SIPEARL

# SiPearl Outlook

## Teratec

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

# Company Update

SIPEARL

# SiPearl corporate overview

## The European Server Processor Solution

**HQ: Maisons-Laffitte (Paris), France**

**Incorporated in June 2019**

**CEO and Founder, Philippe Notton**

**Design centers:**

- France: Maisons-Laffitte, Massy Palaiseau, Sophia Antipolis, Grenoble
- Germany: Duisburg (Düsseldorf )
- Spain: Barcelona

**Key Personnel from Intel, Atos, ST, Marvell, Nokia, Mstar-Mediatek**

**HPC Targeted Architecture based on Arm Neoverse V1 cores**

**+100 employees today, targeting >1,000 in 2025**

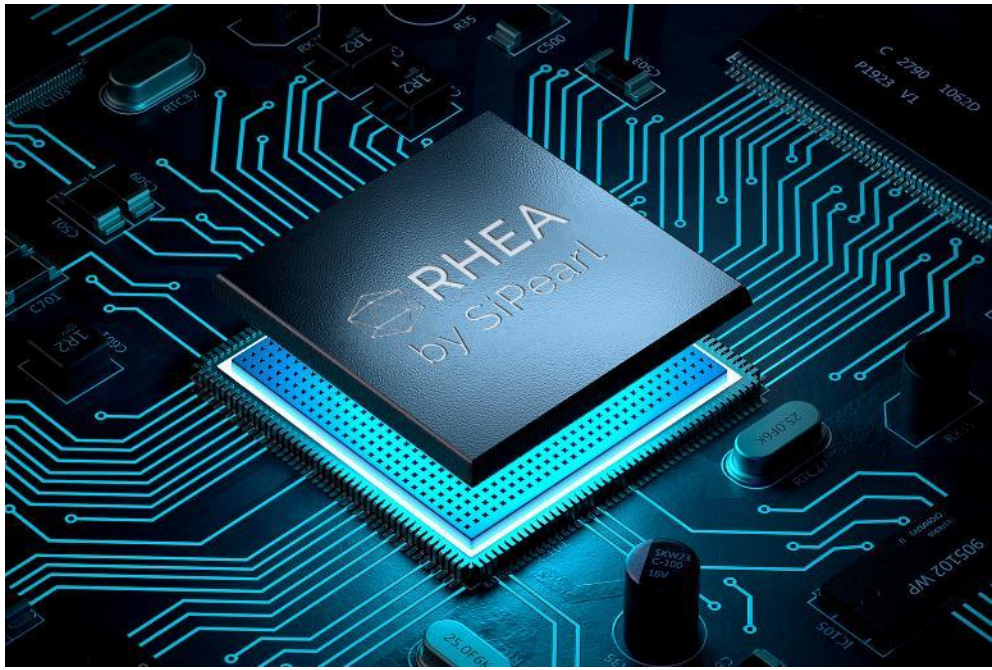


# SiPearl offices

## We are close to our partners and customers




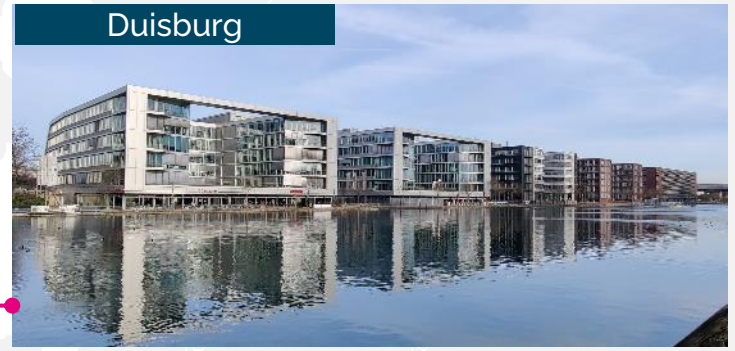






# Sipearl extensions
## 2022 and 2023

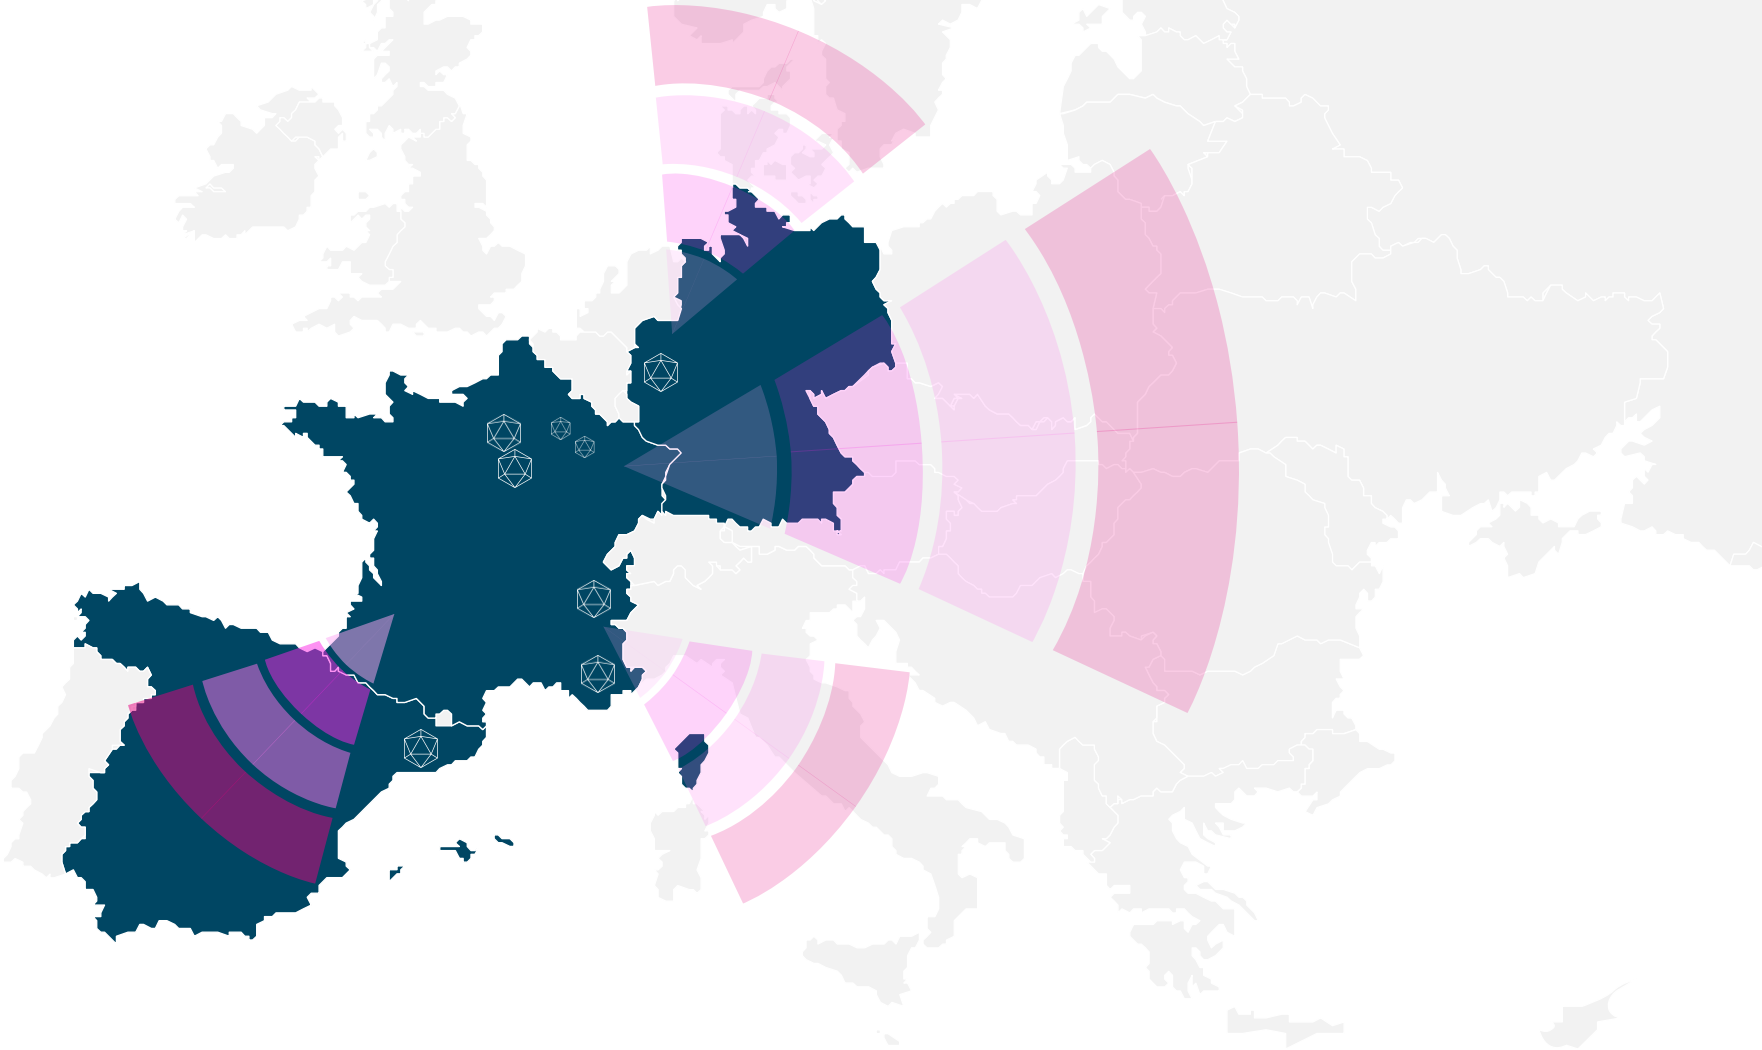

