oneAPI Dev Summit @ISC22

Powering heterogeneous computing with oneAPI on ARM and Ponte Vecchio

Vincent Casillas, SiPearl, VP R&D

Gilles Civario, Intel, HPC software architect

intel®

SIPEARL

Introduction

Vincent Casillas – vincent.casillas@sipearl.com

VP R&D at SiPearl, based in France.

20+ years in software design, working closely with hardware architecture teams in different domains at Valeo, STMicroelectronics and on behalf of Ateme major customers.

Gilles Civario – gilles.civario@intel.com

Senior HPC Software Architect for Exascale at Intel, based in France.

20+ years in HPC with positions held at Dell, ICHEC, Bull and the CEA.

Special interest in code parallelization and optimization

Agenda

Introducing SiPearl

The market gap

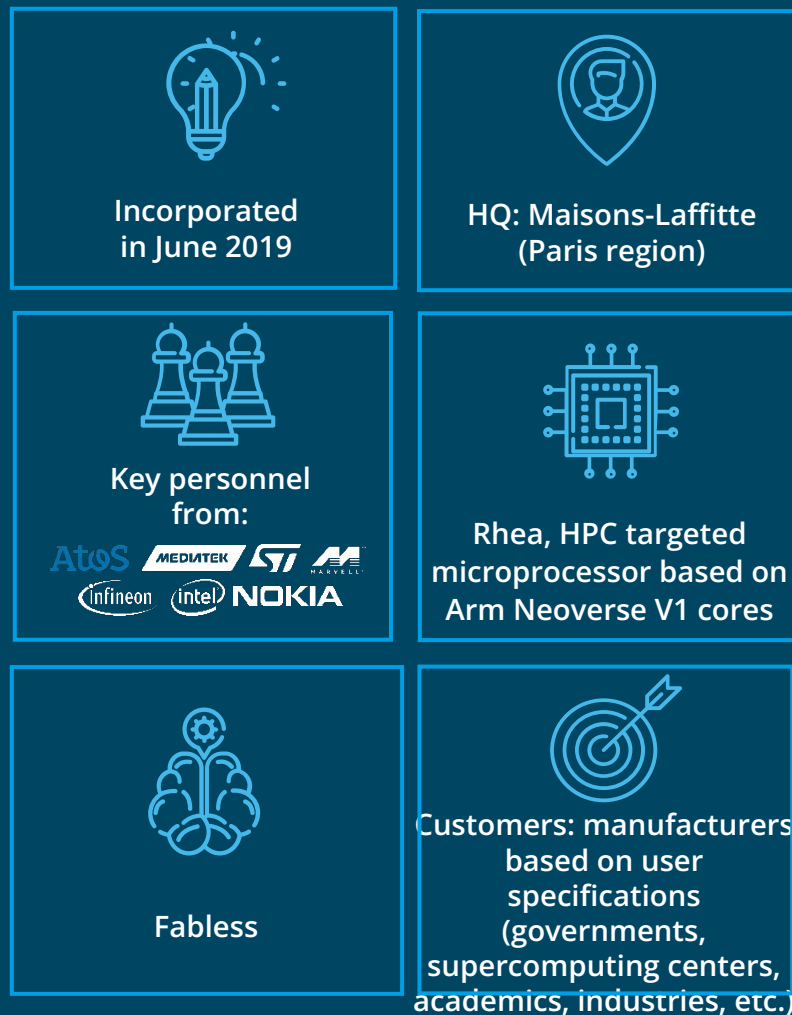Why Ponte Vecchio and oneAPI

Porting the oneAPI stack

Conclusion

SiPearl in a nutshell

SiPearl, the French company designing and bringing to market the high-performance, low-power microprocessor for European supercomputers.
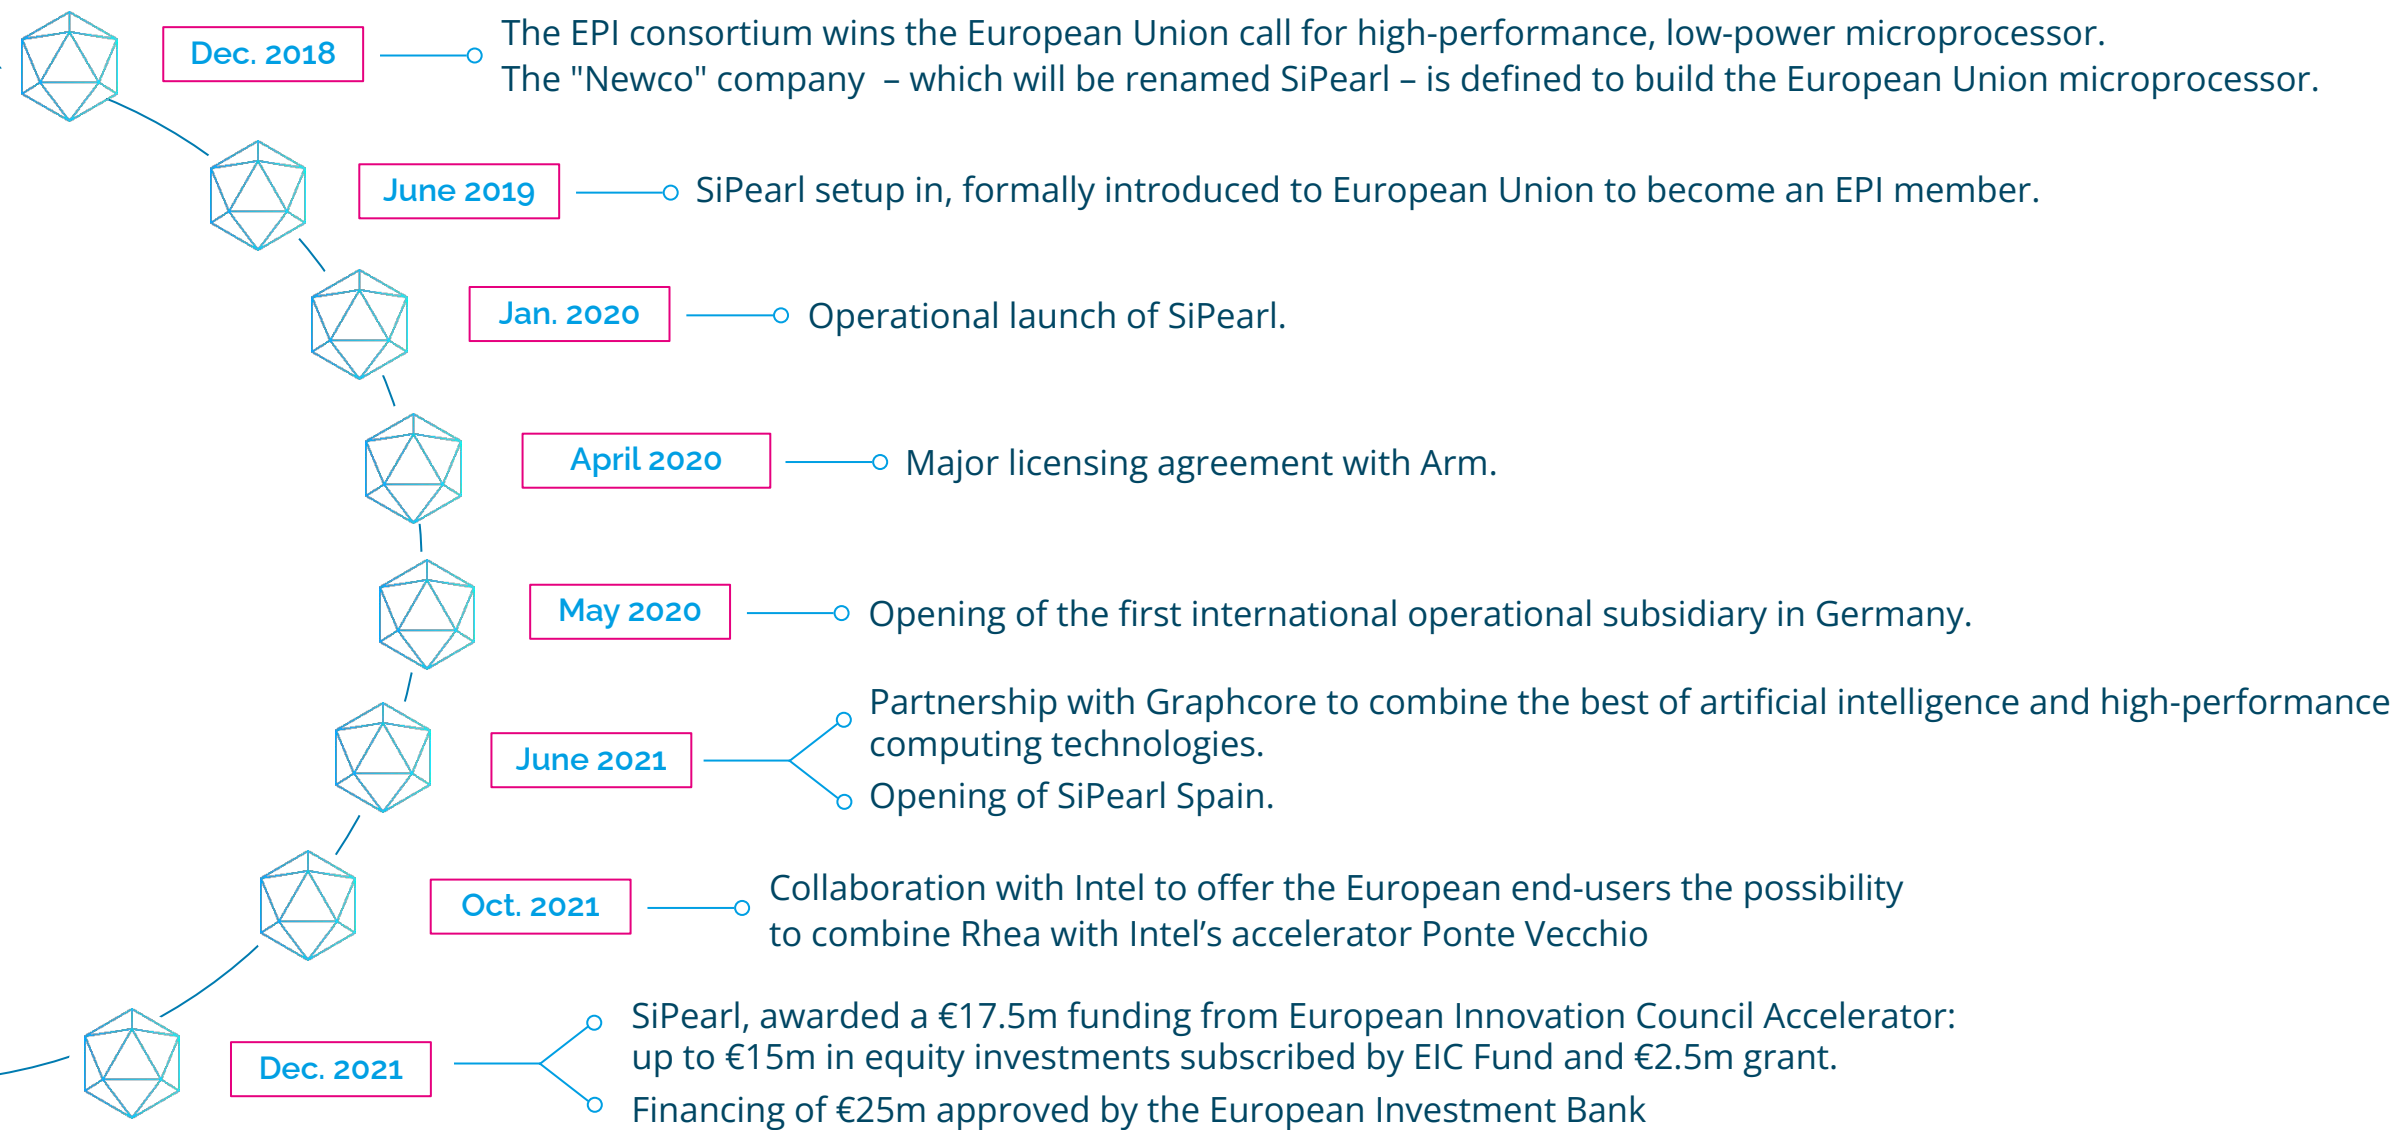
Incorporated in June 2019

HQ: Maisons-Laffitte (Paris region)

Key personnel from:
AtoS MEDIATEK STI MARVELL
infineon intel NOKIA

Rhea, HPC targeted microprocessor based on Arm Neoverse V1 cores

Fabless

Customers: manufacturers based on user specifications (governments, supercomputing centers, academics, industries, etc.)

