



SIPEARL

Rhea the perfect chip for Co-design

ECLAT / Bordeaux

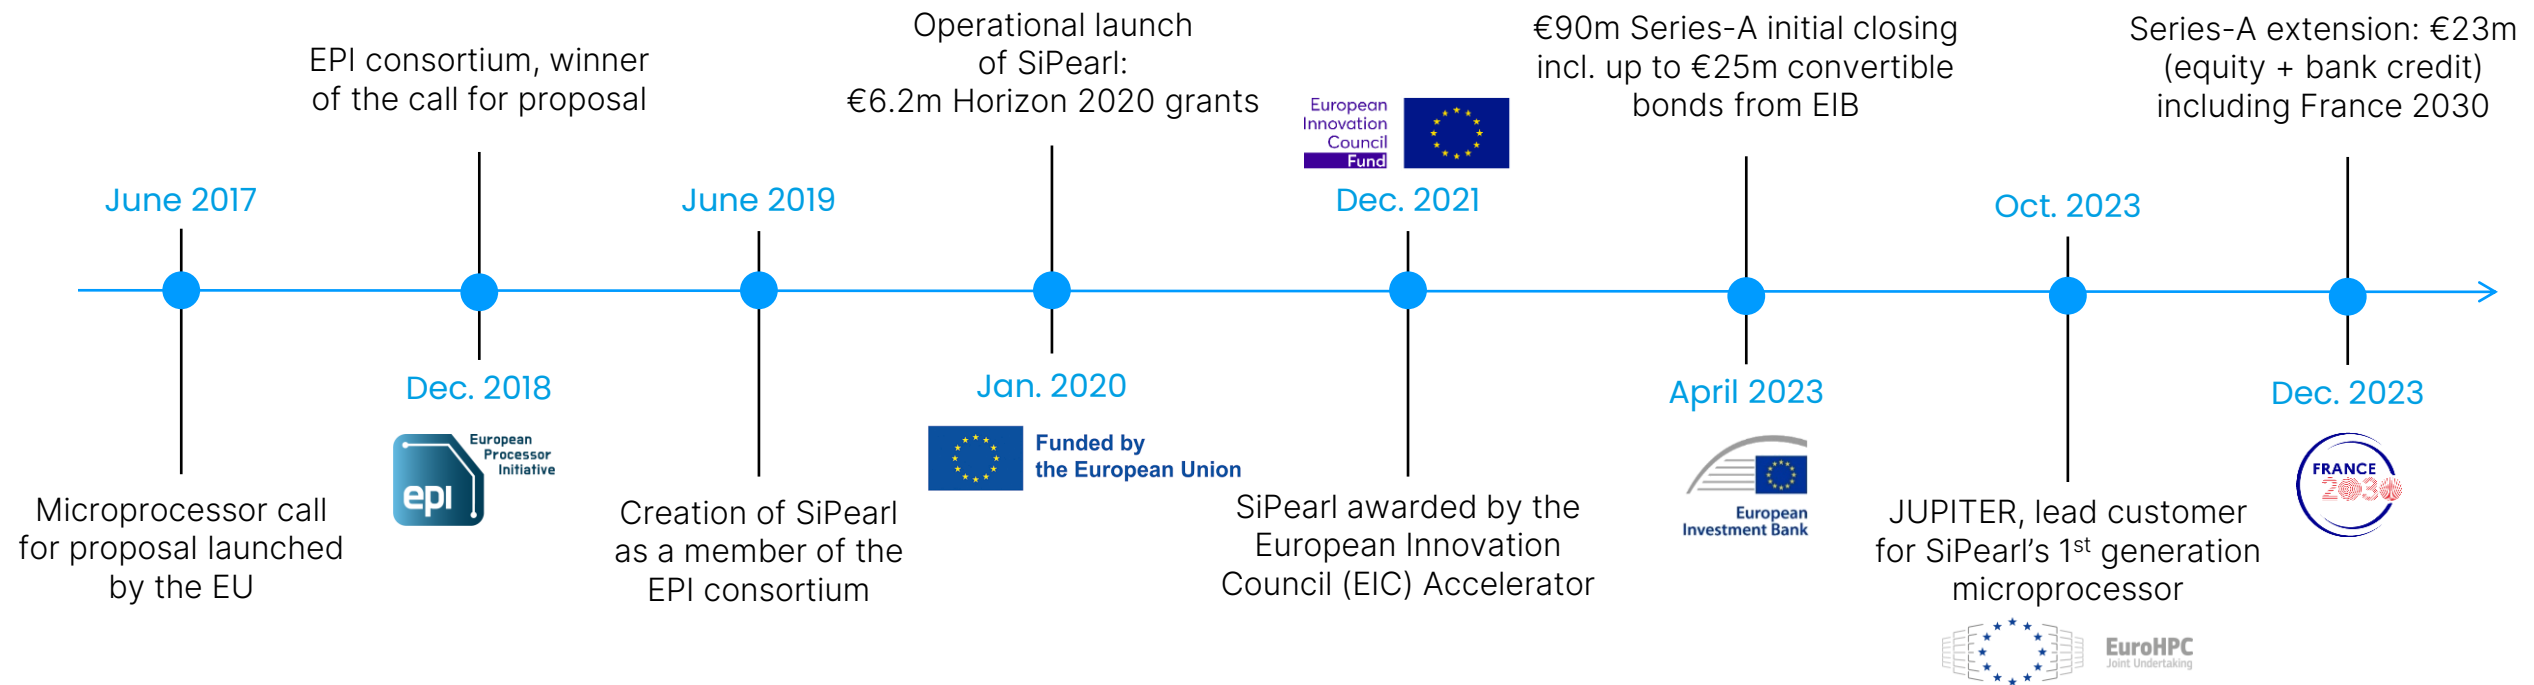
November 2024



SiPearl's history and mission

From a European Union concern to SiPearl's ramp-up

Our common goal: fostering the return of high-performance, low-power microprocessor technologies in Europe.



European Union Initiative

September 2018



Launch of the EuroHPC JU backed by a €8bn budget to deploy in Europe a world class exascale⁽¹⁾ supercomputing infrastructure.

December 2018



Launch of a call for proposals in 2017 for developing a new generation of high-end European microprocessors

- Budget: €150m
- Target: high-performance and energy-efficiency

Coordinated by Bull (Atos Group), the European Processor Initiative (EPI) consortium won this call for proposals. It has currently 30 members:

- Scientists: research institutes, universities and supercomputing centres
- Industry: European leaders, IT, electronics and automotive specialists

June 2019



SiPearl is the private company created within the EPI to launch a strategic industry for Europe.

(1) Exascale: billion of billion calculations per second

Our EPI partners, a powerful ecosystem

Close collaboration with our partners of the EPI consortium

Scientists: research institutes, universities and supercomputing centres.

Industry: European leaders, IT, electronics and automotive specialists.

A joint project involving 200 engineers since December 2018

- Development of elementary hardware and software technological bricks.

Stakeholders

- Privileged access to IP of European leaders and innovative startups.

End-users

- Supercomputing centres.



SiPearl in a nutshell

Building the European high-performance low-power microprocessor



Incorporated

In June 2019



Funded

By the European Union



Arm architecture

Energy-efficiency quick time to market, proven ecosystem



Financing

Series-A to date: €113m
(€105m equity + €8m bank loans)



Key partnerships

Joint-offering with



Identified customers

Server manufacturers based on user specifications: First, EuroHPC ecosystem before going global.