# SIPEARL CORPORATE VISION AND STRATEGY

## HPC Supercomputing
SiPearl entry business: European HPC

2023

## Data Center-Central
Data centers, private and public cloud

2025

## Data Center-Edge
Smaller Data centers, incl around 5G infrastructure

2027+

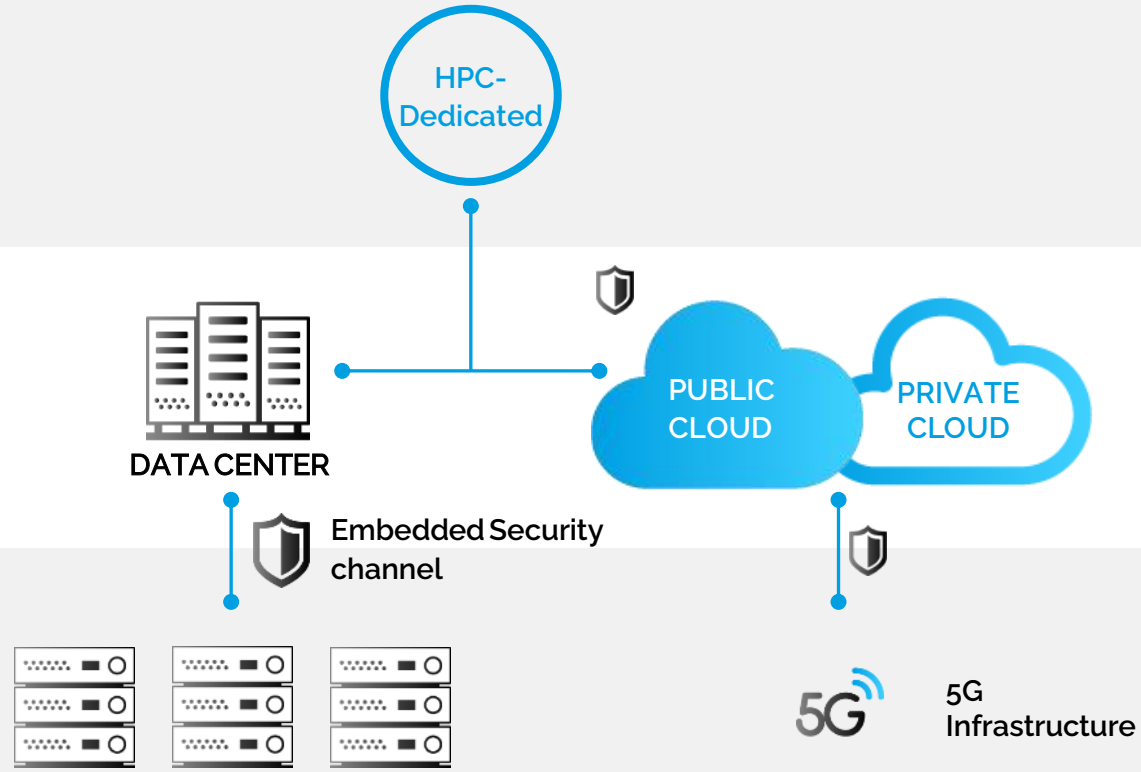


Our business model is sustainable over time

# OVERALL ROADMAP

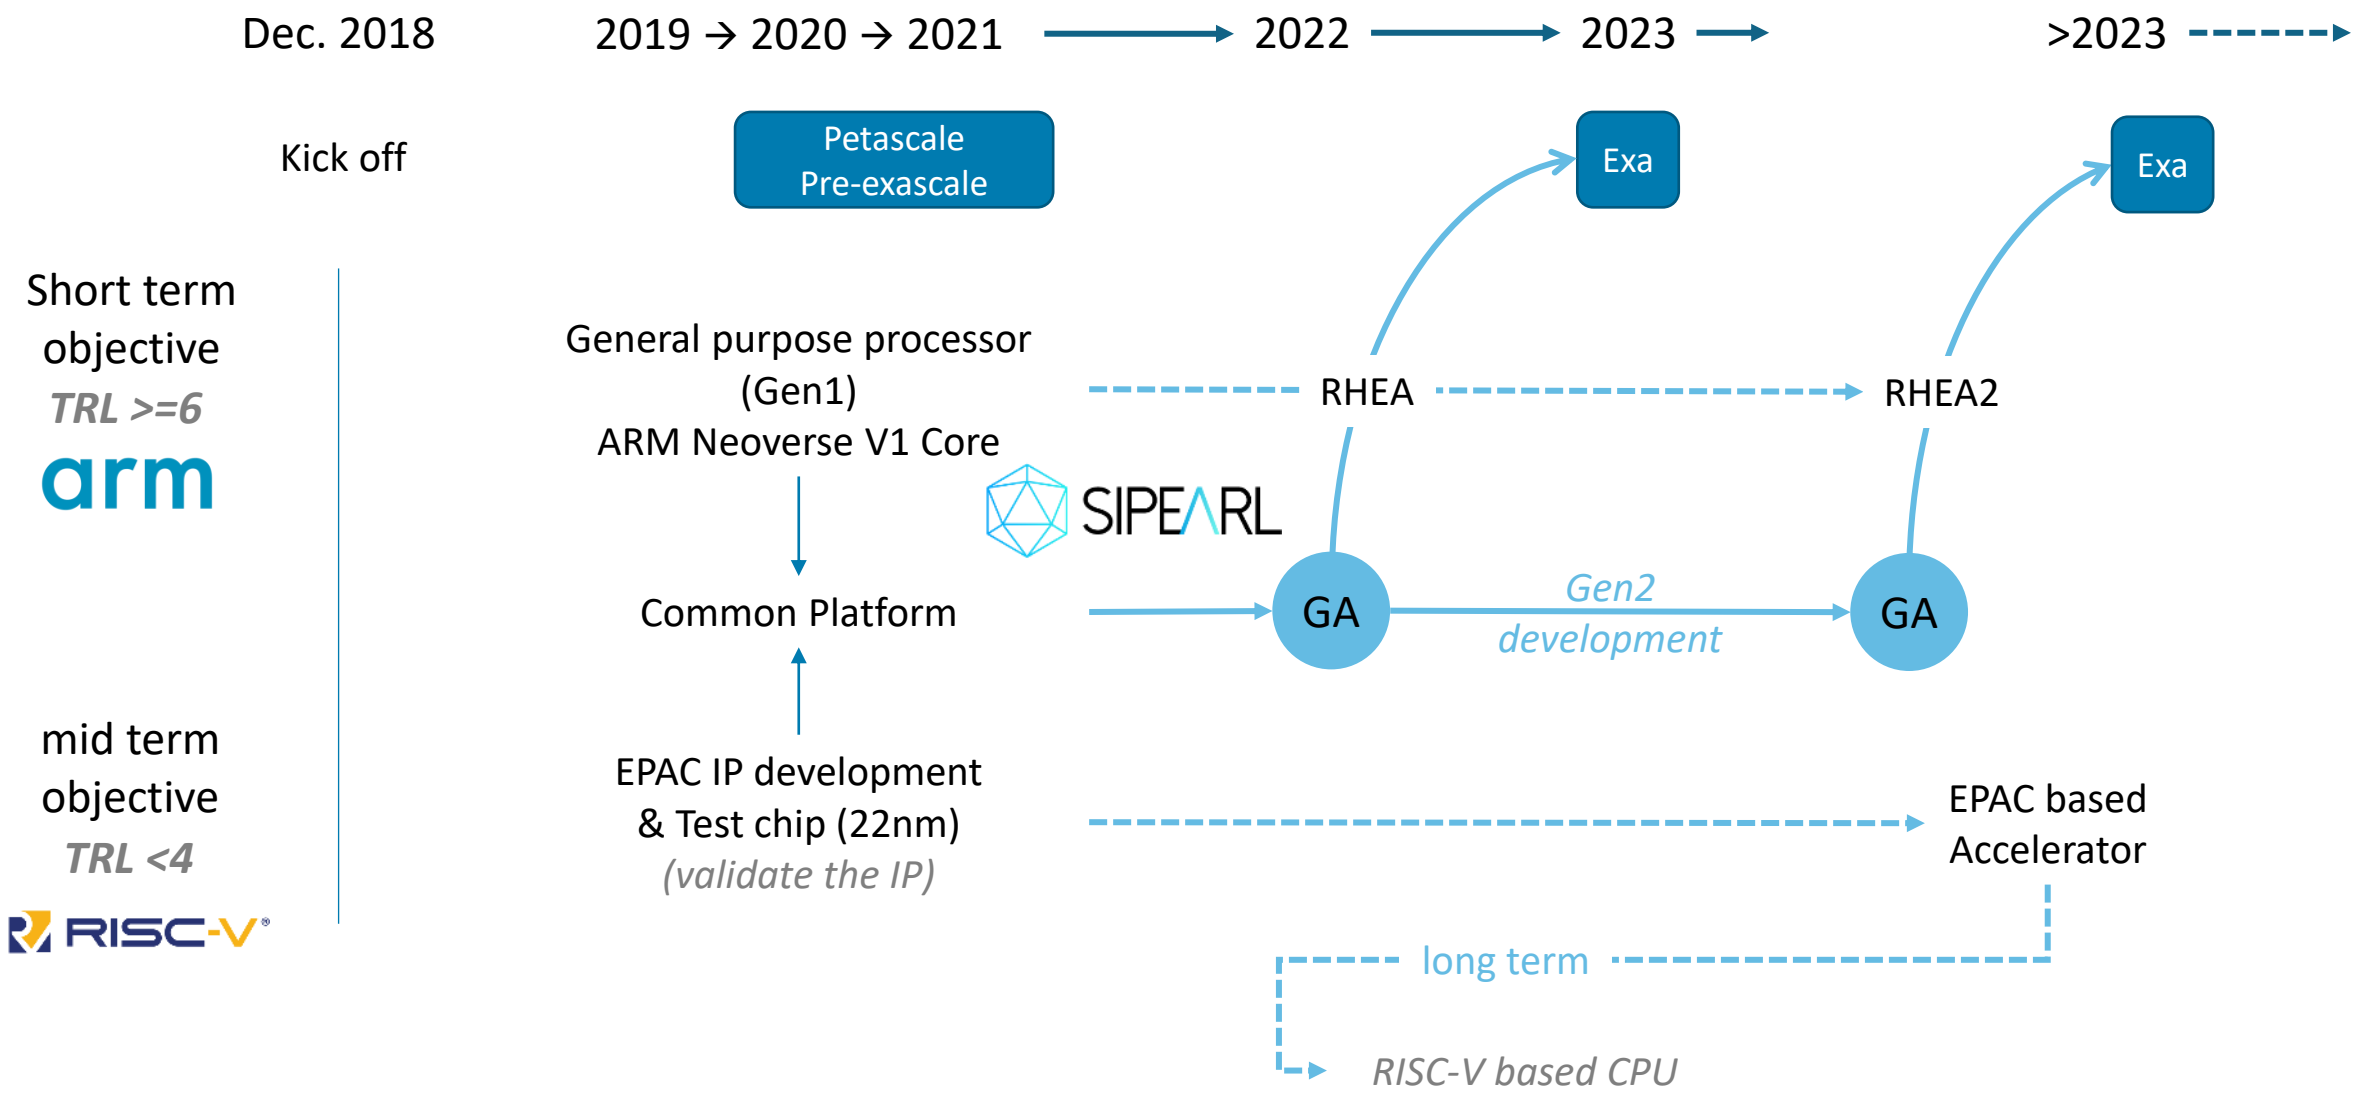




7

# EPI COMMON PLATFORM ENABLES EU ECOSYSTEM

- SiPearl chartered is also to develop the European Ecosystem
- SiPearl shares IP and benefits from IP ecosystem
  - Accelerator development (RISC-V based)
    - AI (tensor)
    - Vector processing
    - Stencil processing
    - FPGA
    - …
  - Packaging
  - IP development
- Staged integration: start with socket-to-socket connections and move into package (multi-chiplets) over time

now Future (10y?)