6 locations in Europe

Duisburg
Maisons-Laffitte
Massy
Barcelona
Sophia Antipolis
Grenoble

104 employees

Spain
Germany22
78

Timeline

Dec. 2018 — The EPI consortium wins the European Union call for high-performance, low-power microprocessor.
The "Newco" company – which will be renamed SiPearl – is defined to build the European Union microprocessor.

June 2019 — SiPearl setup in, formally introduced to European Union to become an EPI member.

Jan. 2020 — Operational launch of SiPearl.

April 2020 — Major licensing agreement with Arm.

May 2020 — Opening of the first international operational subsidiary in Germany.

June 2021 — Partnership with Graphcore to combine the best of artificial intelligence and high-performance computing technologies.
Opening of SiPearl Spain.

Oct. 2021 — Collaboration with Intel to offer the European end-users the possibility to combine Rhea with Intel's accelerator Ponte Vecchio

Dec. 2021 — SiPearl, awarded a €17.5m funding from European Innovation Council Accelerator: up to €15m in equity investments subscribed by EIC Fund and €2.5m grant.
Financing of €25m approved by the European Investment Bank

An intense recruitment strategy

Target: 1,000 employees in 2025

2 May 2022: milestone of 100 employees reached

**Recruiting on 6 locations
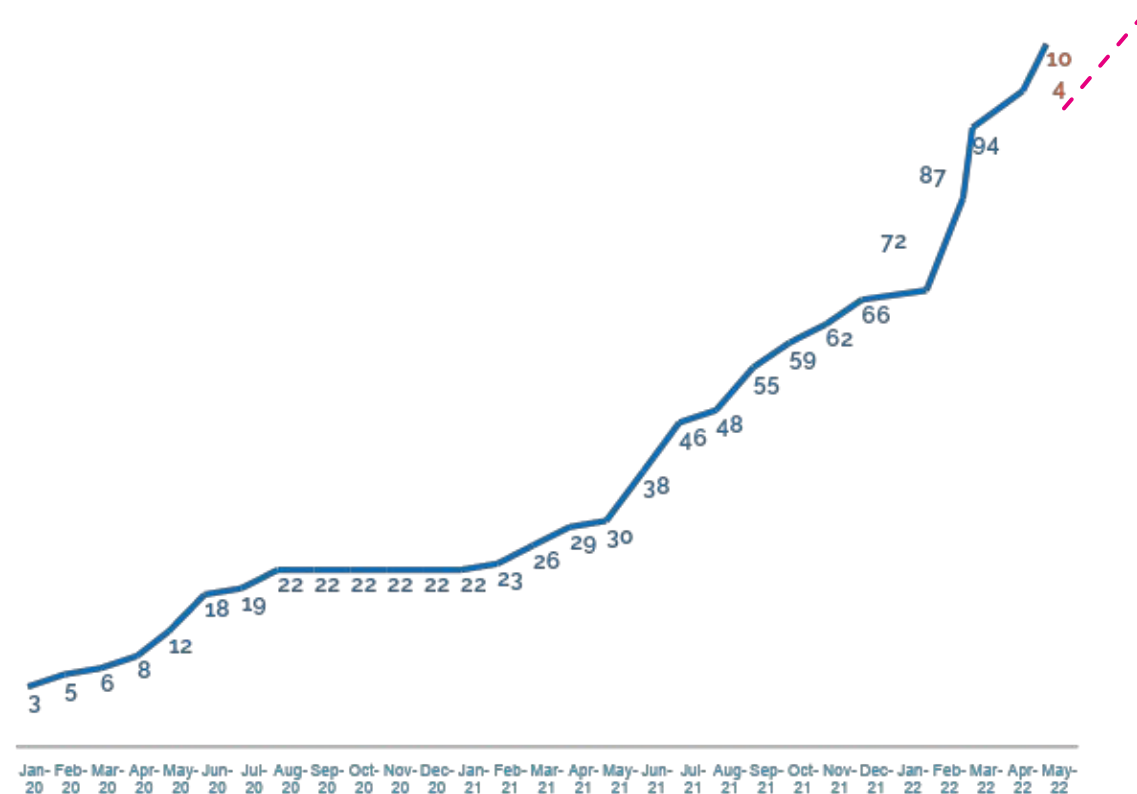in the heart of semiconductor and HPC skill pools**

Software

– Embedded
– Linux Kernel
– Compiler
– UEFI/BIOS
– HPC benchmark
– Software architecture
– Software development

Microelectronics

– RTL design
– SoC design
– Verification
– Virtual prototyping
– Logical synthesis
– PCB design
– DFT
– Layout

We are HIRING!