+200

Employees

from



7 locations

Maisons-Laffitte (HQ), Barcelona, Bologna, Duisburg, Grenoble, Massy, Sophia Antipolis

Rhea1 in 5 key benefits

The European high-performance low-power microprocessor dedicated to HPC and artificial intelligence.



High Performance

To surpass the performance of 10,000,000 desktop computers.



Energy-efficiency

To halve power consumption for equivalent computing power.



Flexibility

Designed to work with any third-party accelerator (GPU, artificial intelligence, quantum).



Backdoor-free security

To protect data with secure end-to-end network transmission.



Sovereignty

To further Europe's technological leadership and independence.

Technology partnerships

with leading providers



Our main partner

SiPearl, the only European licensee to use arm Neoverse V1 platform



EDA software



Validation of semiconductor power integrity, minimization of power consumption



Hardware emulation with Veloce Strato emulation platform

- 128 new generation cards
- x 1,000 simulation speed
- Unique in Europe



Advanced High-Bandwidth Memory solution for Rhea

- Speed
- Energy
- Thermal resistance



Manufacturing initially entrusted to TSMC

Etching: 6nm or better for next generations

World leading industrial partnerships

Our ecosystem to accelerate Europe's adoption of exascale supercomputers

Leading manufacturers

EVIDEN

Europe No1



**Hewlett Packard
Enterprise**

Global leader

Developing joint solutions for HPC applications

Acceleration specialists



Instinct™ accelerators
with ROCm™ open software



GPU Ponte Vecchio
with the open unified software stack oneAPI



Portfolio of world-leading accelerated
computing and networking solutions

JUPITER, lead customer

JUPITER, 1st European exascale supercomputer owned by EuroHPC, operated by Jülich (Germany)

Built by a European consortium

- Eviden: the Atos Group business leading in advanced computing
- ParTec: the German modular supercomputing company


General-purpose Cluster Module of JUPITER to be based on Rhea1

- Very high memory bandwidth
- Extraordinary compute performance and efficiency

JUPITER Key Figures



€273m
expected budget



4 tennis courts
equivalent surface




10 million desktop computers
equivalent power



260km
high-performance cabling



Our business: the high-performance energy-efficient microprocessor dedicated to HPC⁽¹⁾ and AI



Tens of thousands of microprocessors in a supercomputer


RHEA1

HPC and AI microprocessor

80 arm® Neoverse V1 cores
with 2 x 256 SVE each

4 x HBM

4 x DDR5 interfaces



Rhea1, our 1st generation microprocessor

Designed with high-performance energy-efficient Arm Neoverse V1 platform



High performance per watt: Arm ISA power efficiency

Very high memory bandwidth

Built-in HBM

- Ideal performances for Generative AI

Unique memory architecture: High Byte/Flop


Openness

- Arm ecosystem from IoT/edge to HPC and cloud

Fully auditable – backdoor-free

Pre-integration with proven accelerator (AMD, Graphcore, Intel, Nvidia)

Rhea1 will deliver extraordinary performance and efficiency with an unmatched Byte/Flop ratio.