# Rhea
# a Processor for the Exa era

SIPEARL

# 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<br>- Neoverse V1 cores<br>- SVE 256 per core supporting 64/32/BF16 and Int8<br>- ArmVirtualization extensions |
| SoC | - Arm mesh fabric<br>- Advanced RAS support including Arm RAS extensions<br>- Link protection for NoC & high-speed IO<br>- ECC support for selected memory |
| Cache | - Large L3 (Shared Level Cache)<br>- RAS supported for all cache levels |
| Memory | - HBM2e<br>- And DDR5<br>- ECC for memory and link protection for controllers |
| High Speed I/O | - PCIe, CCIX & CXL<br>- 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<br>- Crypto<br>- True random number generation<br>- Made in Europe |

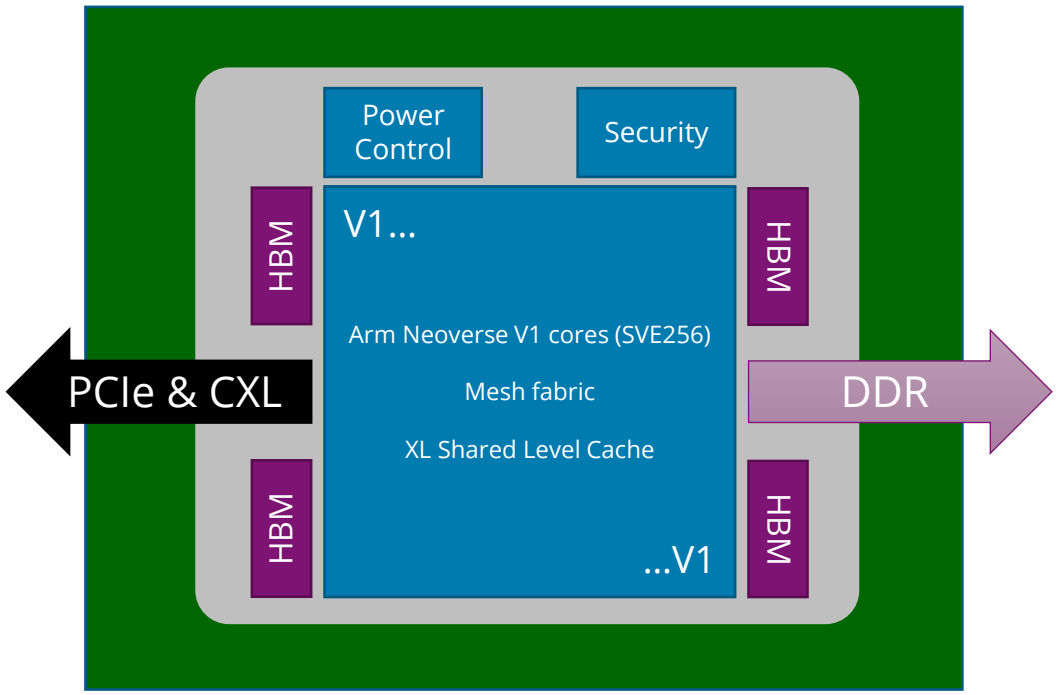


Rhea will deliver extraordinary real compute performance and efficiency with an unmatched Bytes/Flops ratio.