Evolution of the workforce since the operational launch

Rhea, the 1ˢᵗ generation European HPC microprocessor

**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 - Arm Virtualization Extensions
SoC	- Arm Mesh fabric - Advanced RAS support including Arm RAS extensions - Link protection for NOC and high-speed IO - ECC support for selected memory
Cache	- RAS supported for all Cache levels
Memory	ECC for memory and link protection for controllers - HBM2e - DDR-5
High Speed I/O	PCIe or CCIX/CXL - Root and endpoint support
Other I/O	USB, GPIO, SPI, I²C...
Power Management	Power management block to optimize perf/watt across use cases and workloads.
Security Block Support	- Secure boot and Secure upgrade - Crypto - True Random Number Generation

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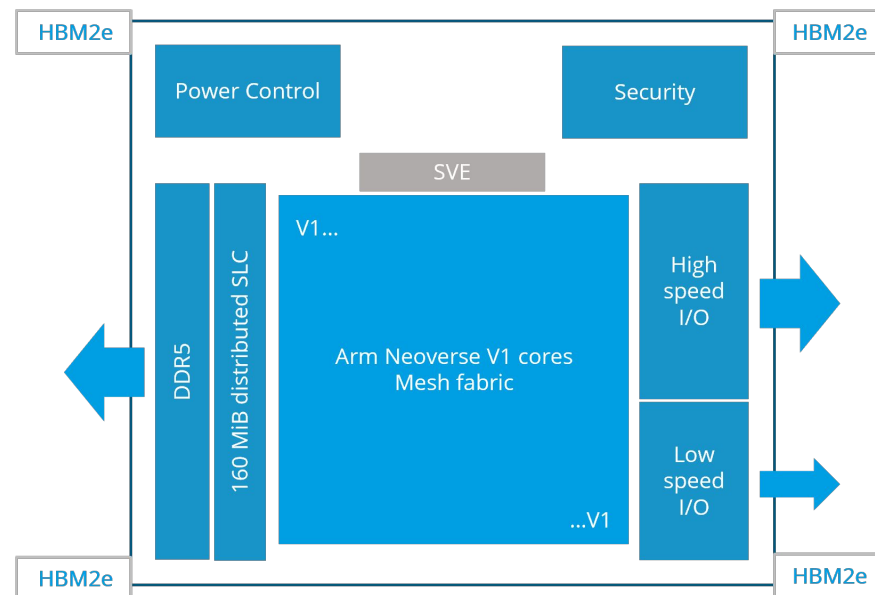
Rhea will deliver extraordinary compute performance and efficiency with an unmatched Bytes/Flops ratio.

SIPEARL

8

SiPearl corporate vision and strategy

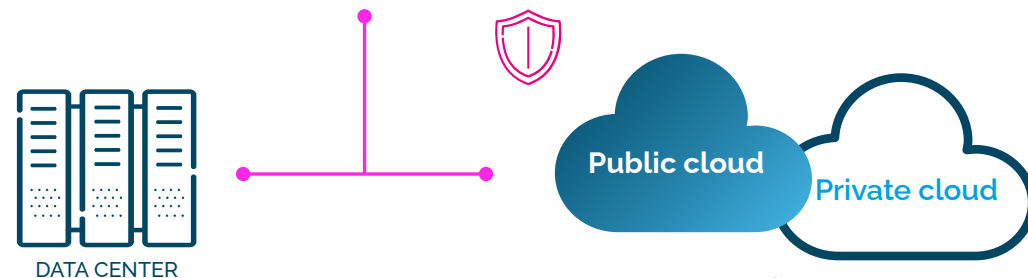## HPC Supercomputing
SiPearl entry business: European HPC

2023

HPC-dedicated

Data Center-Central
Data centers, private and public cloud

2025

DATA CENTER

Public cloud

Private cloud

Embedded Security channel

Data Center-Edge
Smaller Data centers,
including around 6G infrastructure

2027+

6G
6G infrastructure

We are focused on HPC!

The market gap

Addressing the European Exascale

SiPearl's target is primarily the European market

- Technical sovereignty

- Aiming at proposing the community open alternatives to proprietary solutions

But the Rhea CPU cannot cover the full range of needs all by itself

- Need of an accelerator for reaching the Exascale level

- Need of an open programming model for exploiting it

 Intel Ponte Vecchio and the oneAPI open programming model

Targeting Intel Ponte Vecchio

Ponte Vecchio
SOC

>100 Billion Transistors
47 Active Tiles
5 Process Nodes

Compute Tile
Rambo Tile
Foveros
Base Tile
HBM Tile
Xᵉ Link Tile
Multi Tile Package
EMIB Tile

Ponte Vecchio
Execution Progress

A0 Silicon Current Status

> 45 TFLOPS — FP32 Throughput

> 5 TBps — Memory Fabric Bandwidth

> 2 TBps — Connectivity Bandwidth

Currently under deployment at Argonne National
Laboratory for a 2+ EFlops in the Aurora machine

Understanding oneAPI

The two faces of oneAPI

An open cross-industry initiative

- Landing web page: https://oneAPI.io

- Open-source implementation for all components

  - Sources available on github
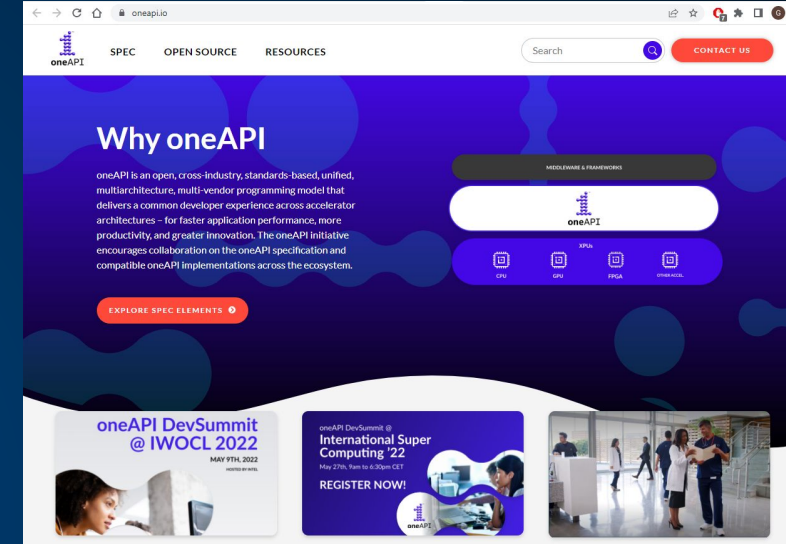
An Intel instantiation of the standard