A peak inside Rhea

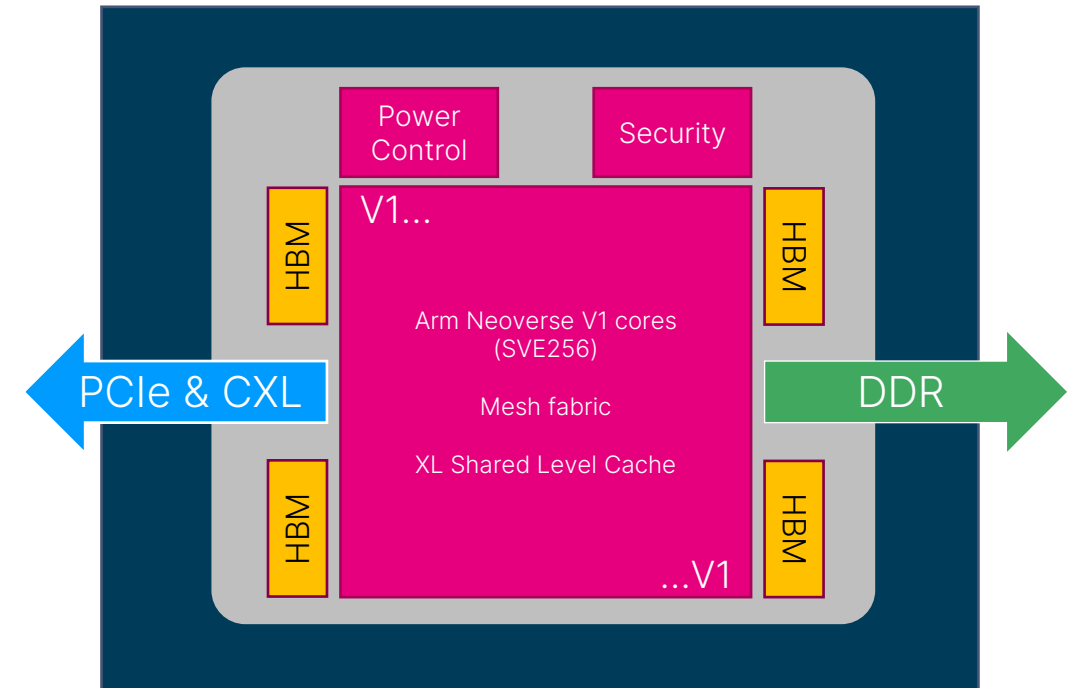
At the heart of Rhea

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



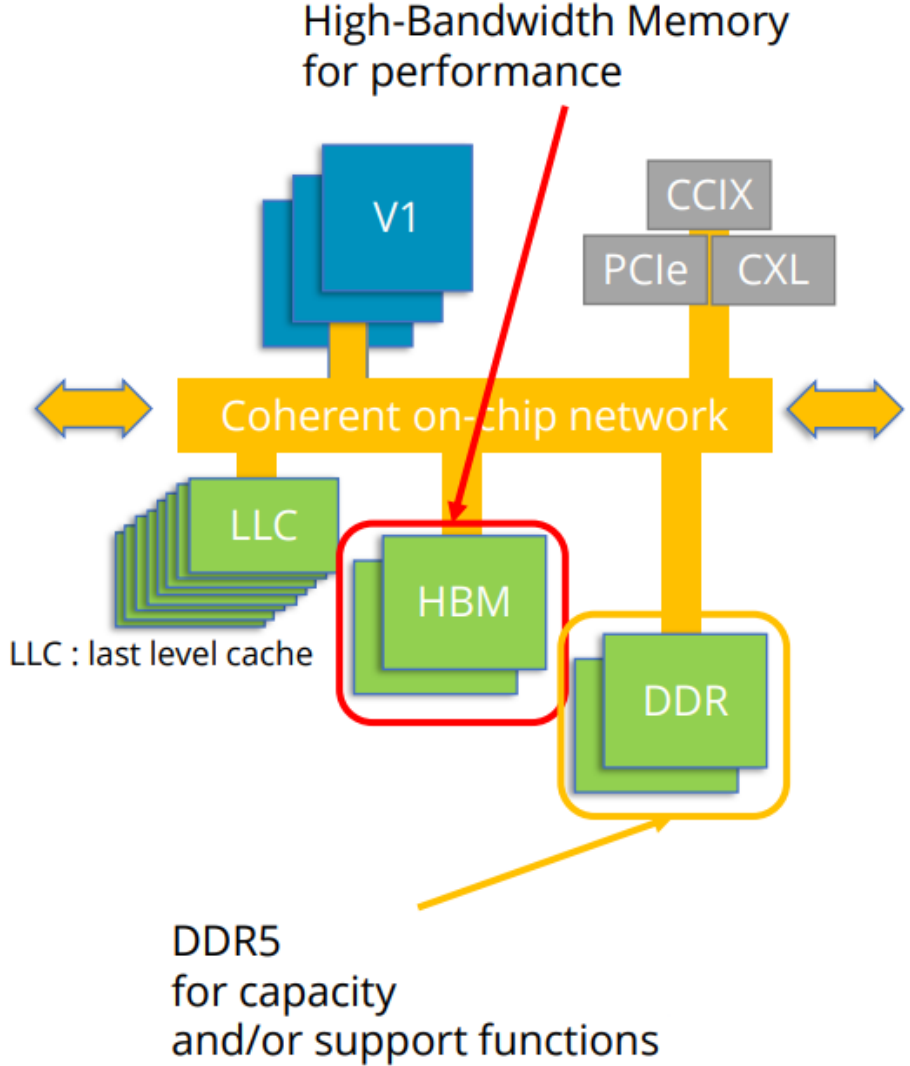
Key features

Core	<ul style="list-style-type: none">- Arm architecture- Neoverse V1 cores- SVE 256 per core supporting 64/32/BF16 and Int8- ArmVirtualization extensions
SoC	<ul style="list-style-type: none">- Arm mesh fabric- Advanced RAS support including Arm RAS extensions- Link protection for NoC & high-speed IO- ECC support for selected memory
Cache	<ul style="list-style-type: none">- Large L3 (Shared Level Cache)- RAS supported for all cache levels
Memory	<ul style="list-style-type: none">- HBM2e- And DDR5- ECC for memory and link protection for controllers
High Speed I/O	<ul style="list-style-type: none">- PCIe, CCIX & CXL- Root and endpoint support
Other I/O	<ul style="list-style-type: none">- USB, GPIO, SPI, I²C
Power Management	<ul style="list-style-type: none">- Power management block to optimize perf/watt accross use cases and workloads.
Security Block Support	<ul style="list-style-type: none">- 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.

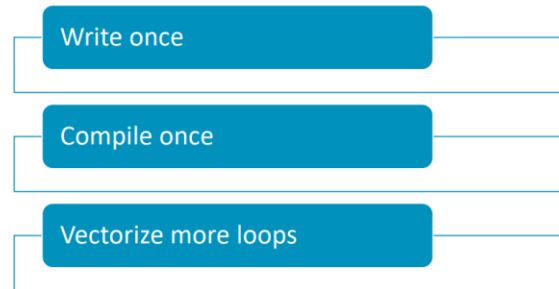
Influential parameters on Rhea



Vector Length Agnostic programming model

VLA

Vector Length Agnostic
programming model



Pros	Considerations
Not thinking about vector length	Should not be writing fixed width
No peel/remainder loops	Generate predicated instructions
Key is loop structures	Cost of generating predicates Near empty loops

SVE main features

Per lane predication

Operations work on individual lanes under control of a predicate register

	1	2	3	4
	5	5	5	5
+				
<i>pred</i>	1	0	1	0
=	6	2	8	4

Vector partitioning & software-managed speculation

First Faulting Load instructions allow memory accesses to cross into invalid pages

	1	2		
	↓	↓		
+	1	2	0	0
<i>pred</i>	1	1	0	0

Predicate-driven loop control and management

Eliminate scalar loop heads and tails by processing partial vectors

```
for (i = 0; i < n; ++i)
```