# PROCESSOR CORES INSIDE RHEA

| | Total: x V1 + 2 M7 Arm, 29 Risc-V, 4x SPUs. | Remarks |
|---|---|---|
| Arm Neoverse V1 cores | Arm Neoverse V1 | Including spare V1s. |
| Arm cortex-M7 cores | 2x Arm cortex-M7 = 2. | for SCP and MCP subsystems. |
| Risc-V in PMS | 1x Ariane + 1x ZeroRiscy = 2. | |
| Risc-V in SEG | 1x Ariane = 1. | SEG for security element. |
| Risc-V in STXs of 2x ERACs | 2x (1x Ariane + 8x Snitch) = 18. | |
| Risc-V in VRPs of 2x ERACs | 2x (4x VRP core) = 8. | VRP core is a modified Risc-V core. |
| SPUs in STXs of 2x ERACs | 2x (2x SPU cores) = 4. | SPU core is a proprietary core. |

- Some additional EU designed IP (power management, clock, cryptography) not counted here
- Not including µC cores used in Synopsys DDR controllers for the PHY training.

| Core | Performance for the core. |
|---|---|
| V1 | 2x 256 SVE = 16 DP FLOPs/cycle; 2.5GHz@N6 |
| Snitch | 1x 64b FPU = 2 DP FLOPs/cycle; >1GHz@N6 |
| SPU | 4x 32b FPU = 8 SP FLOPs/cycle; >1GHz@N6 |

# STENCIL: DESIGNED FOR CLASSICAL HPC MODELS

1. Designed in first place for finite difference and finite elements algorithms (example: CFD, FDTD, O&G)
2. Expanded to support wider class of algorithms while retaining efficiency
3. Ease of programmability as a design goal
   - **Accelerator for physicists rather than computer scientist**
   - **Also applicable to other domains e.g. weather forecasting, CFD and energy**

```
#pragma omp target
{
    #pragma stx loop
    for (int z = stencil_radius; z < dim_z - stencil_radius; z++)
    {
        #pragma stx loop(interleave)
        for (int y = stencil_radius; y < dim_y - stencil_radius; y++)
        {
            #pragma stx loop
            for (int x = stencil_radius; x < dim_x - stencil_radius; x++)
            {
                float dxyz = 0.0f;

                for (int r = 1; r <= stencil_radius; r++)
                {
                    float const weight = parameter (r - 1);

                    dxyz += (   device_pressure_wavefield[z + r][y    ][x    ]
                              + device_pressure_wavefield[z - r][y    ][x    ]
                            ) * weight;

                    dxyz += (   device_pressure_wavefield[z    ][y + r][x    ]
                              + device_pressure_wavefield[z    ][y - r][x    ]
                            ) * weight;

                    dxyz += (   device_pressure_wavefield[z    ][y    ][x + r]
                              + device_pressure_wavefield[z    ][y    ][x - r]
```

**Shown: ISO stencil example source code from our "spu-runtime" repository.**

User can write **natural, complex expressions** as indices for data accesses, "structured data configuratior is made automatically.

**data[x]** (HW-loop variable)
**data[r]** (non-HW loop variable)
**data[x +/- r]**
**data[x][y]** or **data[y][x]** (arbitrary order)
**data[x + (2 * r / 3)][y – ((r + 4) / 2)]**

Fraunhofer ITWM

Native support for plane scheme tripple buffering

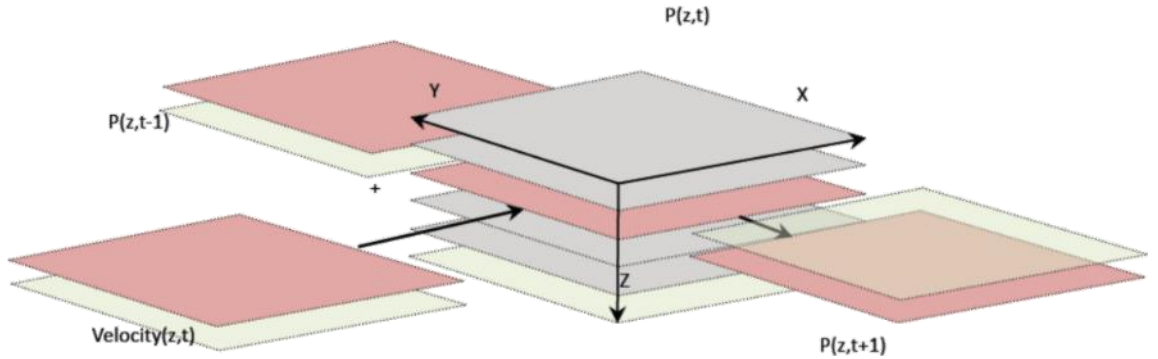


## Usecase: RTM TTI

**Structure**

- Intersecting planes of mixed derivations around a center point
- P- and Q-wavefields
- Velocity wavefield
- 4 additional physical parameter wave fields
- High arithmetic intensity:
  5 - 18 flops/byte

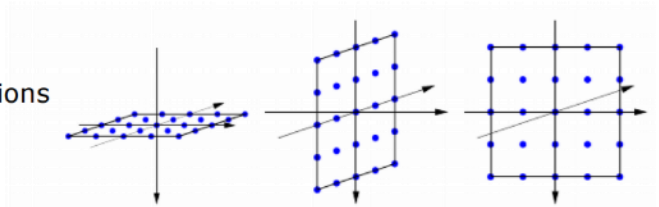

(a) Derivations $\partial_x\partial_z$, $\partial_z\partial_y$ and $\partial_y\partial_y$ (b) Derivations $\partial_y\partial_y$, $\partial_y\partial_z$ and $\partial_z\partial_z$ (c) Derivations $\partial_z\partial_z$, $\partial_z\partial_x$ and $\partial_x\partial_x$

- **Supporting**
  - RTM TTI 1-pass
  - RTM TTI 2-pass
  - Micro application kernel incorporating forward and backward propagation
  - Kernel optimizations included (common subexpression elemination, pre-computation, plane scheme, etc.)



SIPEARL

# Variable Precision processor (VRP)

## Definition

- The Variable Precision Processor (VRP) is a domain specific accelerator for scientific computing, specially tailored for the accurate computation (up to 512 bits fractional parts) of large systems of equations.
- It supports IEEE 754 extendable format in memory with byte-aligned data format to optimize memory usage and computing efficiency.

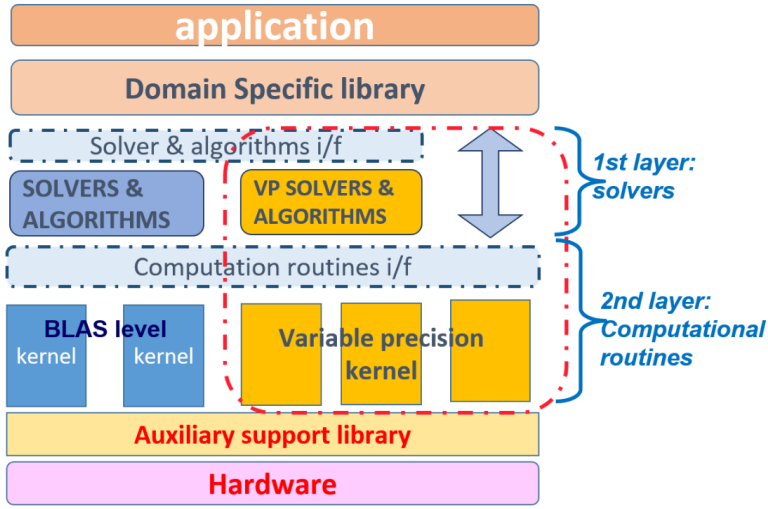


## Motivation

- Reduce conjugate gradient (CG), bi-CG iteration count
- Simplify preconditioning
- Allows direct solvers instead of indirect (matrix with bad conditioning number)
- Generally valid for many other algorithms, in particular for Krylov-based projective resolution.
- More investigation on lanczos based eigenvalue and singular-value solvers

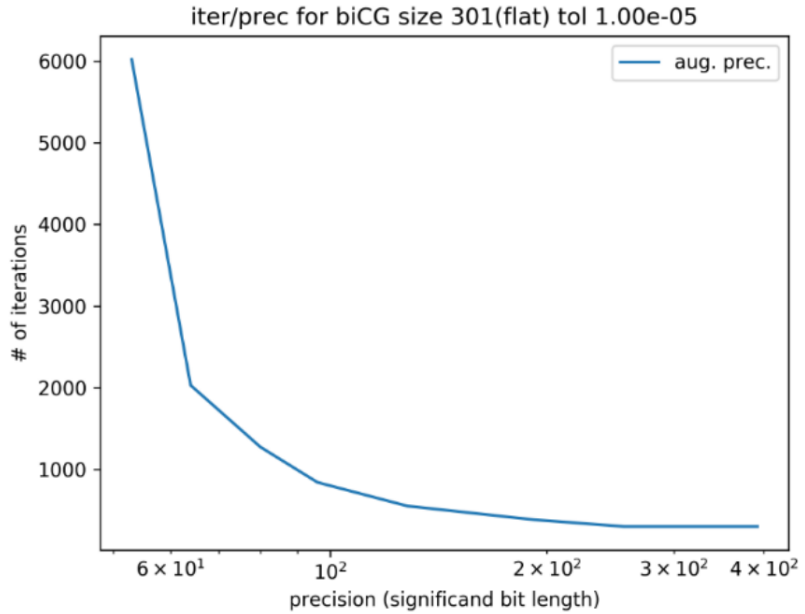


## Performances

- It targets 10x to 100x acceleration of variable precision
- computation (compared to software solutions).