- Landing web page: https://software.intel.com/oneapi

- Binary packages usable free of charge

- Include extra packages from what use to be Intel Parallel Studio
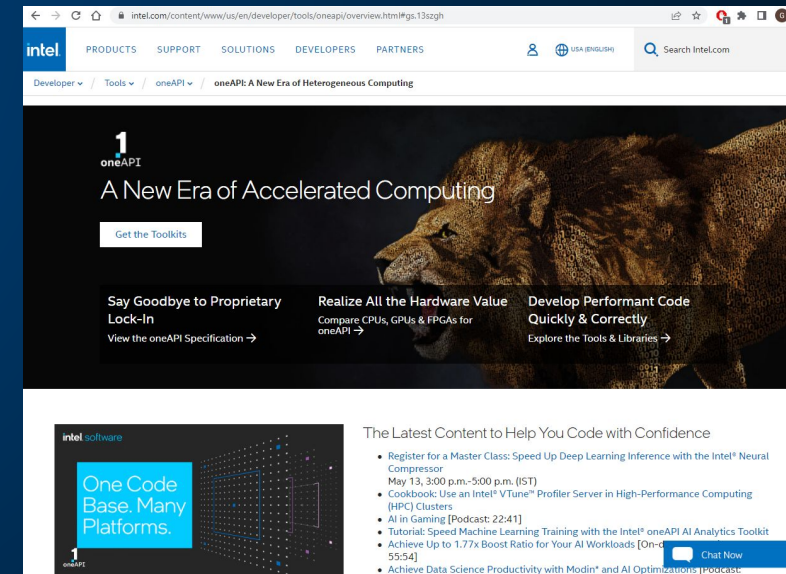
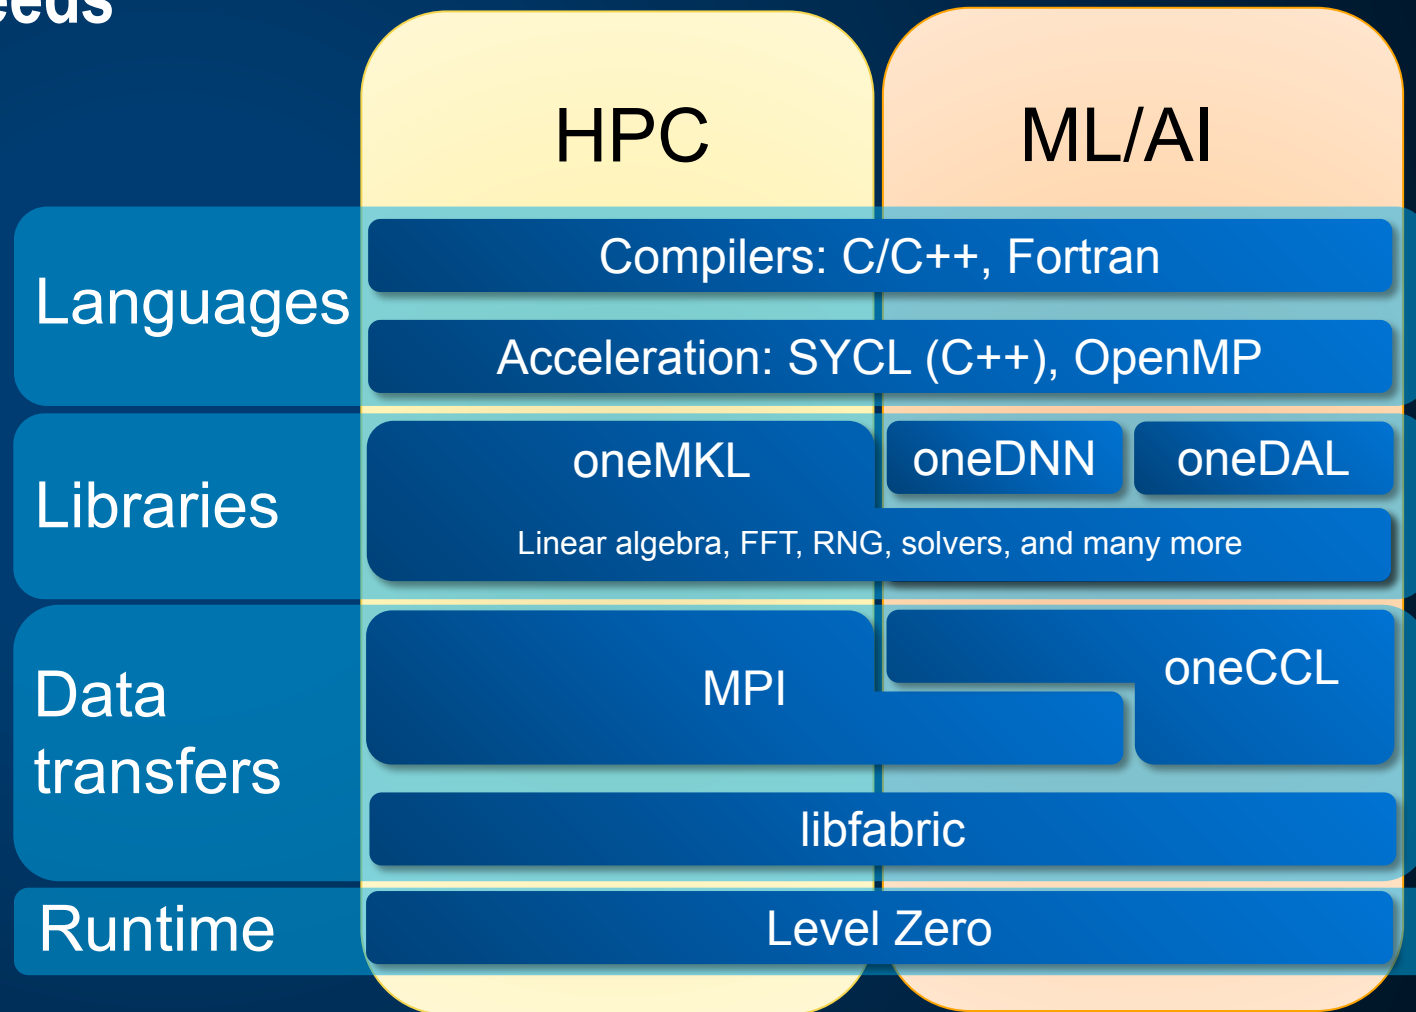 SiPearl leverages the best of both