<i>INDEX i</i>	n-2	n-1	n	n+1
<i>CMPLT n</i>	1	1	0	0

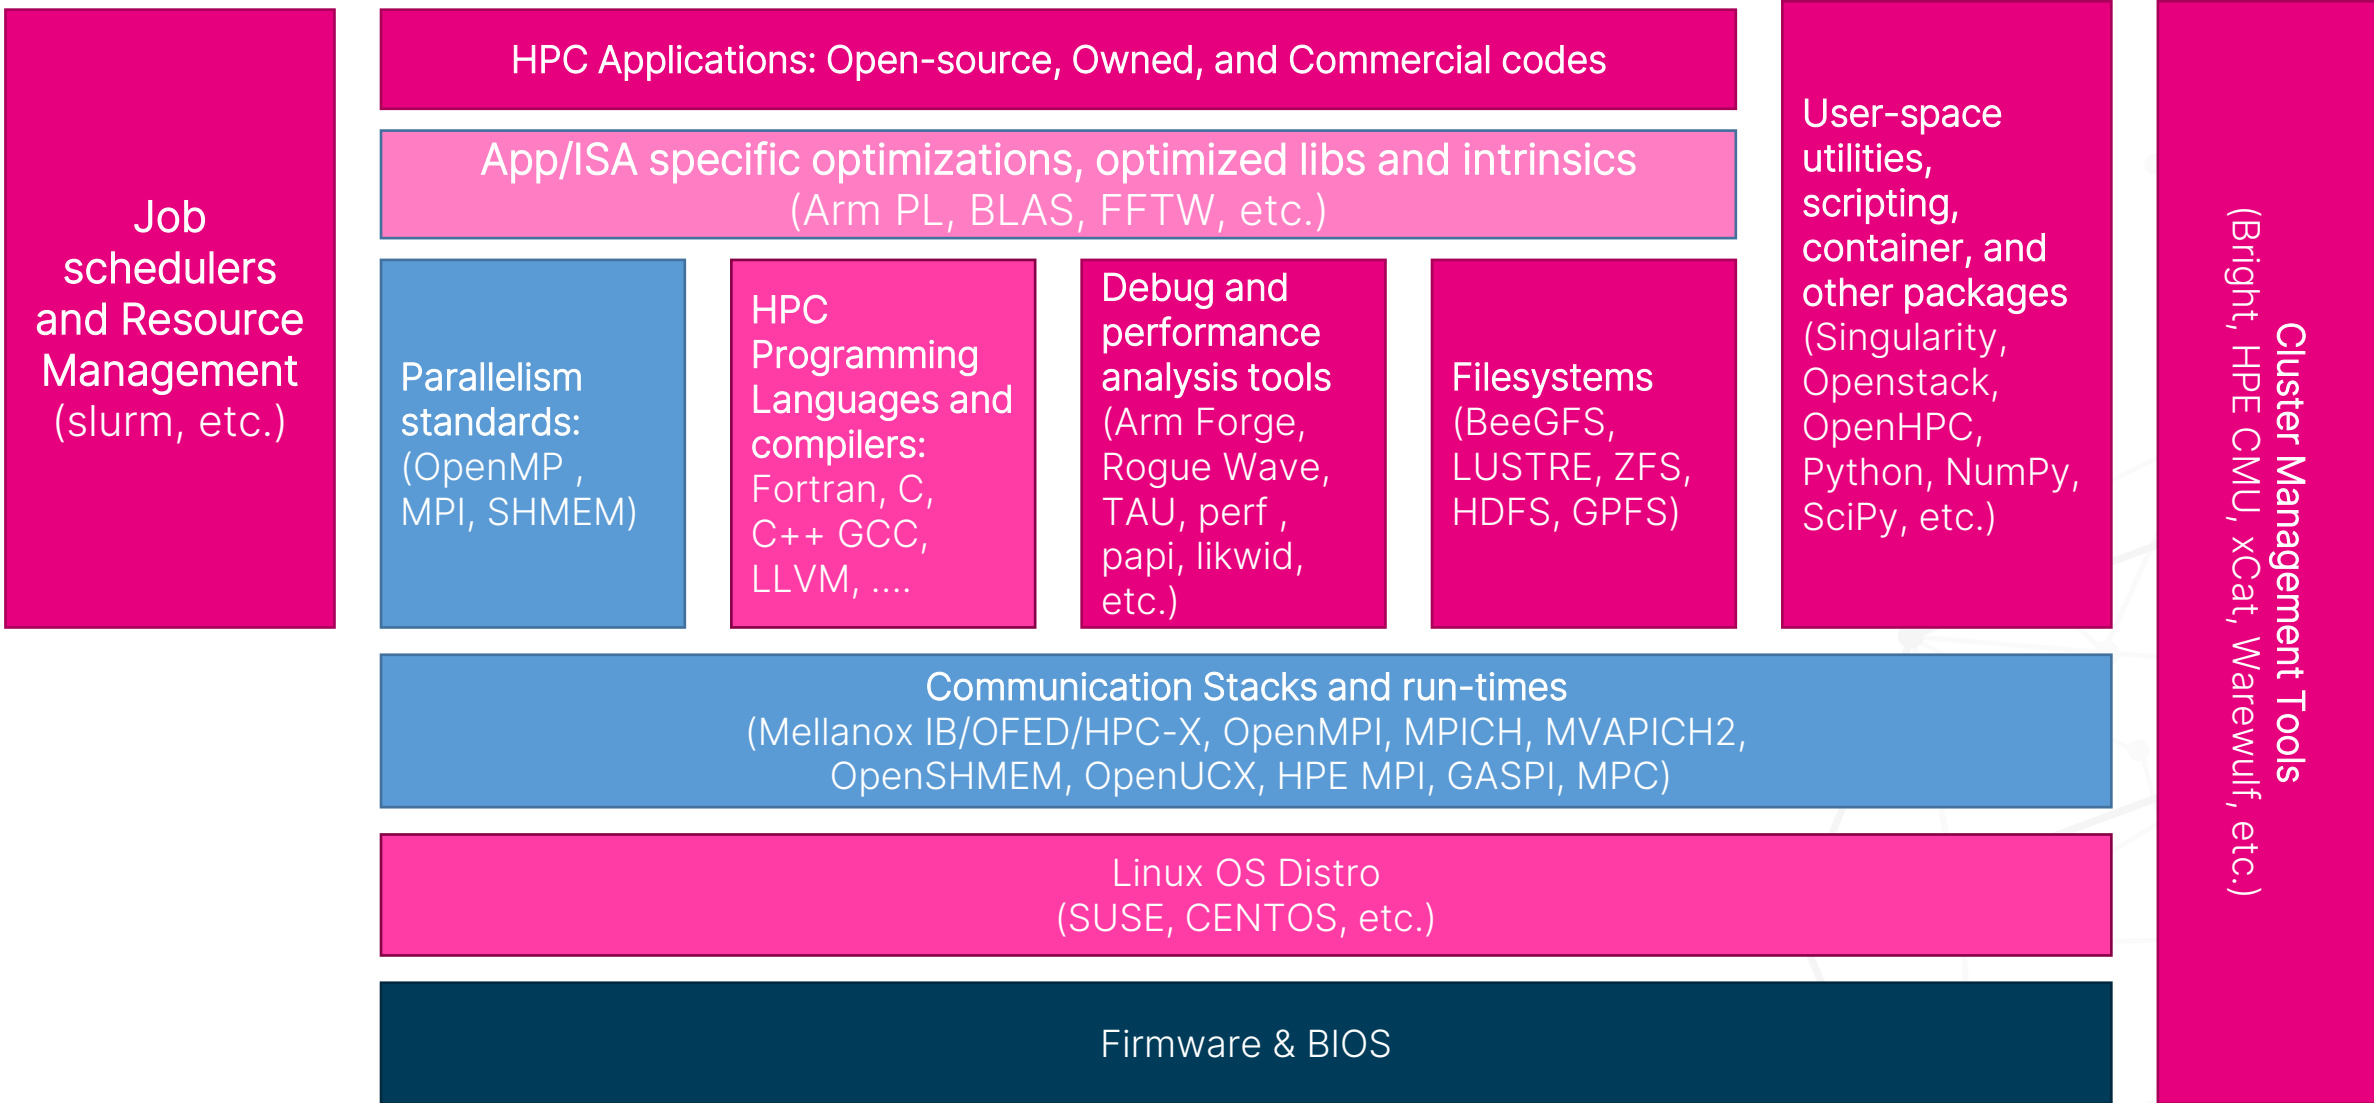


SIPEARL

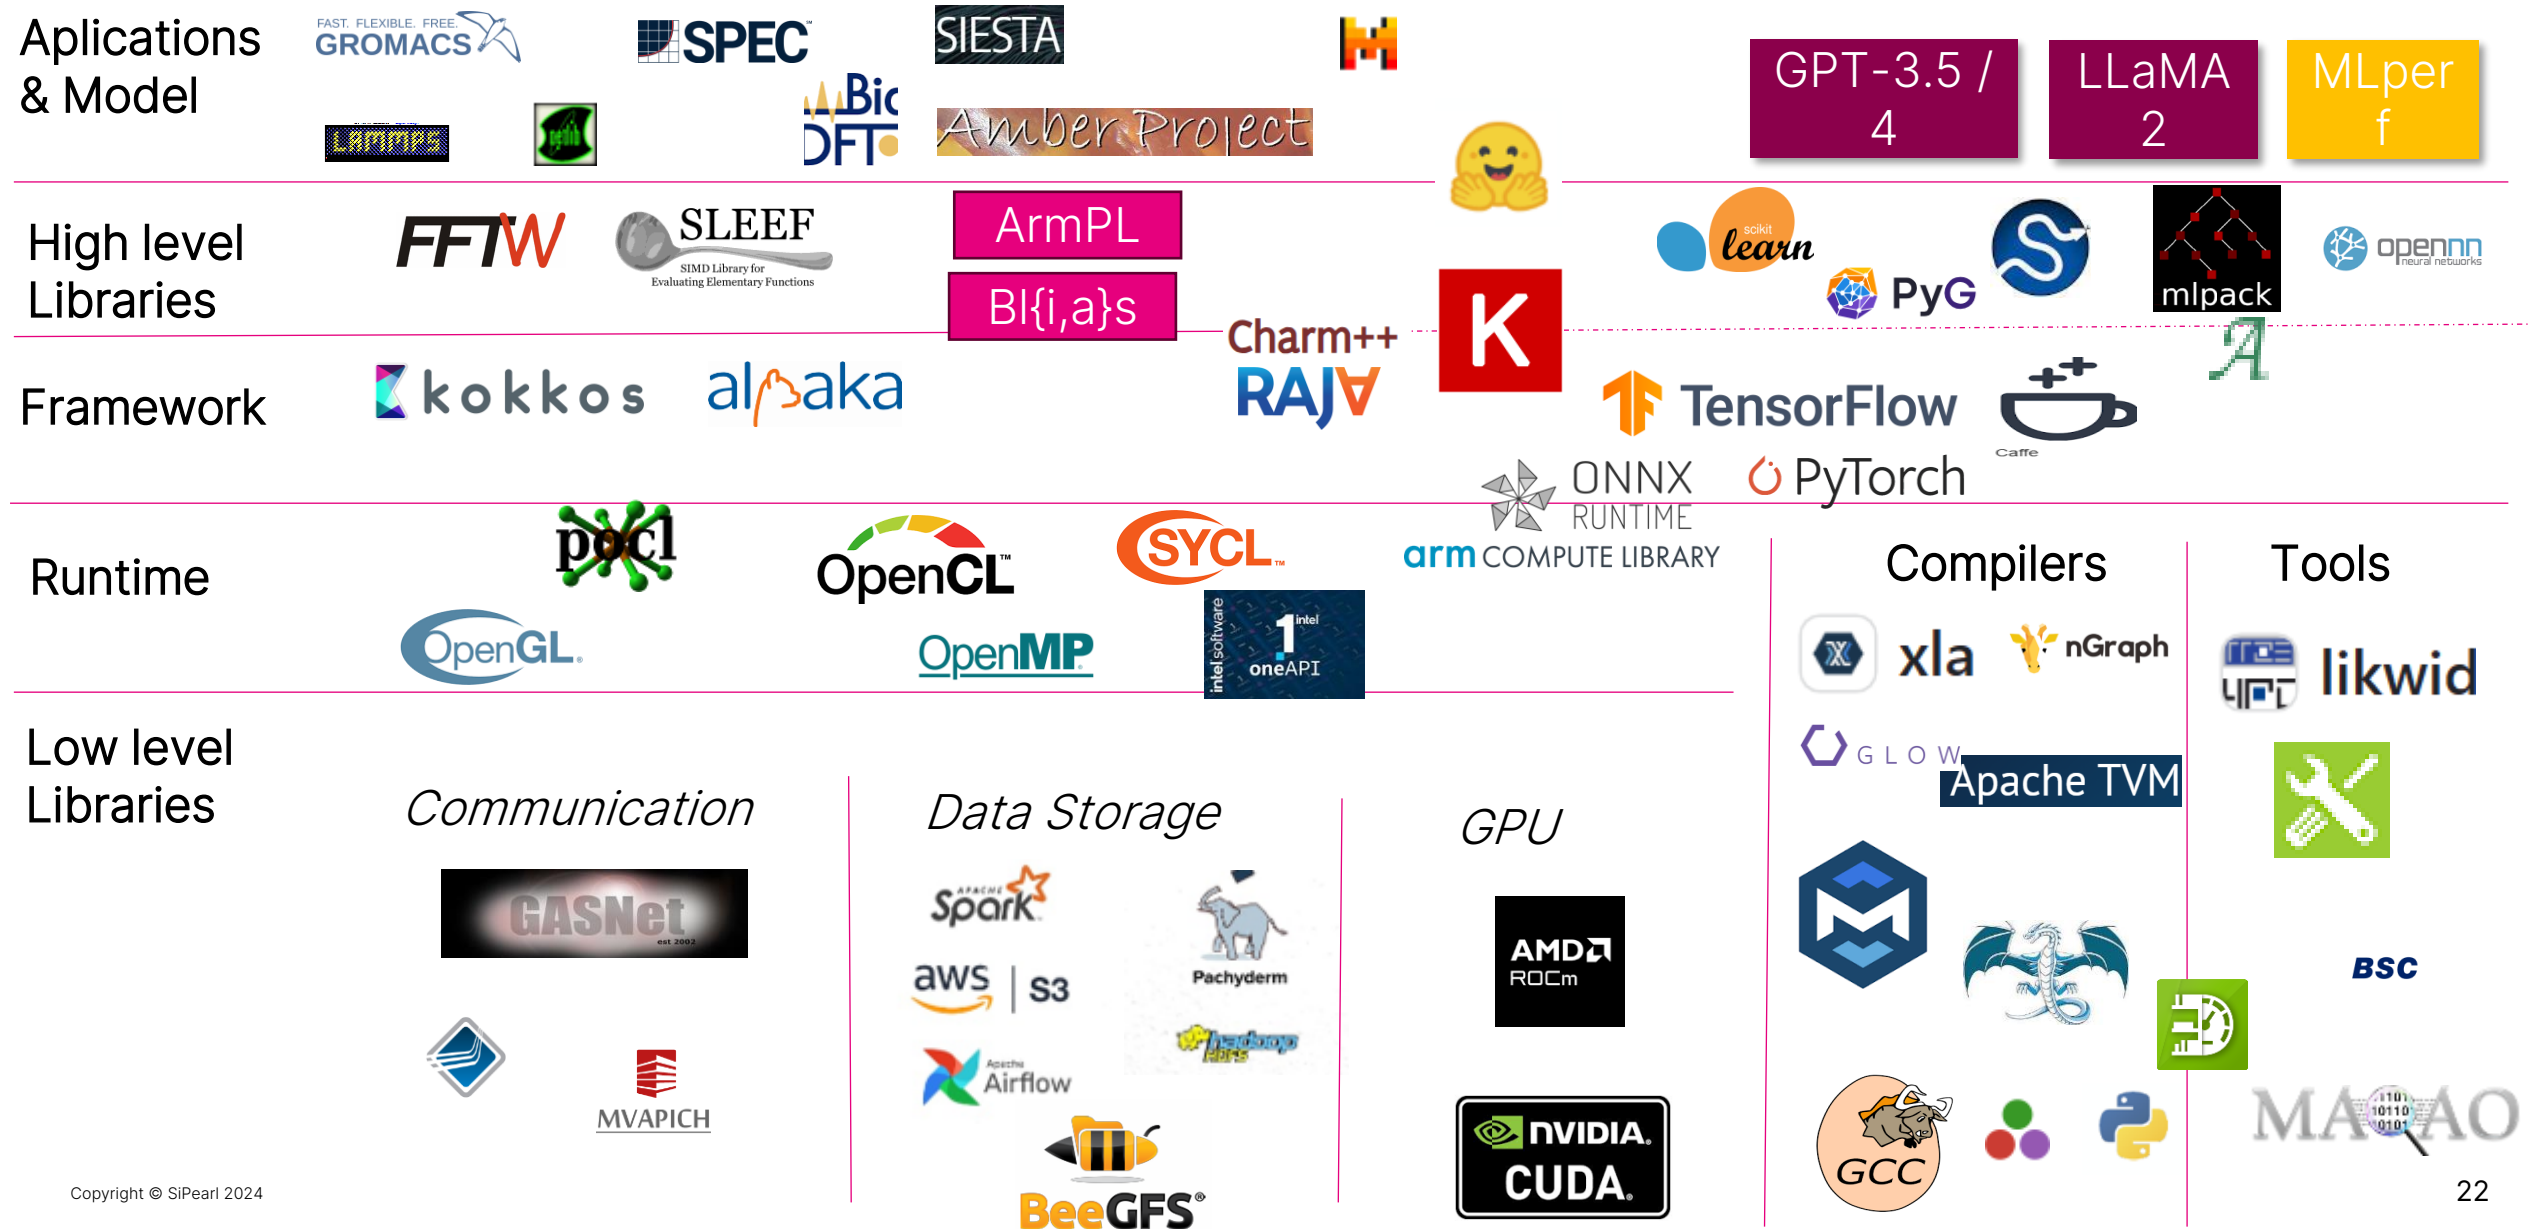
Porting Code for Arm-neoverse

Namely, for Rhea

Bird view of the HPC/AI Stack description



HPC/IA Stack in details



Compilers, how do I call them?



Open-source compiler		Proprietary compilers	
gcc	clang	armclang	nvc
g++	clang++	armclang++	nvc++
gfortran	flang(-new?)	armflang	nvfortran

An ecosystem composed by a variety of actors : ARM, NVIDIA, AWS, Google, Microsoft, SiPearl, ...