# Rhea1 – Memory configurations

## Single socket

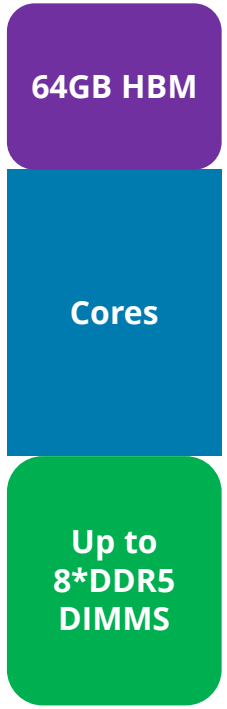


**Two different / independent memory spaces**

## dual sockets

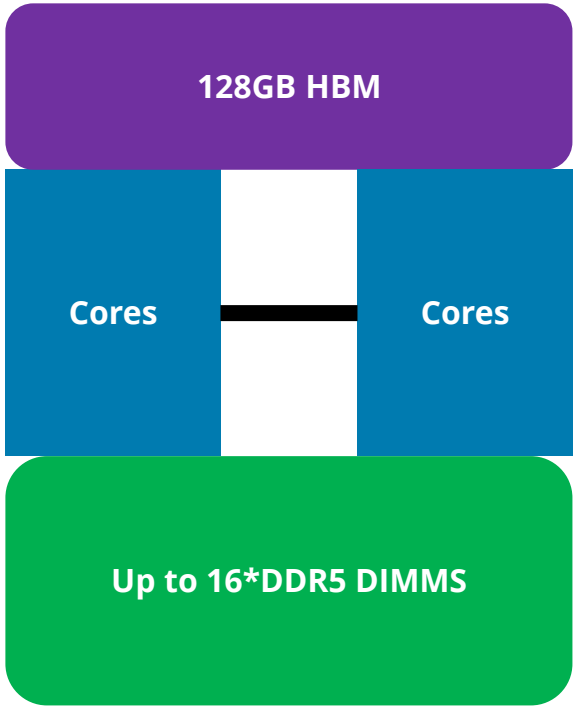


**Two different / independent memory spaces**
- **1 unified (CC-NUMA) HBM space**
- **1 unified (CC-NUMA) DDR space**

# Atos Reference Board with Rhea1

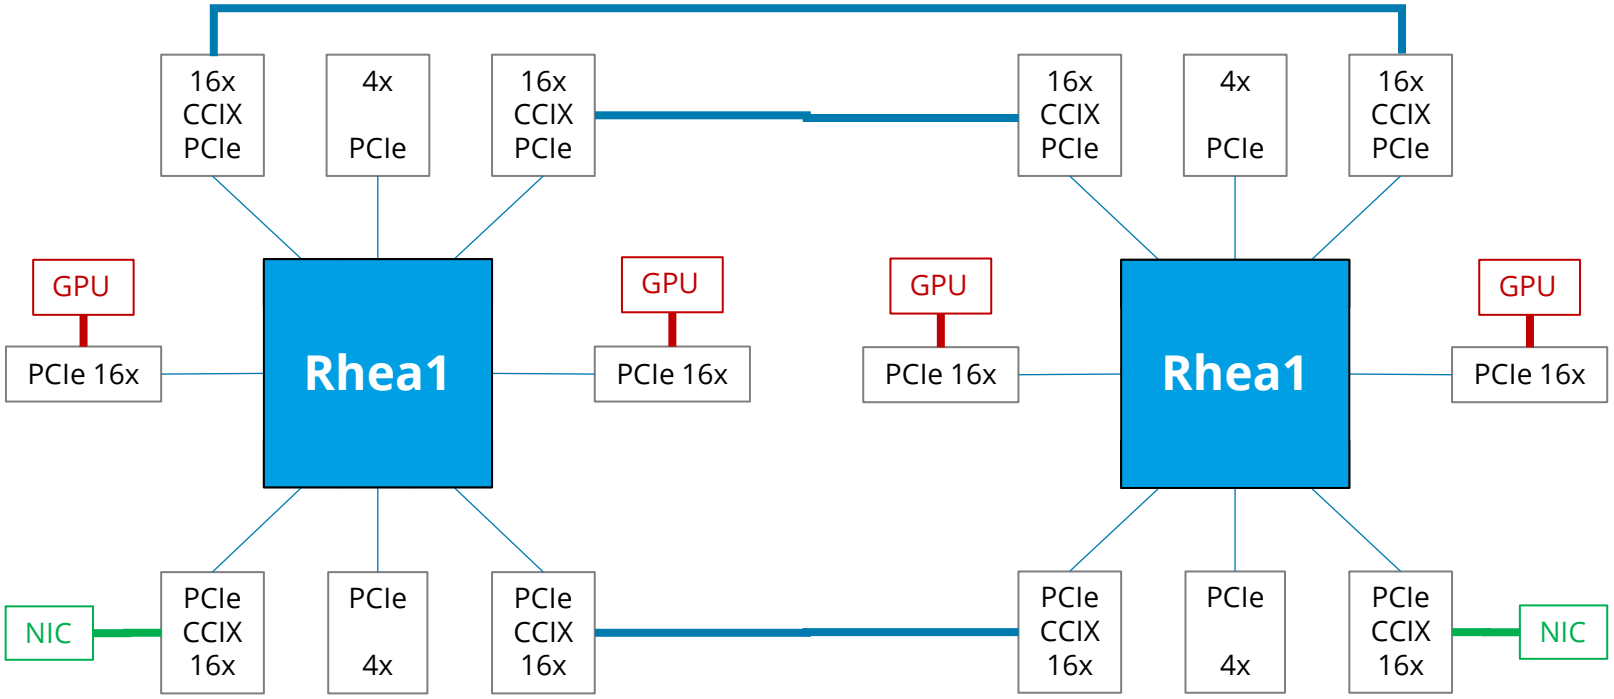


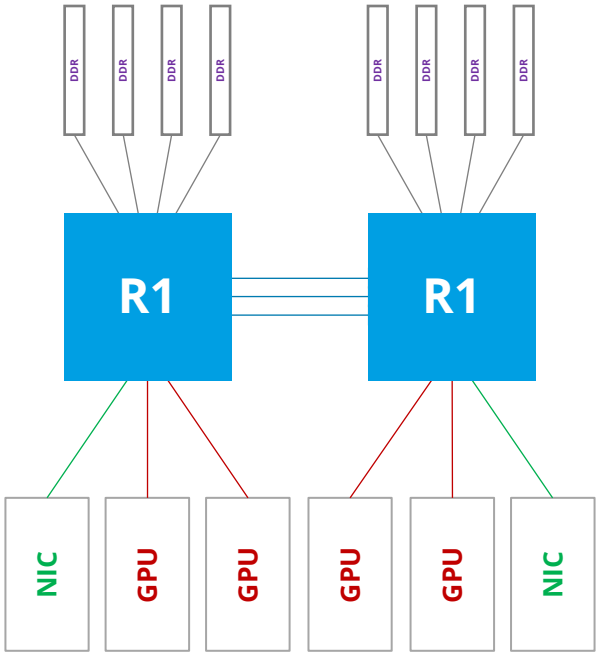


CCIX (16x) ➔ 3

GPU (16x) ➔ 4

NIC (16x) ➔ 2

| | R1a | | | R1b | | |
|---|---|---|---|---|---|---|
| | CCIX 16x to R1b | GPU 16x | NIC 16x | CCIX 16x to R1a | GPU 16x | NIC 16x |
| PCIe 16x 4 CCIX | **3** | | **1** | **3** | | **1** |
| PCIe 16x 2 | | **2** | | | **2** | |
| PCIe 4x 2 | | | | | | |