A fully featured software distribution

Capitalizing on the best of both seeds

Compilers

- New Intel versions based on LLVM
- Same look and feel as on x86
- Support of the most recent standards
  - C 11, C++ 17, Fortran 2008*
  - OpenMP 5.1**, SYCL 2020**

Development tools

- CPU & GPU profilers
- CPU & GPU debugger

	HPC	ML/AI
Languages	Compilers: C/C++, Fortran	
	Acceleration: SYCL (C++), OpenMP	
Libraries	oneMKL	oneDNN / oneDAL
	Linear algebra, FFT, RNG, solvers, and many more	
Data transfers	MPI	oneCCL
	libfabric	
Runtime	Level Zero	

* Growing support of Fortran 2018   ** Some minor features still missing

Compilers porting strategy

Targeting SiPearl Rhea CPUs

- Ensuring full performance on CPU
- Leveraging ARM compilers and performance libraries
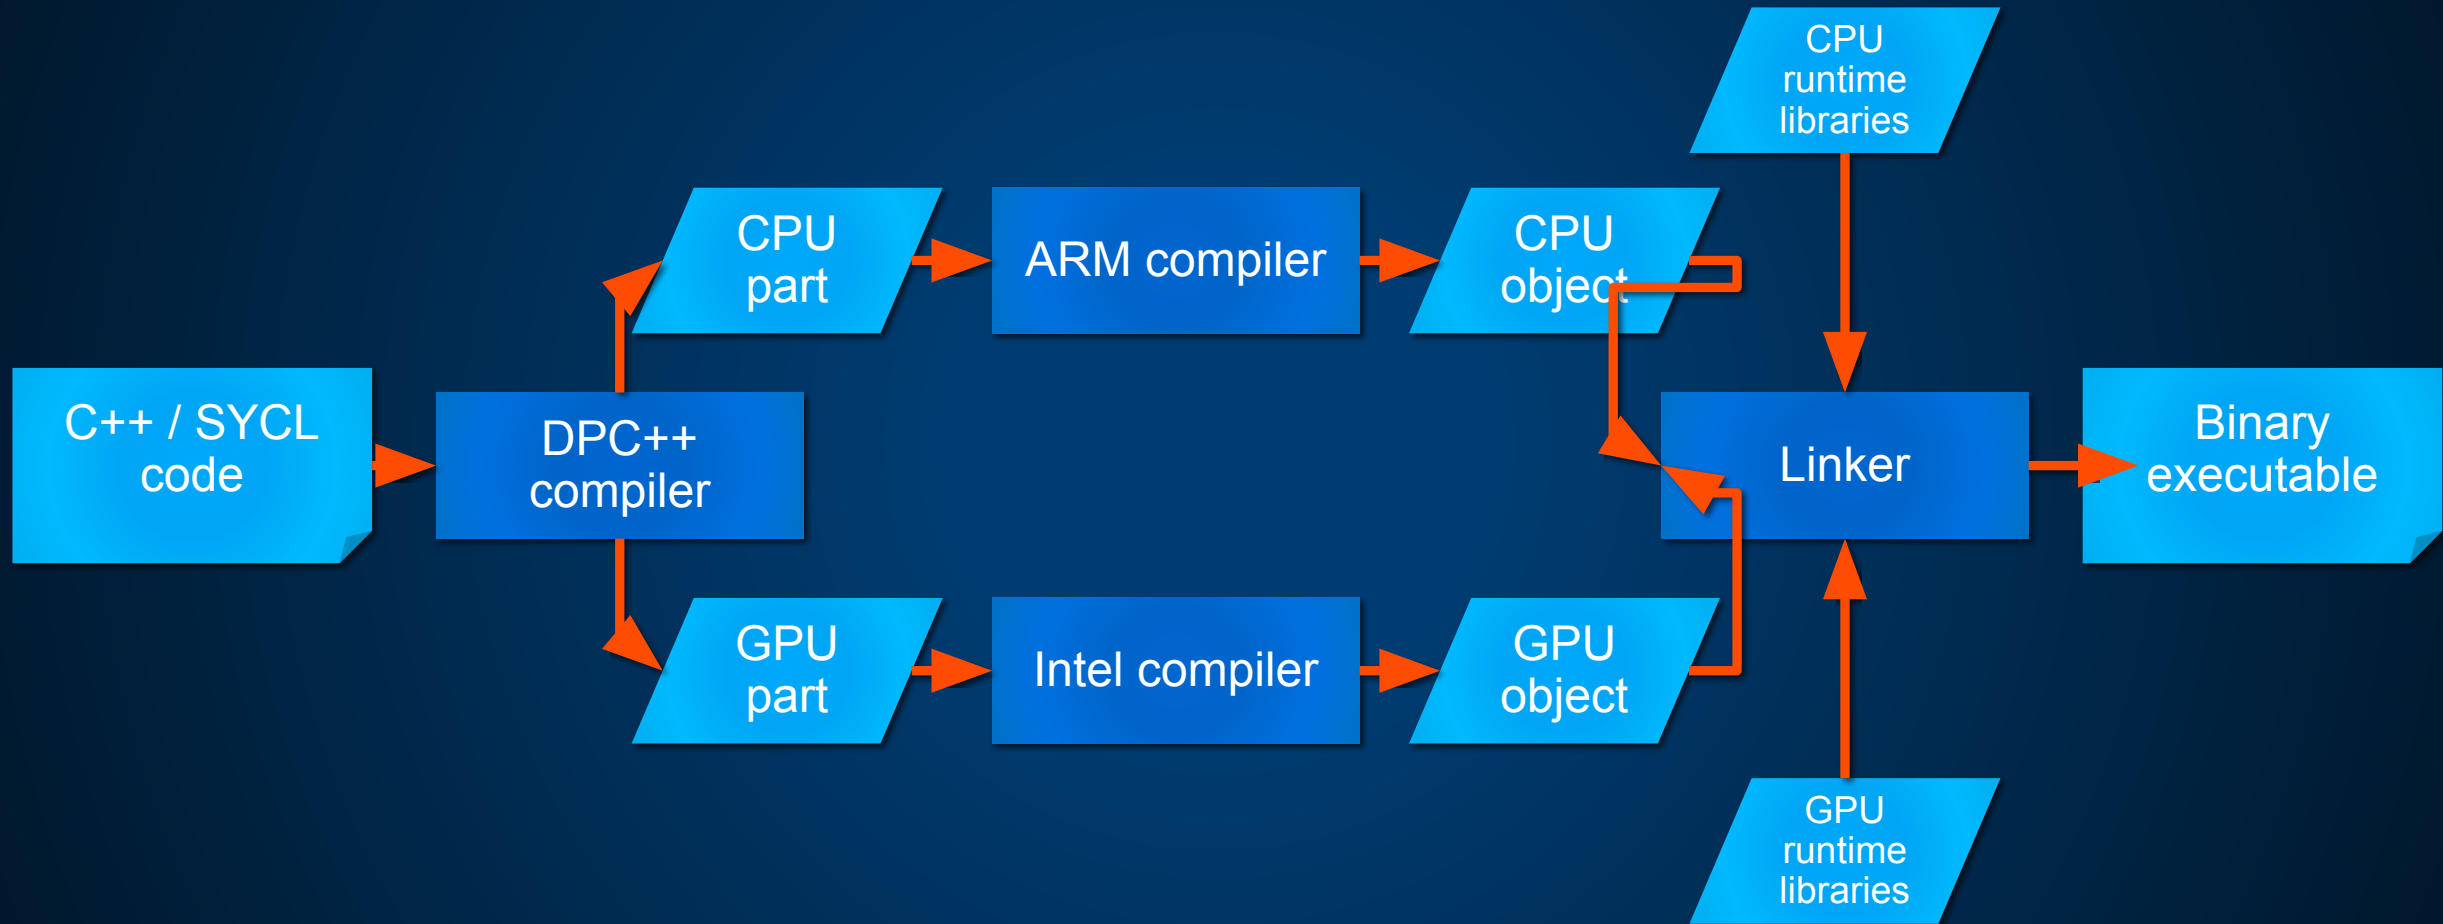
Targeting Intel Ponte Vecchio GPUs