-mcpu=native
or **neoverse-v1**, neoverse-512tvb

Target the detected or specific CPU
(the default mcpu is more generic/slower)

Performance Libraries

Optimized BLAS, LAPACK and FFT for HPC applications



What is a Performance Library

- Optimized BLAS, LAPACK, FFT and math.h routines
- FFTW compatible interface for FFT routines
- Sparse linear algebra and batched BLAS support
- Tuned for Arm Neoverse family of processors

Several competitive Libraries

- ARMPL
- NVPL
- BLAS
- BLIS
- FFTW,
- PETSc
- ..

FLAG	DESCRIPTION
-I\$ARMPL_DIR/include	Include directory
-L\$ARMPL_DIR/lib -larmpl_mp	Link against ArmPL Here OpenMP version (_mp)
-L\$ARMPL_DIR/lib -lamath	Link against libamath
(-lgfortran) -lm	Always finish by linking with -lm

Supported by a wide community



SIPEARL

Profiling

Mandatory to take the best of
your applications


Profiling Tools for Rhea

Name	Typical Use	Typical Scale	Languages	Metrics
Linaro MAP	Initial high-level profile to identify hot spots	Single process to thousands of processes	C, C++, Fortran, Python	Wallclock time, hardware perf, MPI, OpenMP, CUDA, custom
perf	Quick detailed profile of a single process	Single process	Anything, usually compiled	Wallclock time, hardware perf
perf-lib-tools	Math library call profile	Single process	C, C++, Fortran	Count BLAS/LAPACK/FFT calls, call parameters
Maqao	Wide profiling, notably for vectorization	Single node	C,C++,Fortran, Python	Multiple metrics
Carm-roofline	Quick view of performance vs max perf.	Single process/node	C, C++, Fortran	Roofline

Multiple other tools exist HPCToolkit, TAU, ScoreP, Scalasca, LIKWID, CALIPER, ERT, Callgrind, Dimemas, DiscoPoP, Extrae, PAPI, Paraver,...

Examples of profiling

Loop id	Source Location	Source Function	Coverage run_0 (%)	Vectorization Ratio (%)	Vector Length Use (%)
2186	ecrad_ifs_blocked - radiation_ifs_rrtm.F90:594-595	radiation_ifs_rrtm_gas_optics	5.28	28.57	32.14
3717	ecrad_ifs_blocked - random_numbers_mix.F90:185-192	random_numbers_mix_initialize_random_numbers	3.68	0	14.58
3500	ecrad_ifs_blocked - radiation_two_stream.F90:655-695 [...]	radiation_two_stream_calc_ref_trans_sw	3.53	92.65	100
2123	ecrad_ifs_blocked - radiation_aerosol_optics.F90:648-656	radiation_aerosol_optics_add_aerosol_optics	2.42	0	25
3517	ecrad_ifs_blocked - radiation_adding_ica_sw.F90:111-116	radiation_adding_ica_sw_adding_ica_sw	2.38	90.48	96.43
3503	ecrad_ifs_blocked - radiation_two_stream.F90:630-647 [...]	radiation_two_stream_calc_ref_trans_sw	2.36	96.3	100
3507	ecrad_ifs_blocked - radiation_adding_ica_sw.F90:137-144	radiation_adding_ica_sw_adding_ica_sw	2.11	100	100
2514	ecrad_ifs_blocked - radiation_mcica_lw.F90:313-316	radiation_mcica_lw_solver_mcica_lw	2.11	88.89	91.67
2120	ecrad_ifs_blocked - radiation_aerosol_optics.F90:688-696	radiation_aerosol_optics_add_aerosol_optics	2.09	0	25
2102	ecrad_ifs_blocked - radiation_aerosol_optics.F90:756-770	radiation_aerosol_optics_add_aerosol_optics	2.09	0	25
2439	ecrad_ifs_blocked - radiation_mcica_sw.F90:290-294	radiation_mcica_sw_solver_mcica_sw	1.71	90.91	93.18
3866	ecrad_ifs_blocked - srtm_gas_optical_depth.F90:317-320	srtm_gas_optical_depth	1.59	0	25
2121	ecrad_ifs_blocked - radiation_aerosol_optics.F90:713-717	radiation_aerosol_optics_add_aerosol_optics	1.57	0	25
3496	ecrad_ifs_blocked - radiation_two_stream.F90:385-400	radiation_two_stream_calc_no_scattering_transmittance_lw	1.5	100	100
2125	ecrad_ifs_blocked - radiation_aerosol_optics.F90:673-677	radiation_aerosol_optics_add_aerosol_optics	1.3	0	25
3791	ecrad_ifs_blocked - rrtm_gas_optical_depth.F90:180-180	rrtm_gas_optical_depth	1.21	25	31.25
2479	ecrad_ifs_blocked - radiation_cloud_generator.F90:332-388 [...]	radiation_cloud_generator_generate_column_exp_ran	1	0	23.08
2235	ecrad_ifs_blocked - radiation_ifs_rrtm.F90:509-509	radiation_ifs_rrtm_gas_optics	0.92	100	100