# Performance

SIPEARL

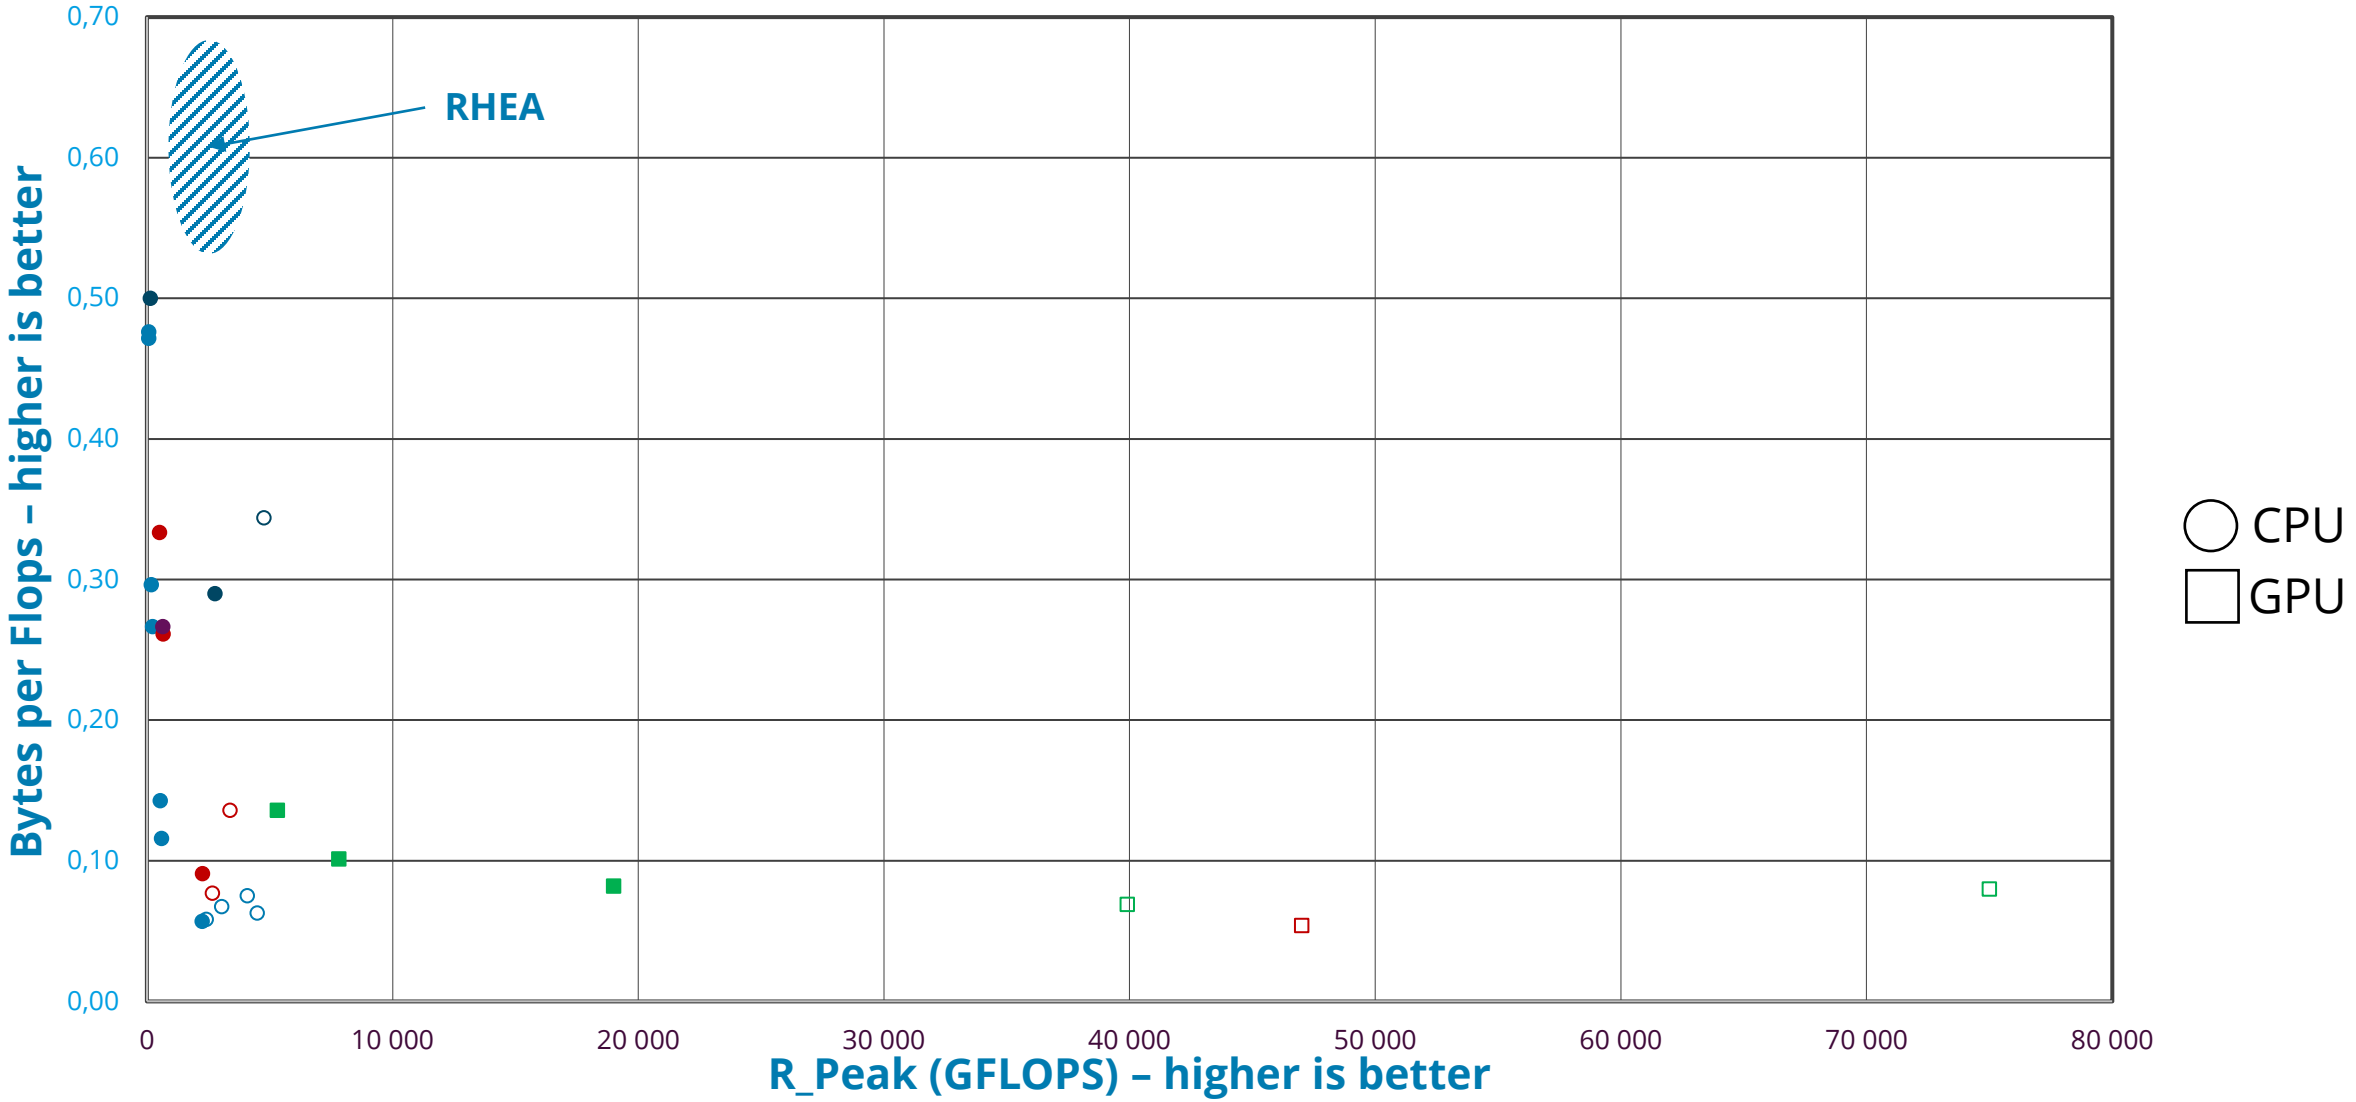

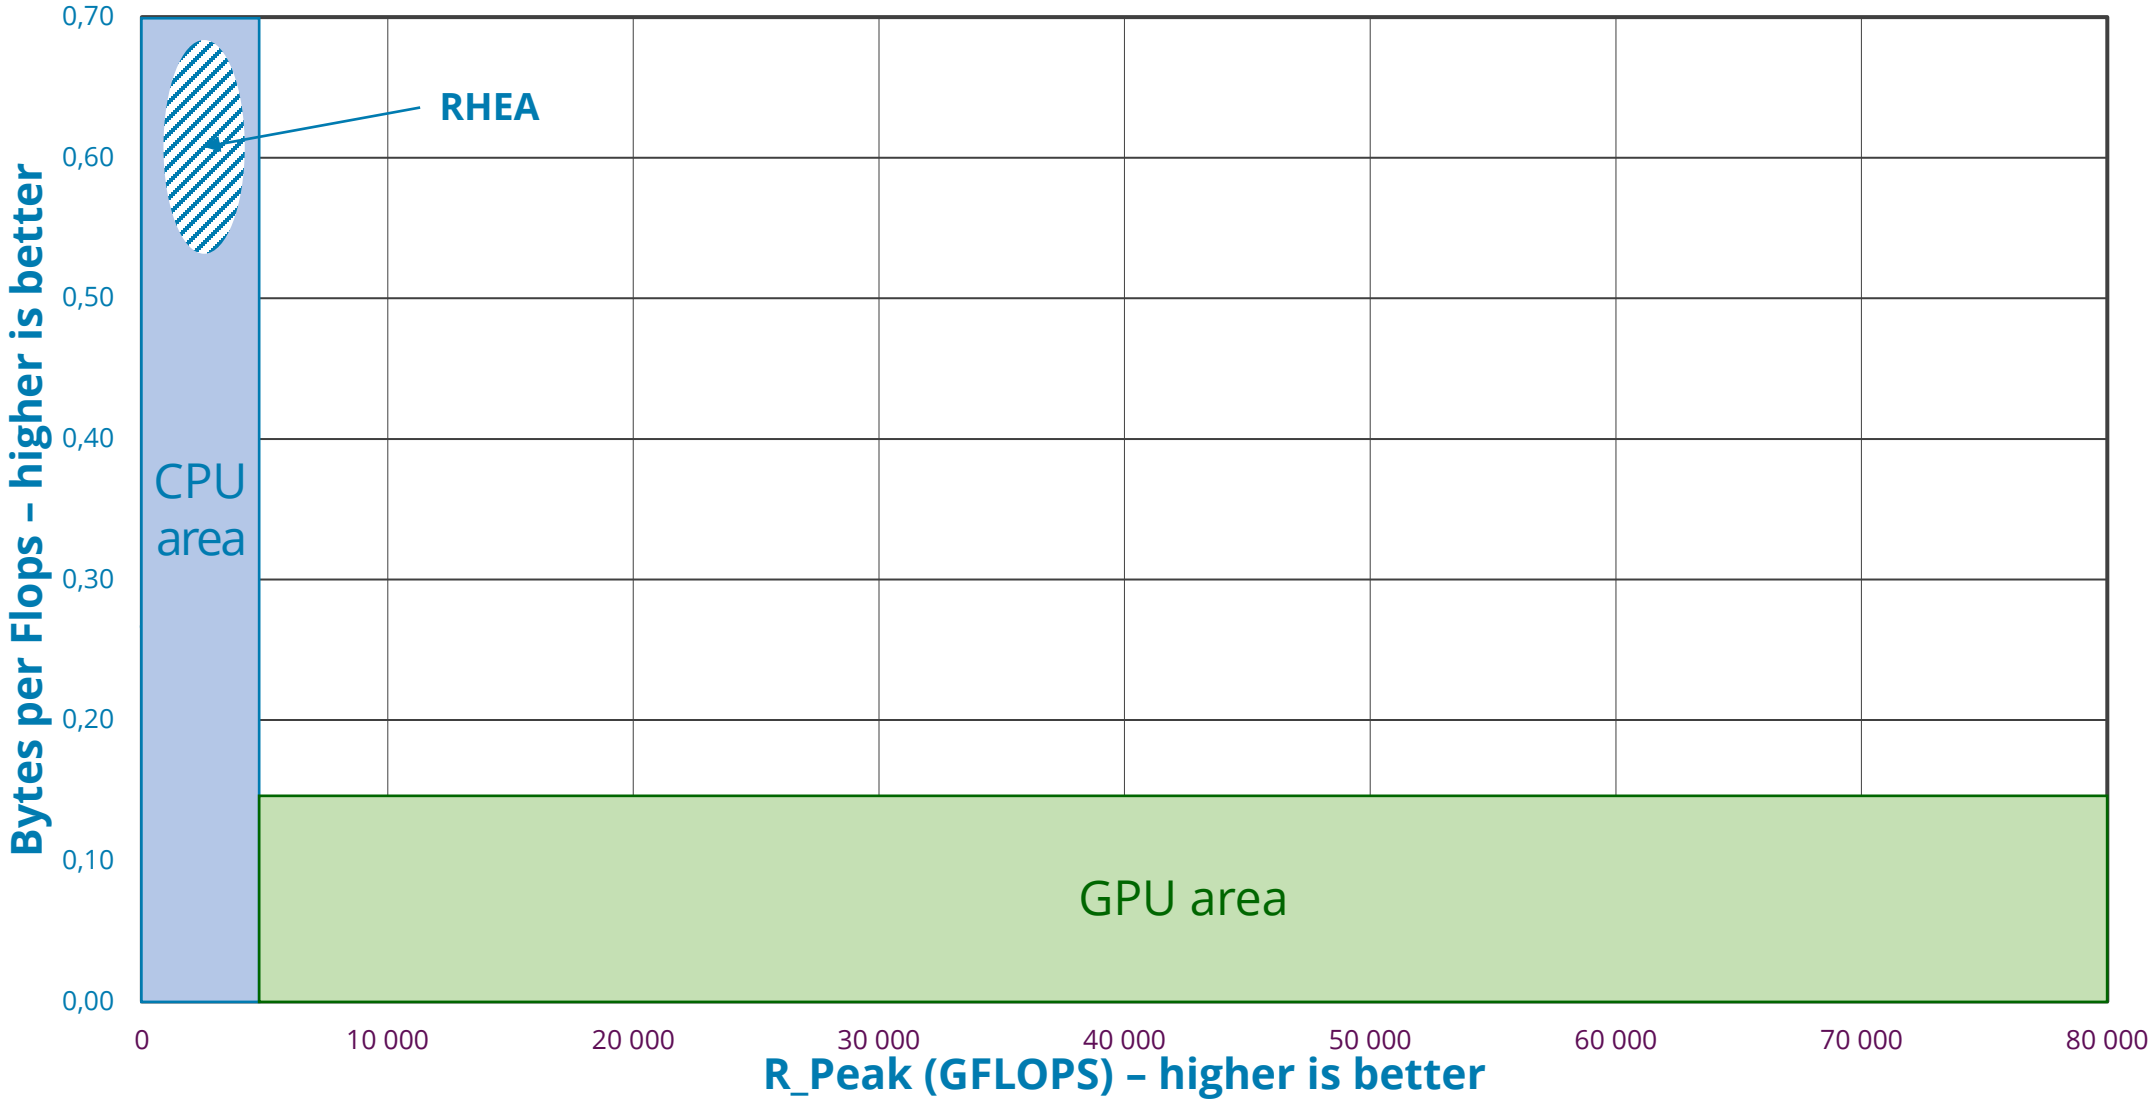

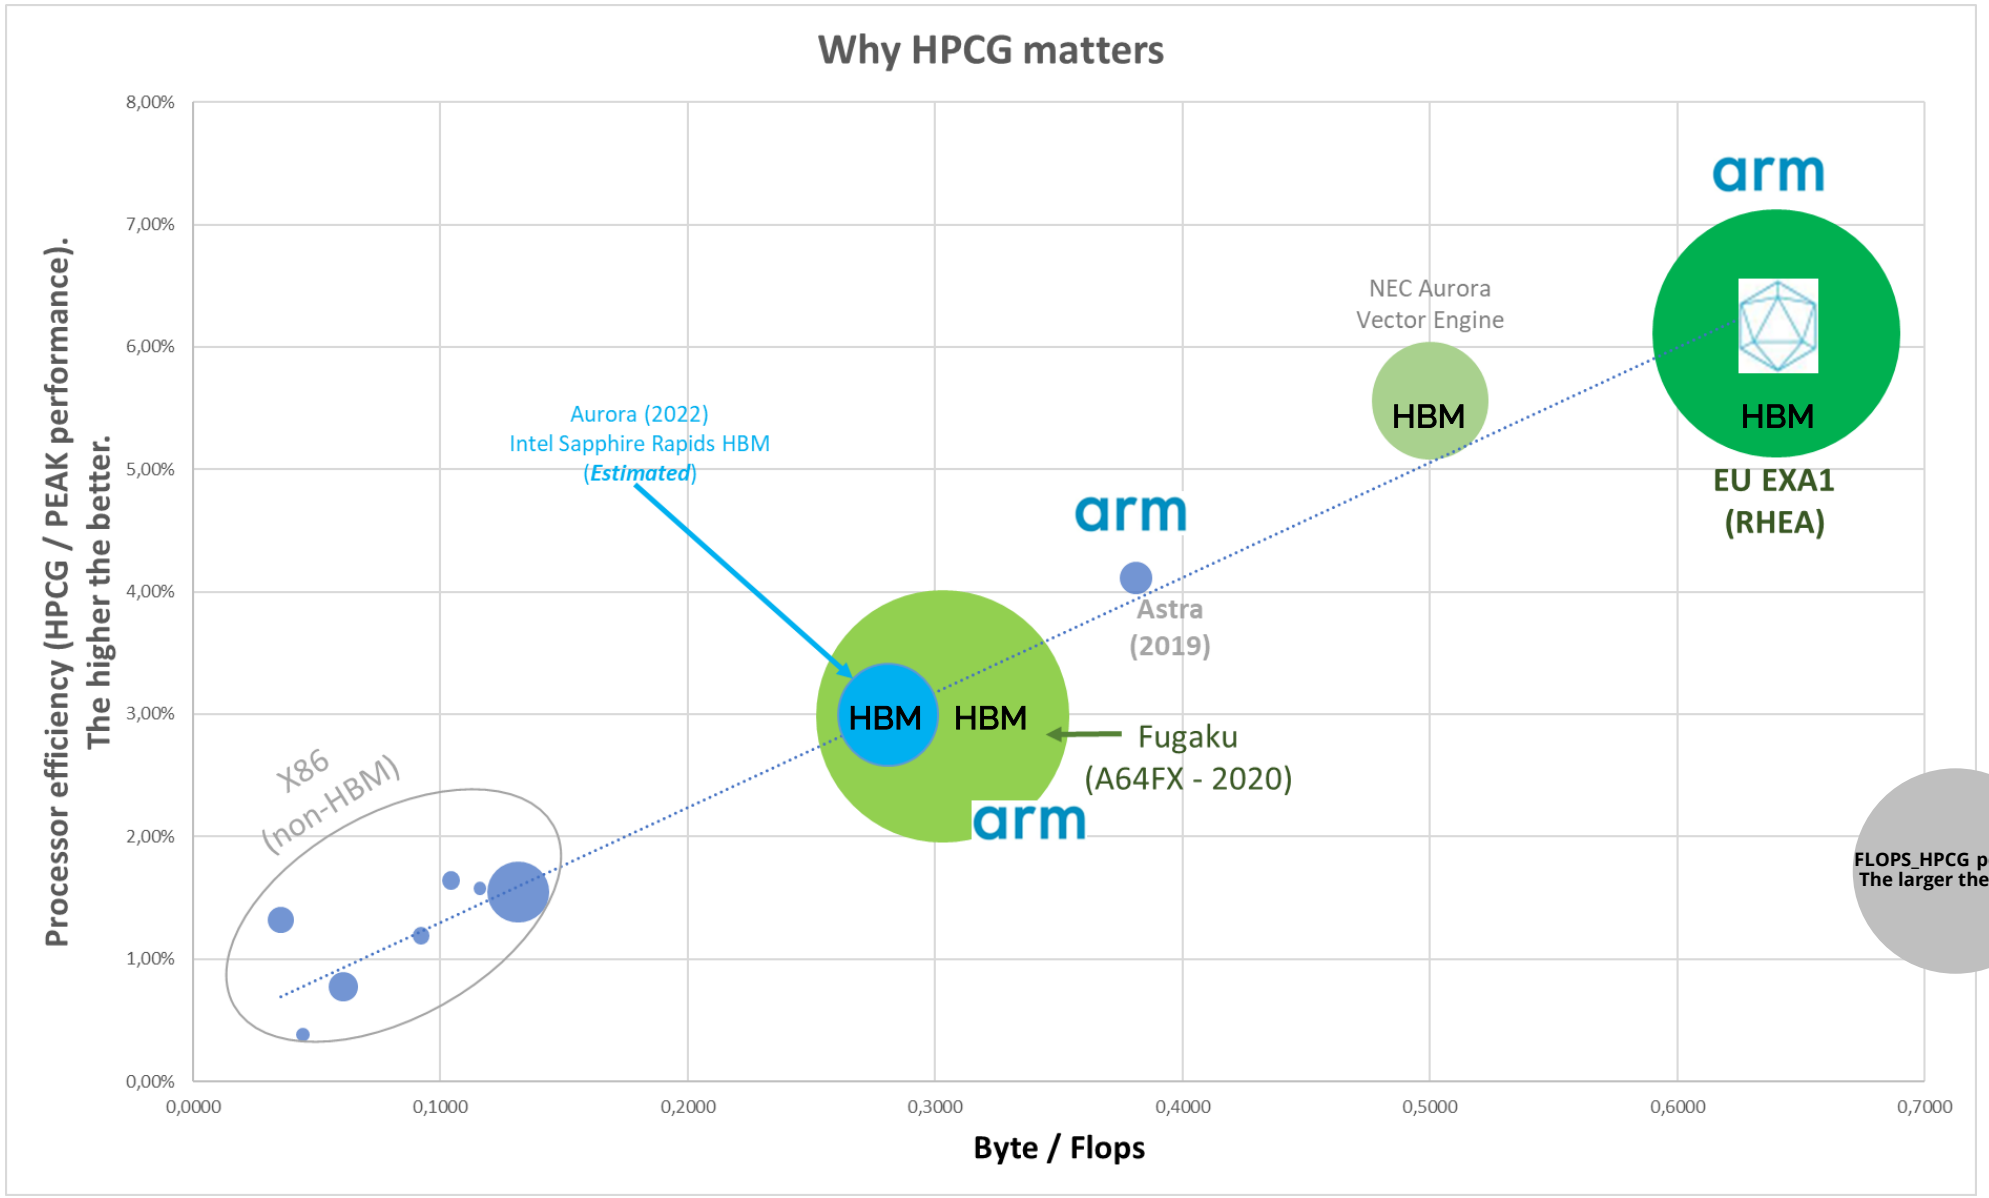

# INSTRUCTION LATENCY (SMALLER IS BETTER)

best in class | Sort-of-ex-aequo

| VFP64, full width | Fujitsu | Arm | Intel | Intel | AMD | AMD |
|---|---|---|---|---|---|---|
| **Latency** | A64FX | Neoverse V1 | Broadwell | Skylake-X | Rome / Zen2 | Milan / Zen 3 |
| Add | 9 | 2 | 3 | 4 | 3 | 3 |
| Mul | 9 | 3 | 3 | 4 | 3 | 3 |
| FMA | 9 | 4 (2 if chained) | 5 | 4 | 5 | 4 |
| Div | 43 | 7 to 15 | 19-23 | 24 | 13 | 13,5 |
| Sqrt | 43 | 7 to 16 | 27-29 | 28-29 | 20 | 20 |
| **Throughput** | | | | | | |
