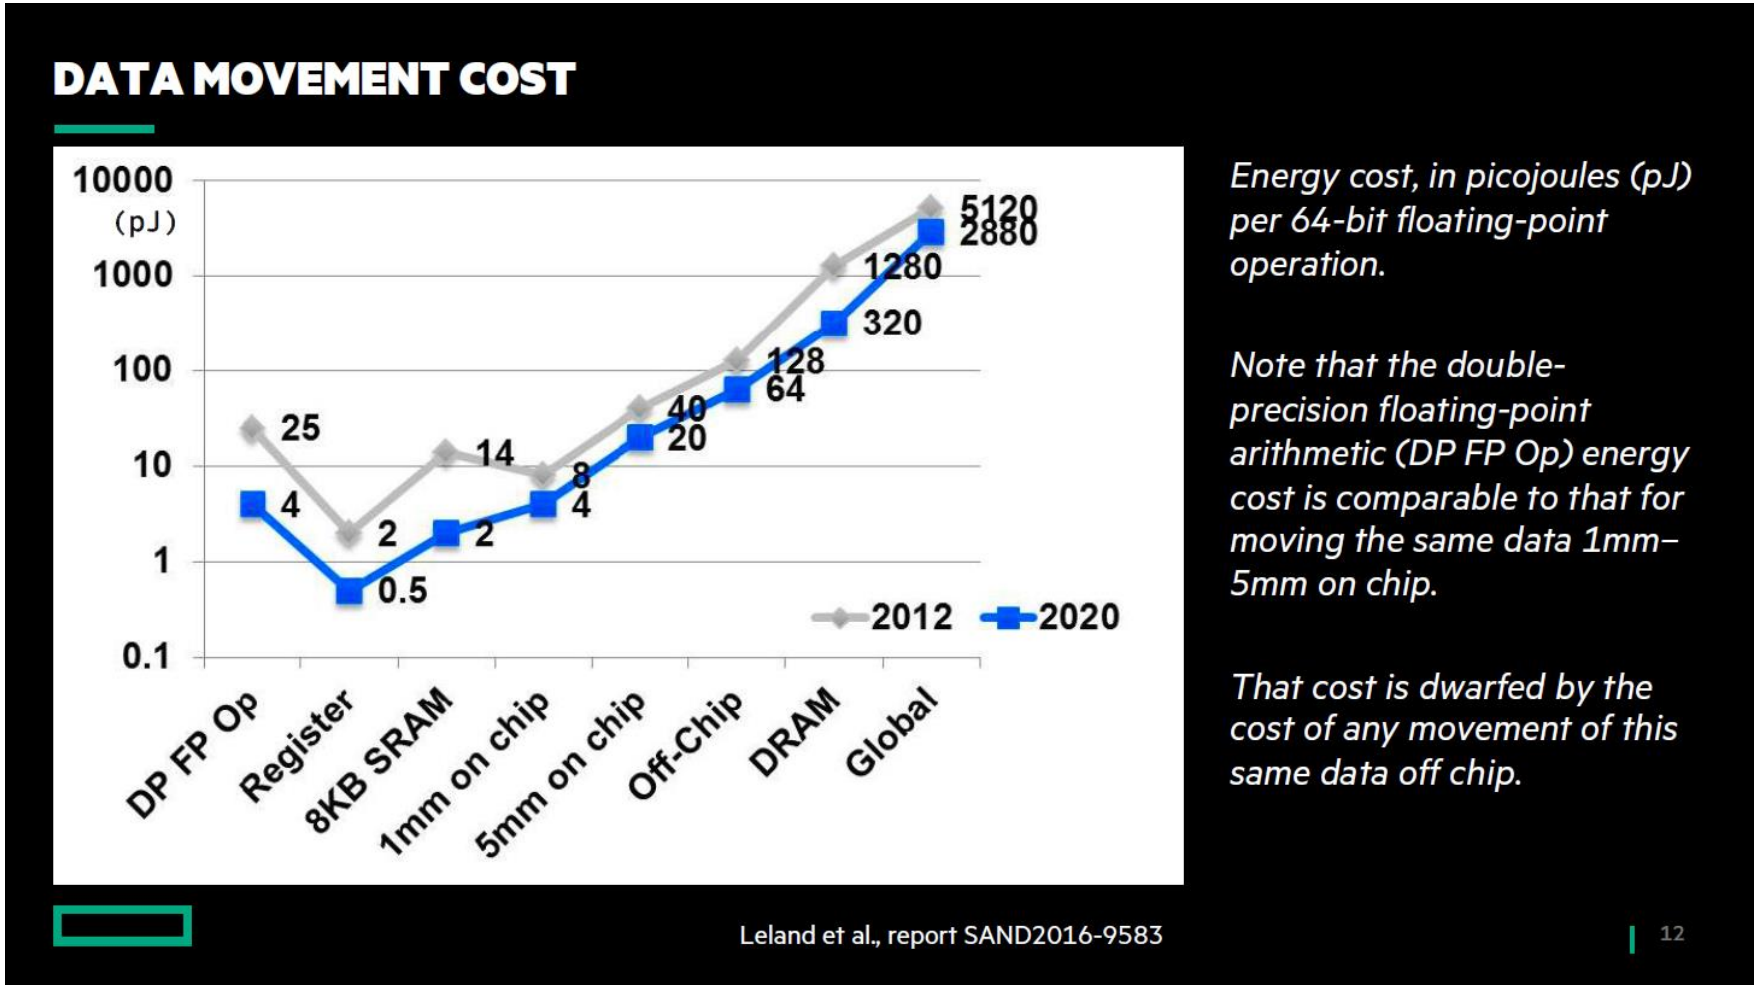

Source: Utz-Uwe Haus, Head of HPE HPC/AI EMEA Research Lab
2021-03-03, TRR154/MINOA conference "Trends in Modelling, Simulation and Optimisation: Theory and Applications"