SIPEARL

Co-design with Rheia

... Lead by collaborative projects

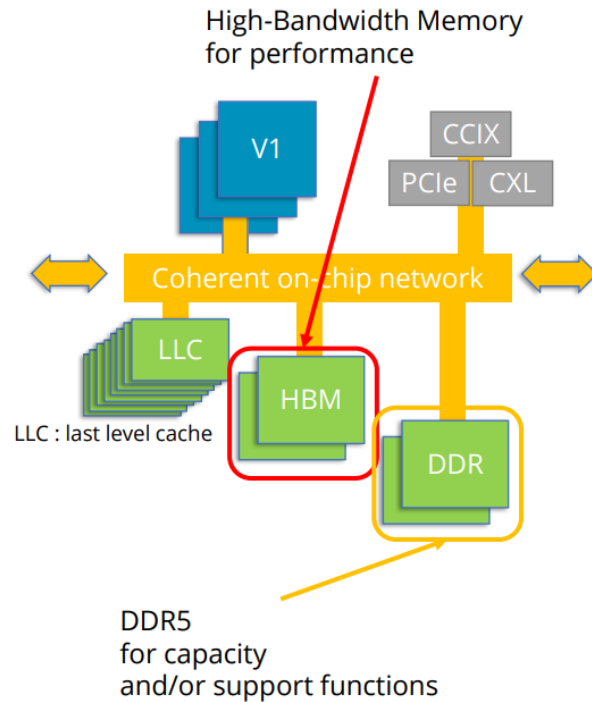
Challenges with Memory Management

Identify « hot » data

- Tools for identifying pattern
- Tools for identifying memory consumption

Exploit data tiering

- Tools for automatic migration
- Rewrite applications to only fill HBM

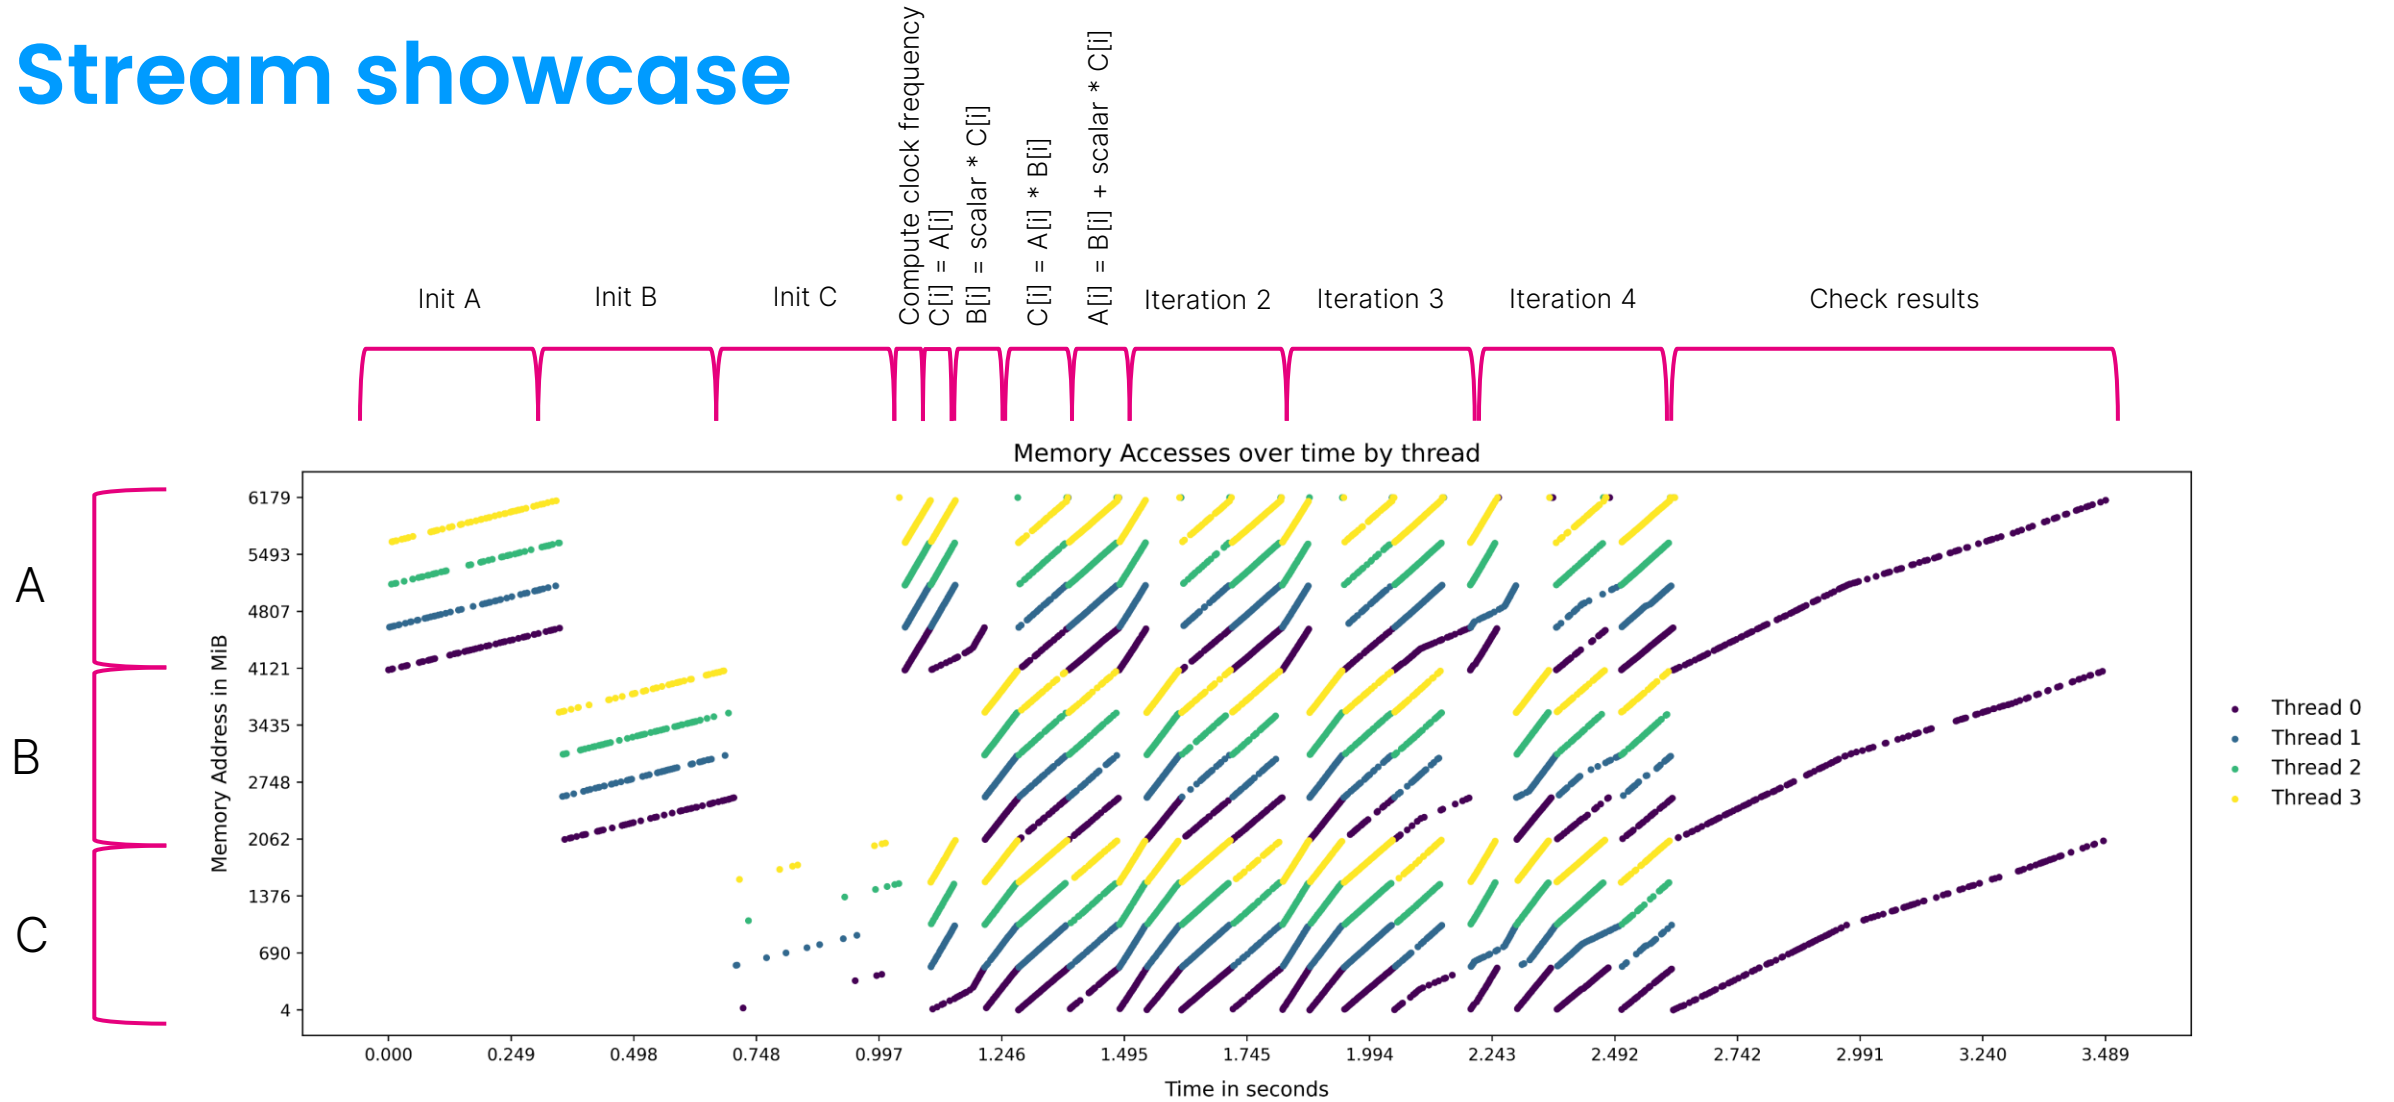


```
void compute(...)  
{  
    #memory-alloc(HBM)  
    int *a = malloc (...)  
  
    #memory-alloc(DDR)  
    int *b = malloc (...)  
  
    int *c = malloc (...)  
    int *d = malloc (...)
```

```
// Depending on runtime parameters will migrate dynamically data  
//     if a,b,c fit in HBM migrate everyone  
//     Otherwise migrate c and potentially a or b  
# memory-schedule {a; b; c}  
for (...)  
{  
    c[i] = a[i]*c[i-1] + b[i]*c[i+1];  
}
```

```
}
```