| Add | 2 | 2 | 1 | 2 | 2 | 2 |
| Mul | 2 | 2 | 2 | 2 | 2 | 2 |
| FMA | 2 | 2 | 2 | 2 | 2 | 2 |
| Div | 1/43 | 1/14 to 1/7 | 1/16 | 1/16 | 1/5 | 1/(4.5) |
| Sqrt | 1/43 | 1/14 to 1/7 | 1/28 to 1/16 | 1/24 to 1/18 | 1/9 | 1/9 |
| Max SIMD | SVE [512] | SVE [256] | AVX2 [256] | AVX-512 | AVX2 [256] | AVX2 [256] |
| | NEON, scalar have the same throughput | Neon, Scalar have twice the throughput (4x128 instead of 2x256) | XCC-based cores; SSE, AVX, Scalar have the same throughput | XCC-based cores; SSE, AVX, Scalar have the same throughput | SSE, AVX, Scalar have the same throughput | SSE, AVX, Scalar have the same throughput |
| | https://github.com/fujitsu/A64FX/blob/master/doc/A64FX_Microarchitecture_Manual_en_1.4.pdf | https://developer.arm.com/documentation/pjdoc466751330-9685/5-0 | https://www.agner.org/optimize/instruction_tables.pdf | https://www.agner.org/optimize/instruction_tables.pdf | https://www.agner.org/optimize/instruction_tables.pdf | https://www.agner.org/optimize/instruction_tables.pdf |

# 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 109* 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 program under specific grant agreement no.826647).

* as of June 15th 2022

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