- Porting Intel optimized sources
- Leveraging Intel's decades of expertise in compilers, libraries and tools

Combining the best of both worlds for maximizing performance

- Ability to integrate deeply both branches thanks to their common LLVM base
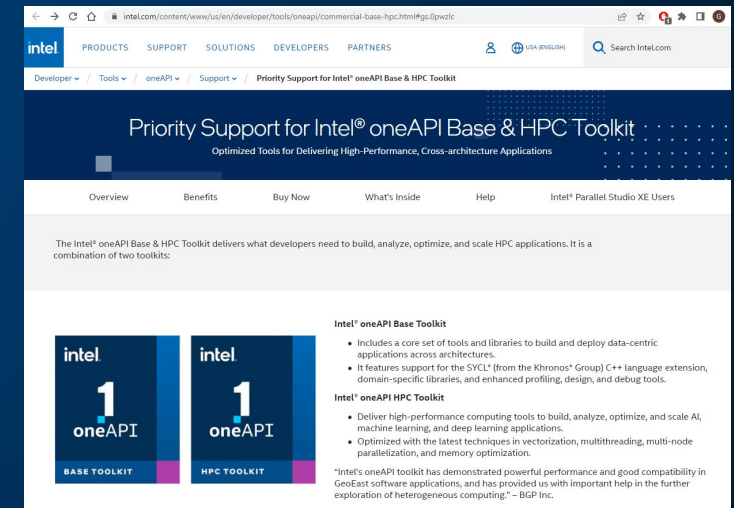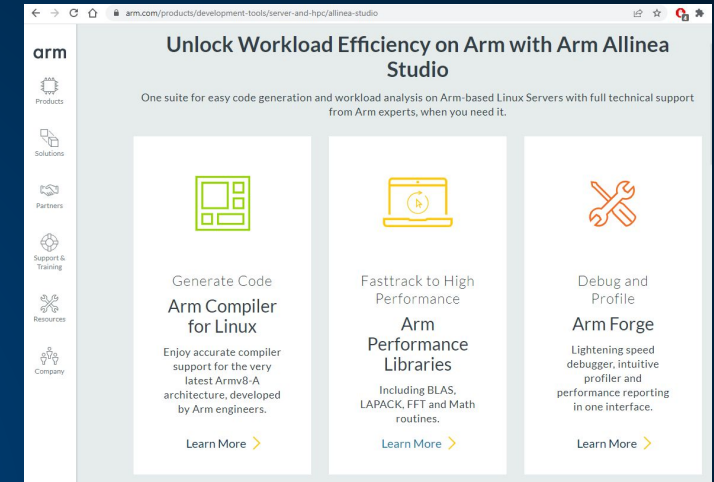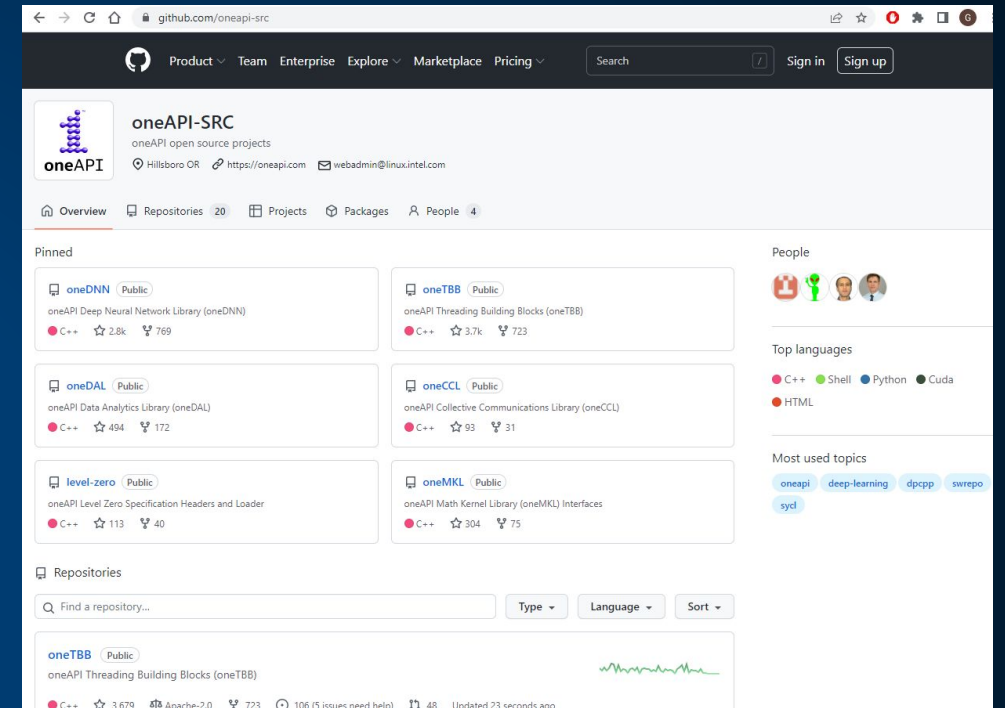
The compilation flowchart