Performances: R_Peak, Bytes/FLOPS metrics

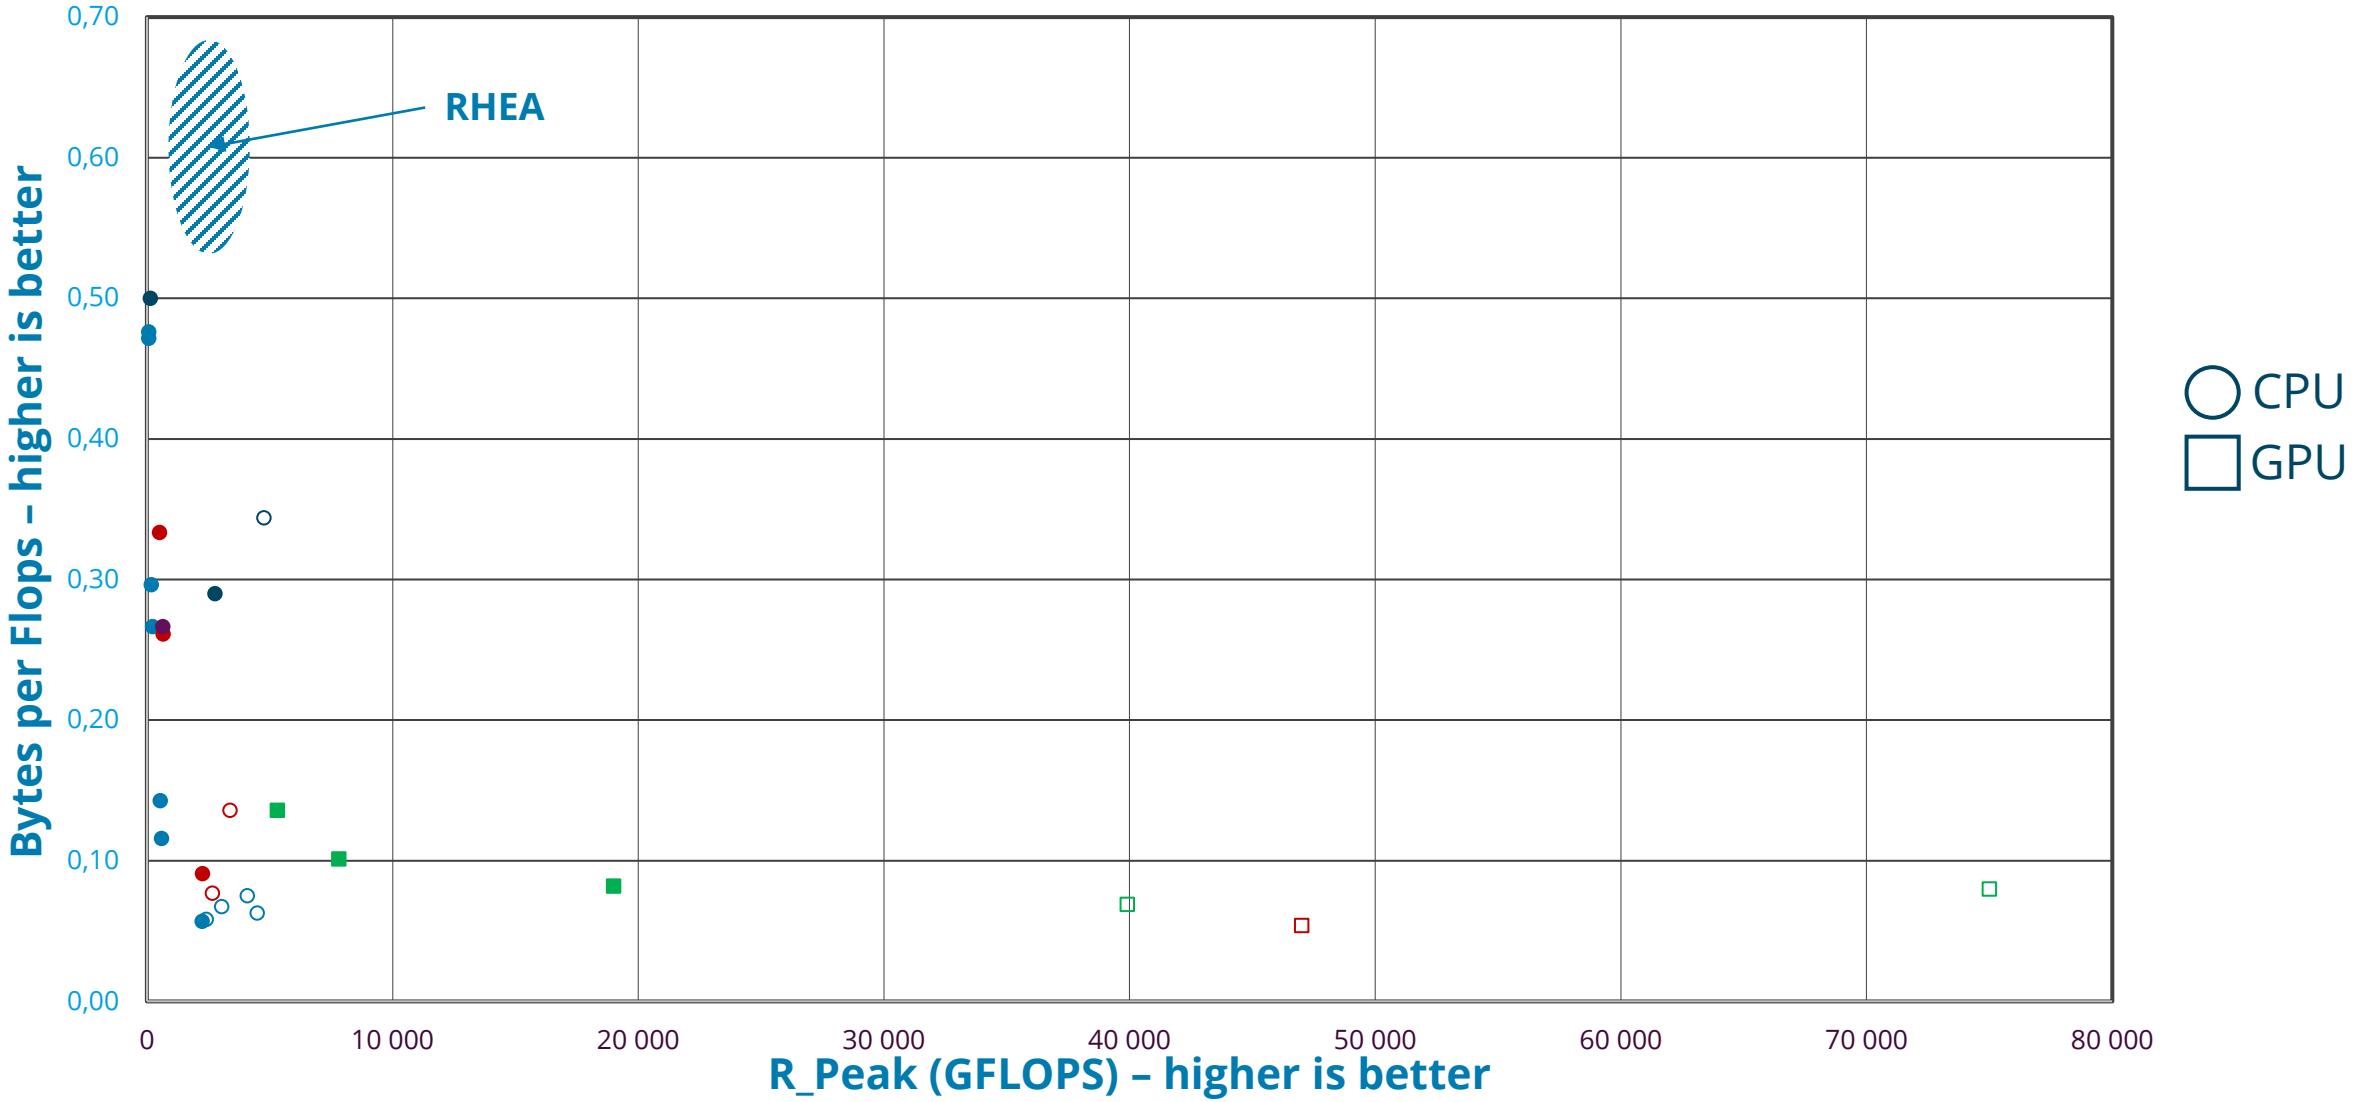

Performances: R_Peak, Bytes/FLOPS metrics