Stream showcase



Challenges with Vectorization

VLA Programming

- Rewrite application/Libraries to be vector-size agnostic
- Vector Length (VL) is a hardware implementation choice from 128 up to 2048 bits
- Predicate-driven loop control and management
- Extended floating-point horizontal reductions:
- SVE ACLE

VLA

Vector Length Agnostic
programming model

Write once

Compile once

Vectorize more loops

```
void example01(int *restrict a, const int *b, const int *c, long N)
{
    long i = 0;
    //Pseudo code of vectorization
    auto predicate = /*...*/
    for (; i < N; i += 4)
    {
        <predicate[0]>? a[i]   = b[i]   + c[i];
        <predicate[1]>? a[i+1] = b[i+1] + c[i+1];
        <predicate[2]>? a[i+2] = b[i+2] + c[i+2];
        <predicate[3]>? a[i+3] = b[i+3] + c[i+3];
    }
}
```

Plateform Seine DEMO



SiPearl involved in core European projects to ensure sovereignty

Cloud



Developing an open-source software ecosystem needed to optimize the efficiency of EPI hardware and facilitate the integration of SiPearl's microprocessors in the cloud.



Developing the 1st all-European RISC-V cloud server infrastructure, significantly enhancing Europe's open strategic autonomy.



Developing a custom cloud installation with the guarantee that an entirely European solution can be deployed reproducibly.

Centres of Excellence



Making some of the most used HPC application suites in engineering and manufacturing work on exascale EuroHPC supercomputers based on SiPearl's microprocessors.



Developing materials modelling, simulations and discovery technologies, and making them accessible to a vast community of researchers.



Promoting scientific and technological progress in key areas such as magnetic confinement fusion, industrial plasmas, medical applications...

And also regional projects:
Emopass (France), FlexFMM (Germany)

AERO – Two interesting pilots

- **High-Performance Algorithms for Space Exploration (UNIGE)**, complex scientific computing use case by the Gaia project and ESA in two fronts
- harness the capabilities of hardware acceleration in the EU cloud (using EPI technologies) especially with the integration of its existing codebases (SOS characterization) with TornadoVM and the rest of the acceleration services that will be developed within **AERO**
- demonstrate how the EU cloud with its hardware acceleration capabilities can serve the EU scientific community towards faster data processing and delivery.
- **HPC/Cloud Database Acceleration for Scientific Computing (SED)**
- assist large scale data processing infrastructure providers, such as SEDNAI, to innovate towards achieving higher performance via heterogeneous hardware acceleration in the context of the heterogeneous EU processor and cloud (using EPI technologies).
- demonstrate how the EU cloud can be utilized by HPC/Cloud technology providers to encourage migration and rapid adoption by developing a rich and diverse software stack that can be used out-of-the-box.

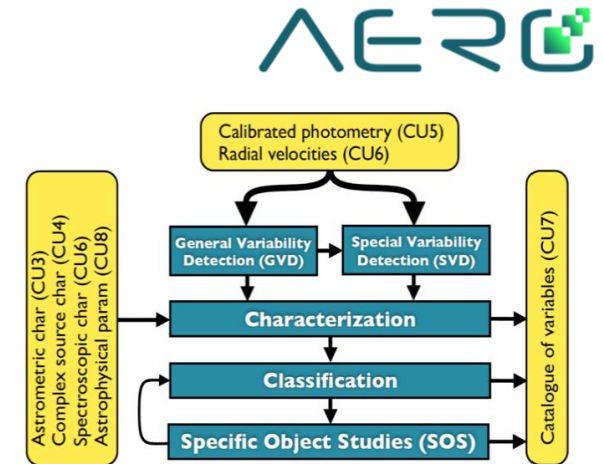


Figure 8: Data processing flow of CU7 of GAIA.

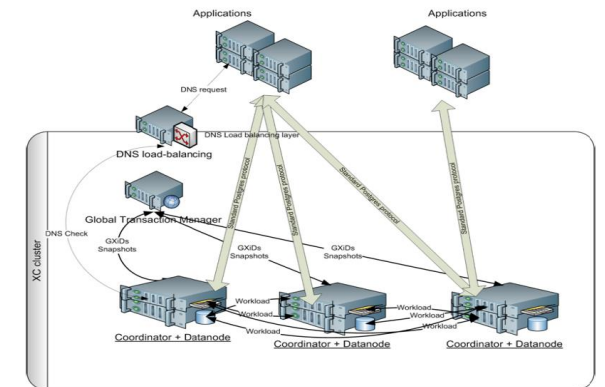


Figure 9: GAIA data processing architecture.

Main objective is to port and optimize those workloads on Rhea

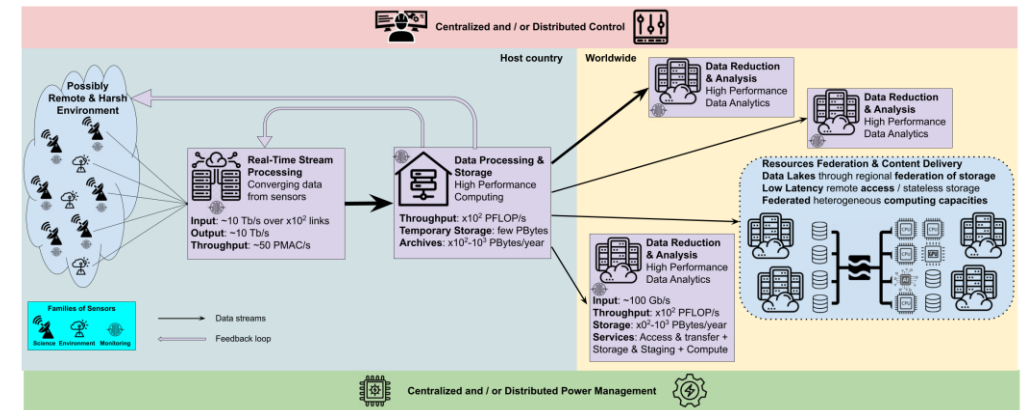
Next Project

ODISSEE – Start in 2025

- Deploy an HW platform based on Rhea and experiment several optimisation (energy efficiency by adding some AI on power management)
- 10+ partners including Observatoire de Paris, CNRS, INRIA, BsC, GENCI, SKA, etc.

STREAMS – Proposal done waiting for the answer

- Objective is to use Rhea platform for part of the processing and optimize for real time processing
- Partnership between Observatoire de Paris, SiPearl , CNRS



Main objective are to demonstrate Rhea capabilities on real use cases at scale