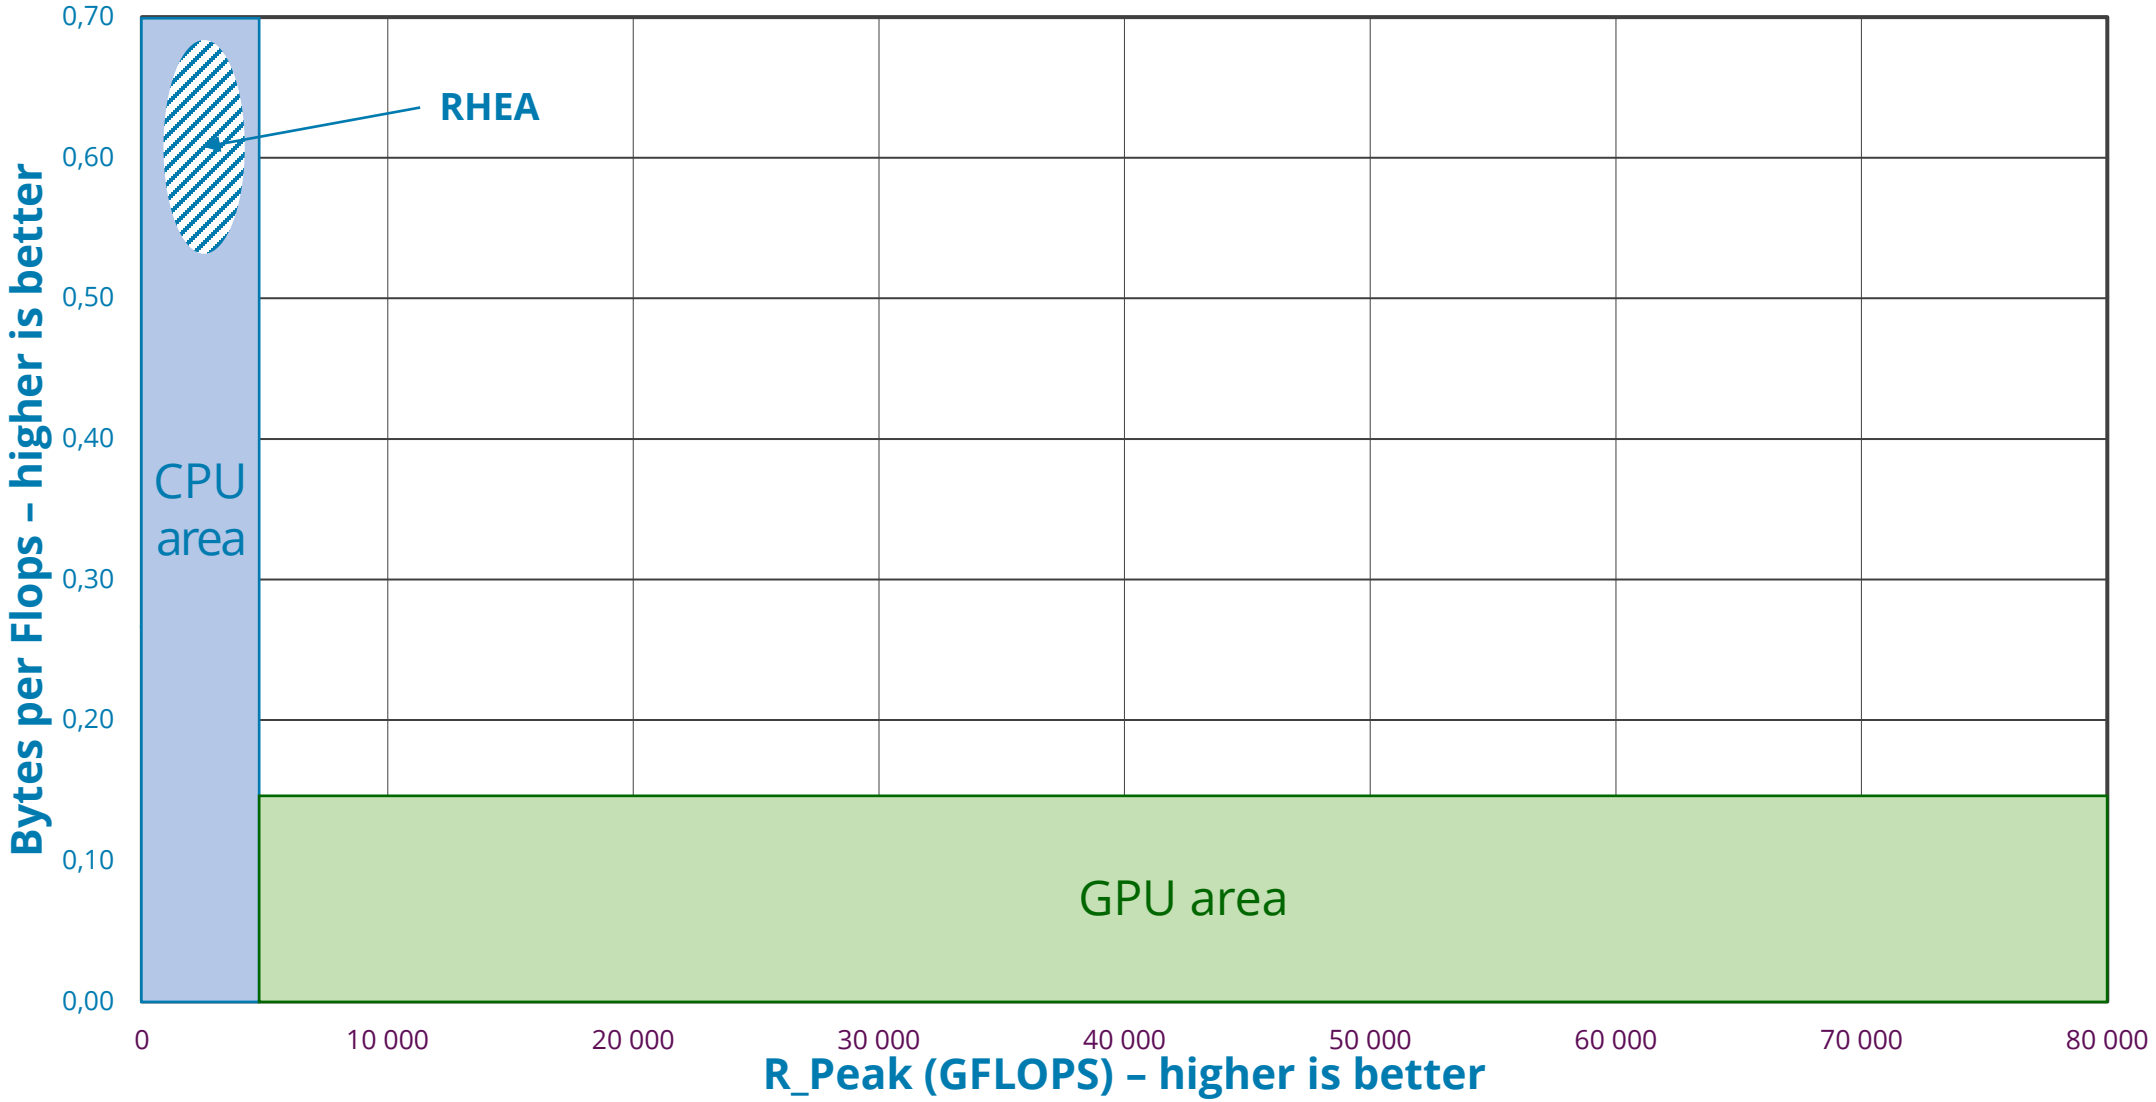

BYTES/FLOPS – HPCG/R_Peak and Power consumption metrics

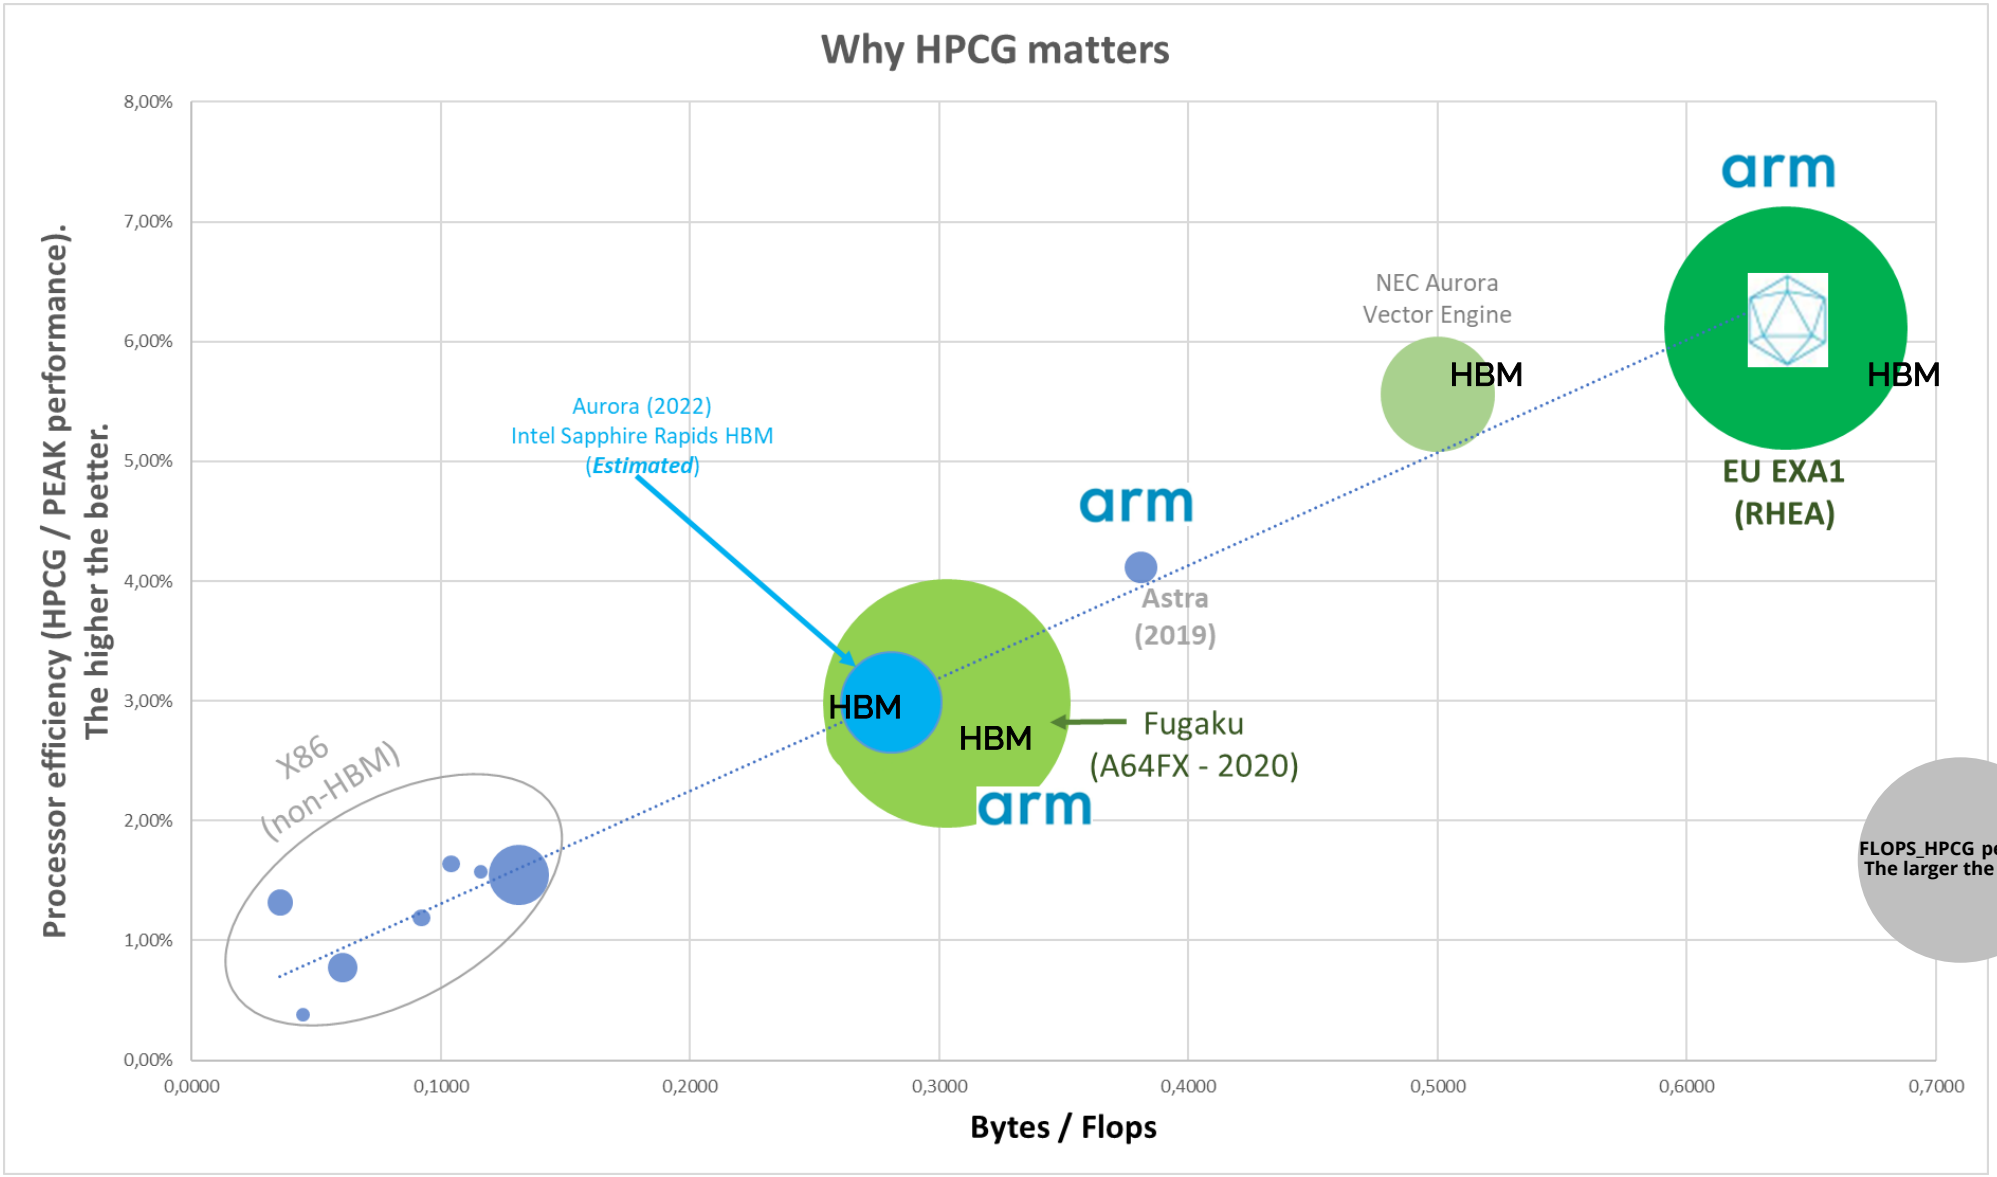

Bytes/FLOPS, Power consumption per FLOPS_HPCG

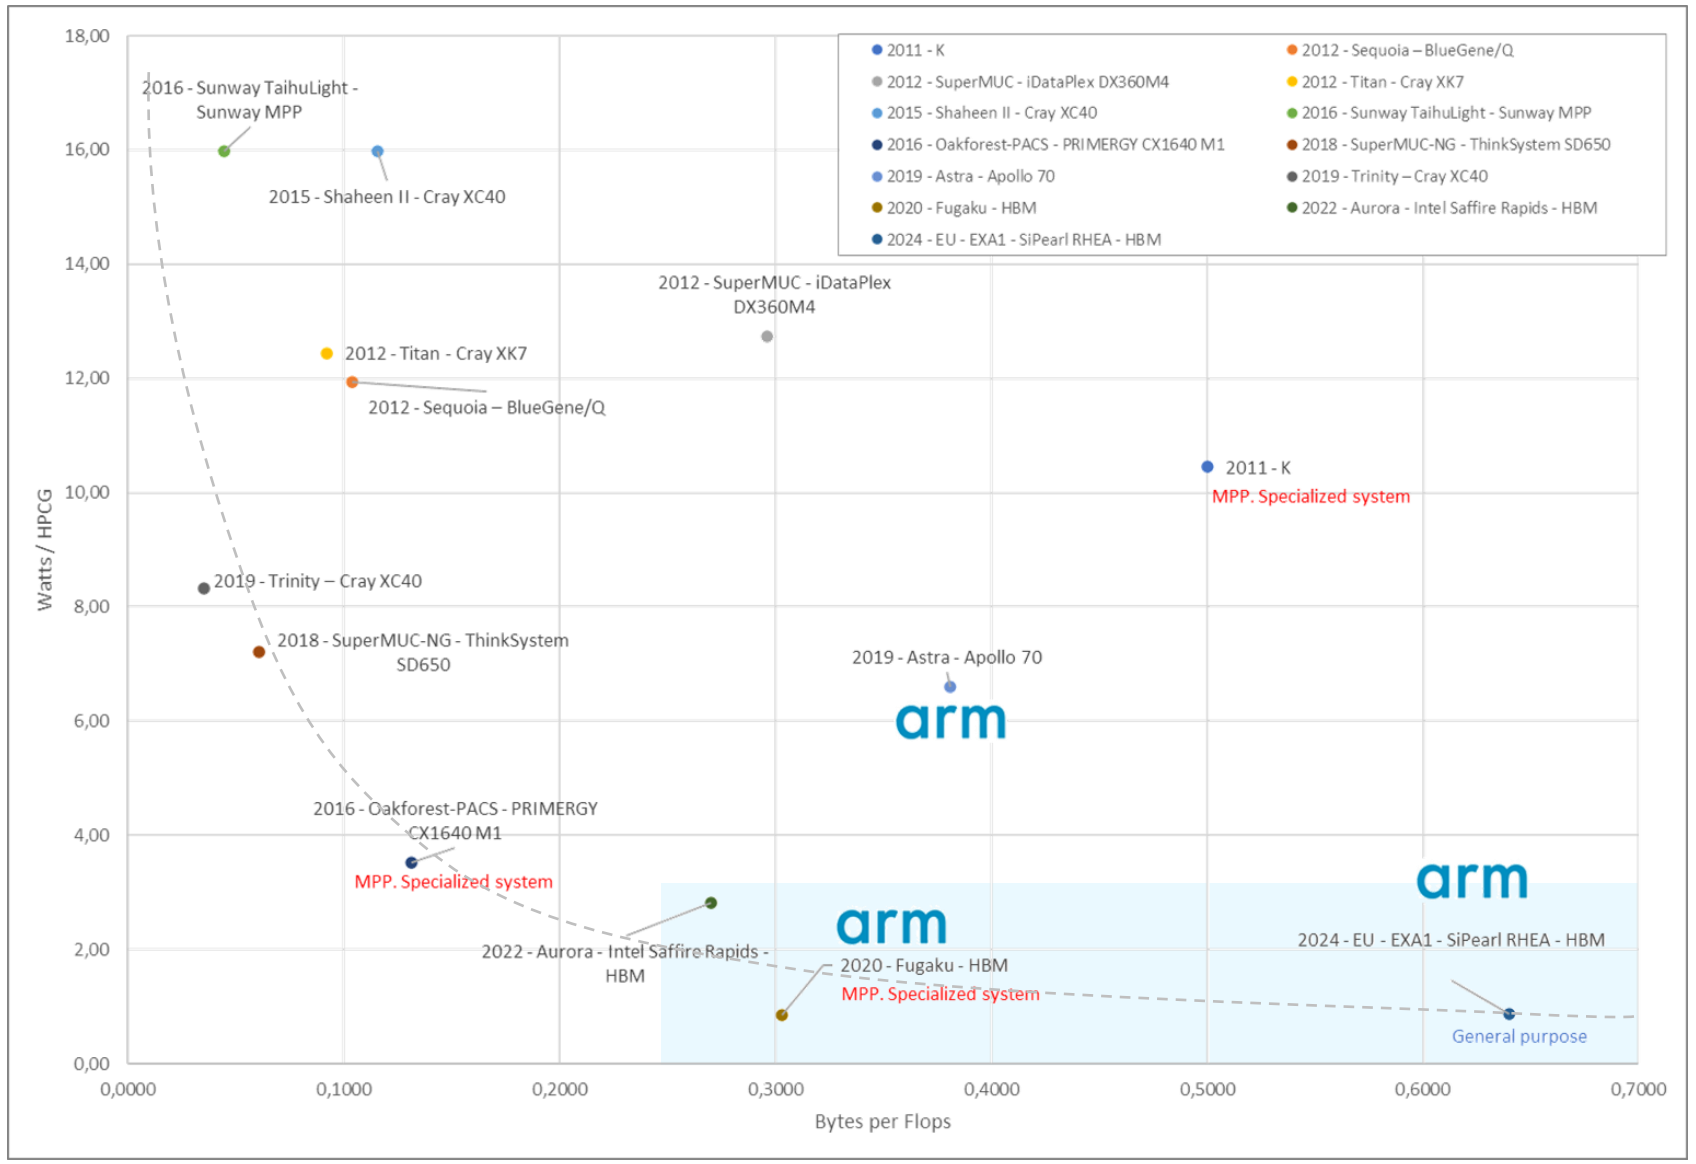

- The lower the better.
- Measures the power consumption on “real” applications.

- ARM is better
- HPCG is a must.
  - Warning: Higher CAPEX
  - **much** better OPEC
  - Optimal TCO

Mem bandwidth versus Peak DP performances – CPU & GPU

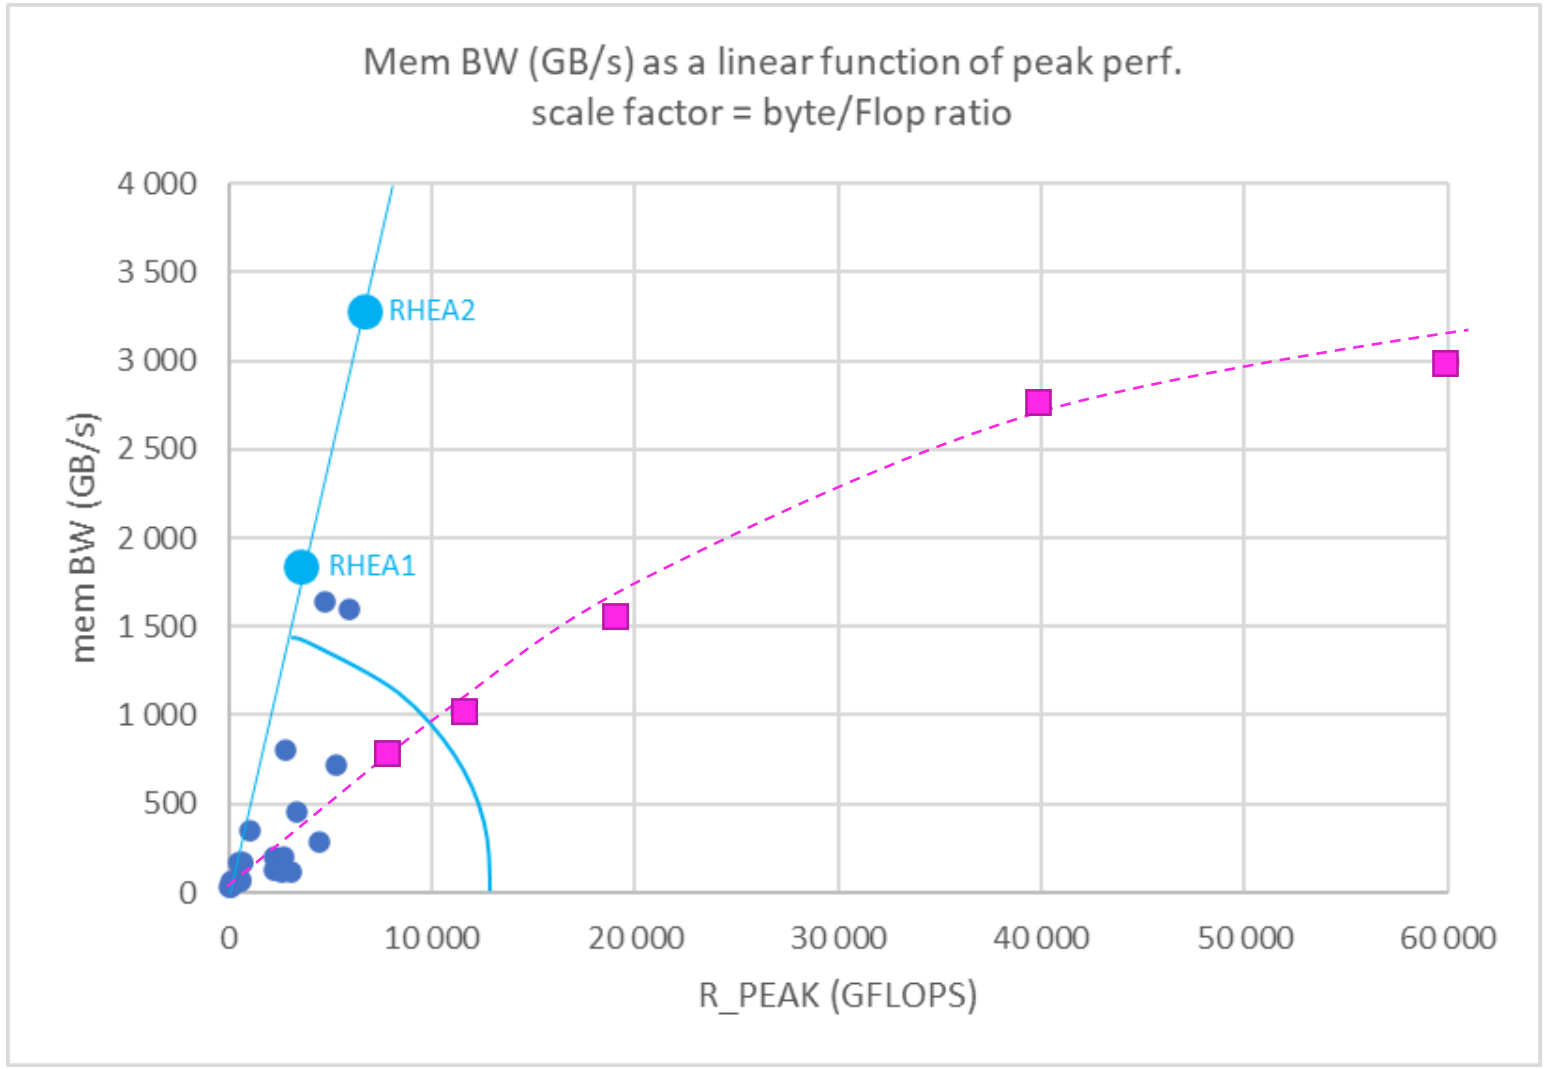

Memory hierarchy (ies) & metrics

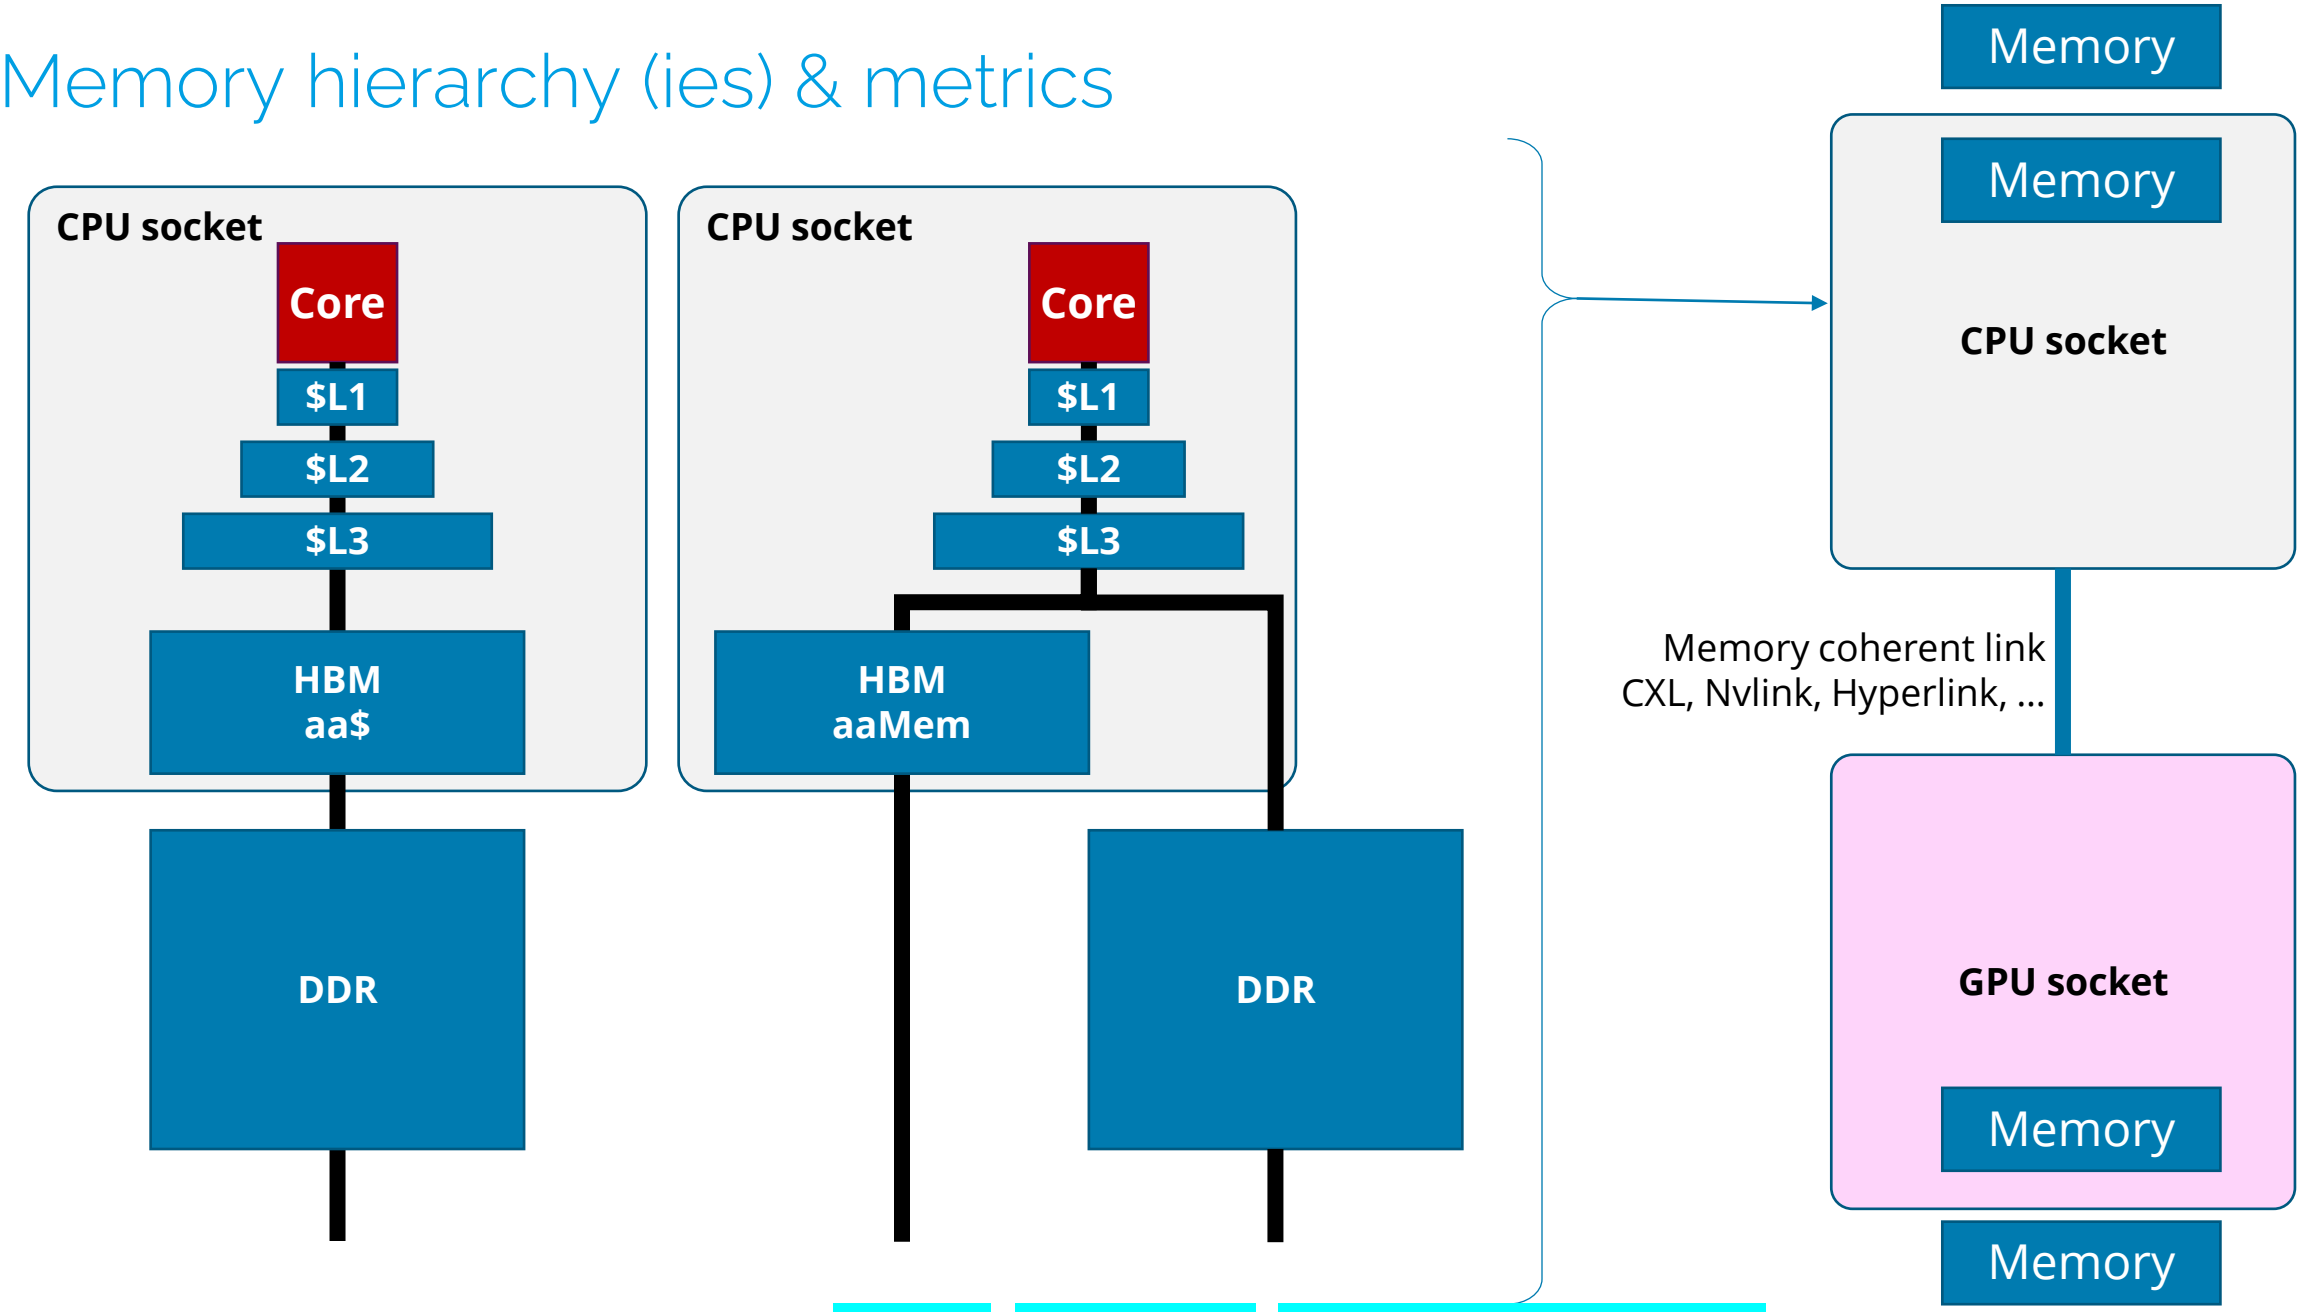

Fundamental metrics: latency, bandwidth, pJ per transfered Byte. Any other?

Memory hierarchy (ies)

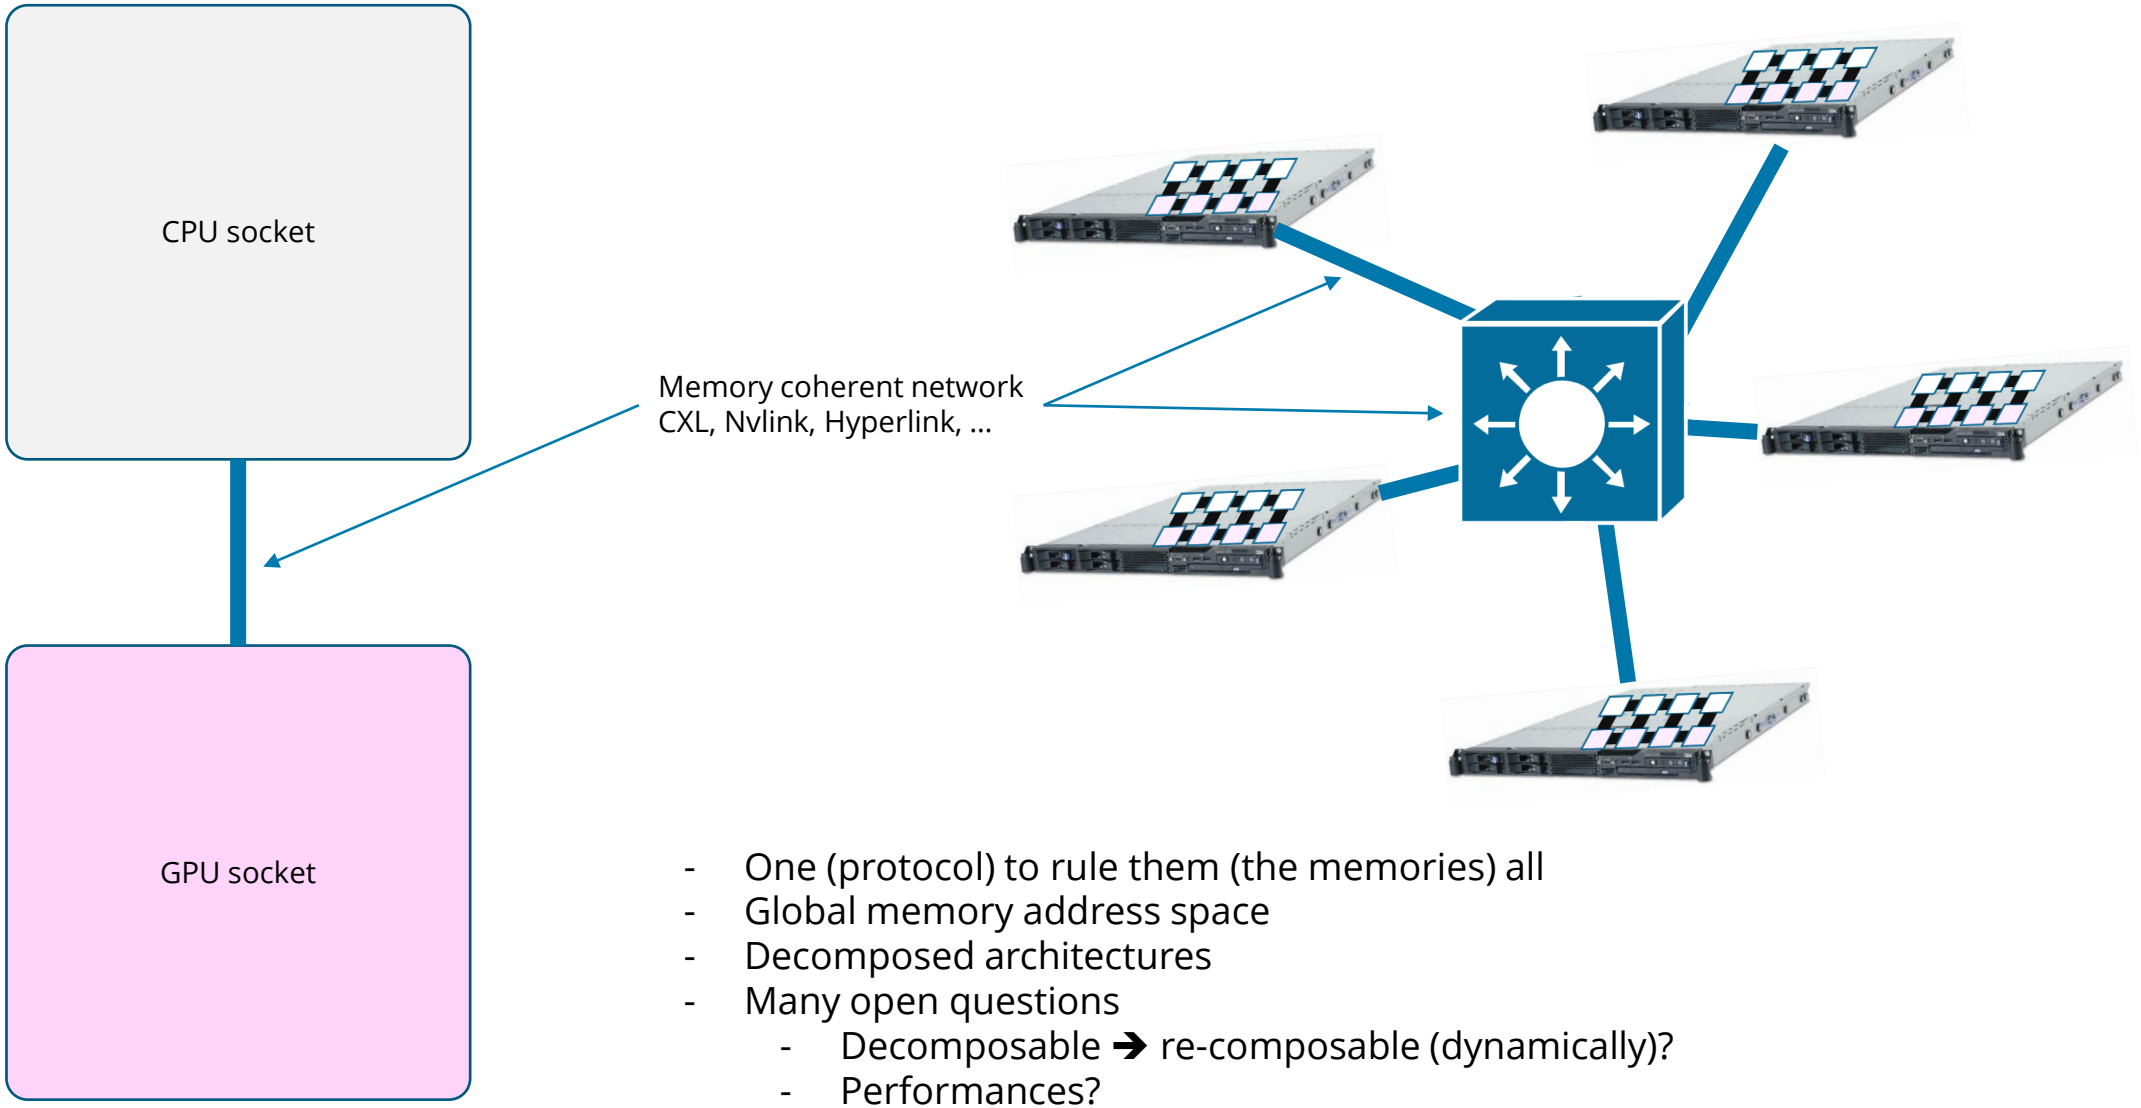

- One (protocol) to rule them (the memories) all
- Global memory address space
- Decomposed architectures
- Many open questions
  - Decomposable ➔ re-composable (dynamically)?
  - Performances?
  - Δ+ energy cost?

Metrics for the CPU

Metric	Description	Remark
FLOPS	- FP64 (of course) - Peak - Linpack - HPCG	- Linpack is related to Peak (% of peak) - HPCG is related to mem BW and Bytes per Flops ratio
Memory BW	From L1/L2/L3 to DDR (and/or HBM)	- No impact on peak or Linpack - Fundamental parameter for HPCG
Bytes per Flops	B/F: how many bytes per FP operation per cycle?	The HW perspective. The higher the better. Was 13-15 in Cray C90 (1991) & T90 (1995). Now, from <0.01 to 0.3 The most simple metric we can imagine to measure data movement. In the chip.
Flops per Bytes	F/B: computational intensity.	Two perspectives - The HW perspective. The lower the better. Inverse of F/B. - The SW perspective. The higher the better. Apps have less memory dependency.
Latency	Depending on what you measure, it could be in cycle, ns, µs, ms	Interesting to know. Relevant for better cache utilization.
Hit rate	hit divided by total accesses	- Between cores, caches, from cache to memories. - Properly defined. Very much related to latency. - The higher the better (0%: never in cache / 100% always in cache). - This is a very important metric for microprocessor designers. - Is it used by application developers and end-users?

Energy metrics

Metric	Description	Remark
Watts 	TDP, W_Linpack (green500), W_HPCG, W_Idle Linpack per watt (Green500) HPCG per watt	- **The definition is unclear** about what should be measure (perimeter). - Green500 does not measure energy per data movement - W per HPCG and HPCG per watt are missing metrics.
Energy (J or watts per hours) 	How watts per hour (per day, per month, per year, per job …)	- This is the metric the people should always consider (because of OPEX). - **Energy to solution would be more relevant.**

La société française EAS(*) (Energy Aware Solutions) développe l'outil EAR qui permet

1) de récolter des métriques de performance et d'énergie
2) de prendre des décisions pour réduire automatiquement la consommation énergétique des applications selon des critères définis par l'utilisateur.

SiPearl est engagé avec EAR sur des projets et a l'intention d'en faire un des outils de base de son portefeuille pour répondre aux problématiques de mesure et d'optimisation énergétiques.

(*) https://www.eas4dc.com/

About… SiPearl

Created by Philippe Notton, SiPearl is designing the high-performance, low-power microprocessor for European exascale supercomputers. This new generation of microprocessors will enable Europe to set out its technological sovereignty in strategic high performance computing markets such as artificial intelligence, medical research or climate modelling.

SiPearl is working in close collaboration with its 27 partners from the European Processor Initiative (EPI) consortium - leading names from the scientific community, supercomputing centres and industry - which are its stakeholders, future clients and end-users.

SiPearl employs 106 people in France, Germany and Spain. Its first range of microprocessors, Rhea, will be launched at the end of the year.

The company is supported by the European Union (funding from the European Union's Horizon 2020 research and innovation programme under specific grant agreement no.826647).

Contact
Jean-Marc Denis
Chief Strategy Officer
jean-marc.denis@sipearl.com