Combining the best tools for each parts

Open-source packages

Working directly on the sources from github

- Porting and testing

- Optimizing when needed

- Finding and fixing bugs sometimes

- Contributing back as per the licenses

Making sure that all the expected components of the software stack are ported and fully functional on the Rhea + PVC solution

Debugging and profiling

Working on Intel sources for the tools

- Debugger: Intel distribution of gdb

  – Debugging capability for both the CPU and the GPU

  – All you can expect from a modern debugger

- Design assistant: Intel Advisor

  – Roofline analysis and offloading advisor

- Profiler: Intel VTune Profiler

  – Optimization of the offloading and tuning of the GPU code
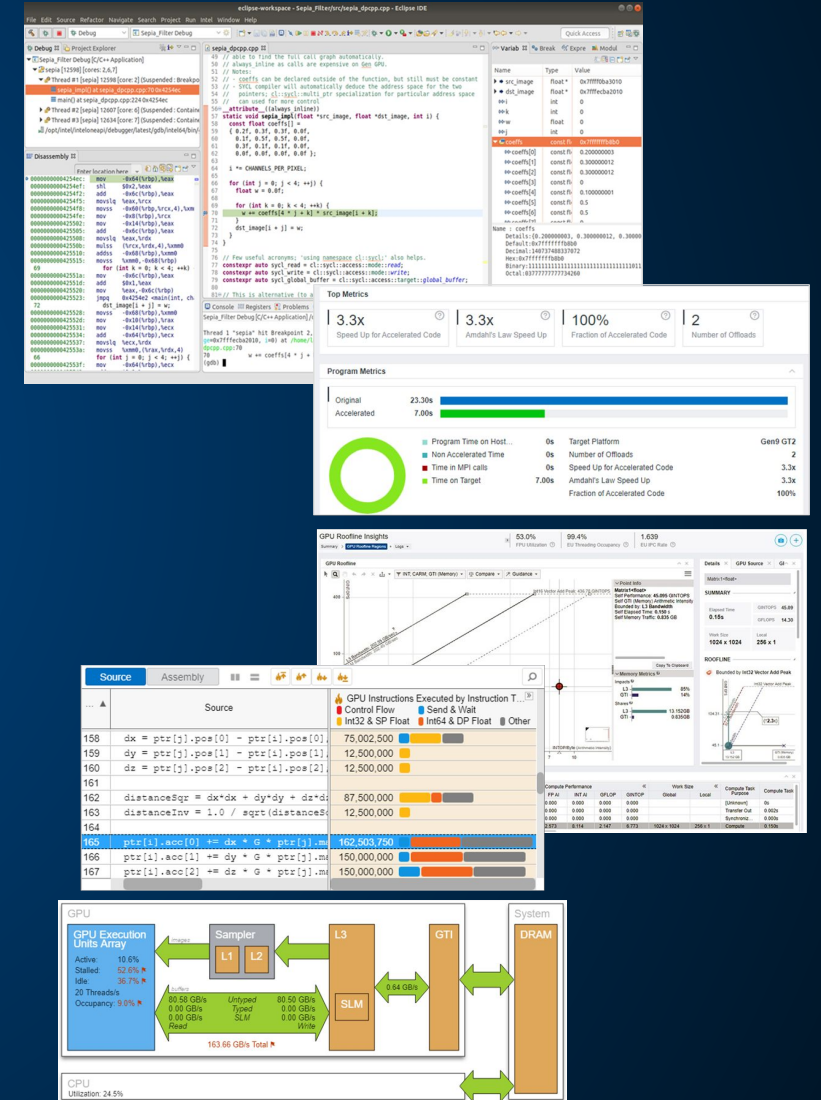
Collaborative work

Close collaboration between Intel and SiPearl

- Regular project meetings
- Full transparency of respective developments

Tracked advancement of the project

- Steady pace of progression
- Steady improvements in features and performance

We will be ready, will you?

Current status

Divers, runtime and low-level libraries

- Initial porting done

- Fully functional

- Performance optimizations as needed

Compilers

- Initial porting done

- Fully functionals

- Integration work in progress

Libraries

- Too many to detail them all

- Work progressing as expected

- Integration very promising
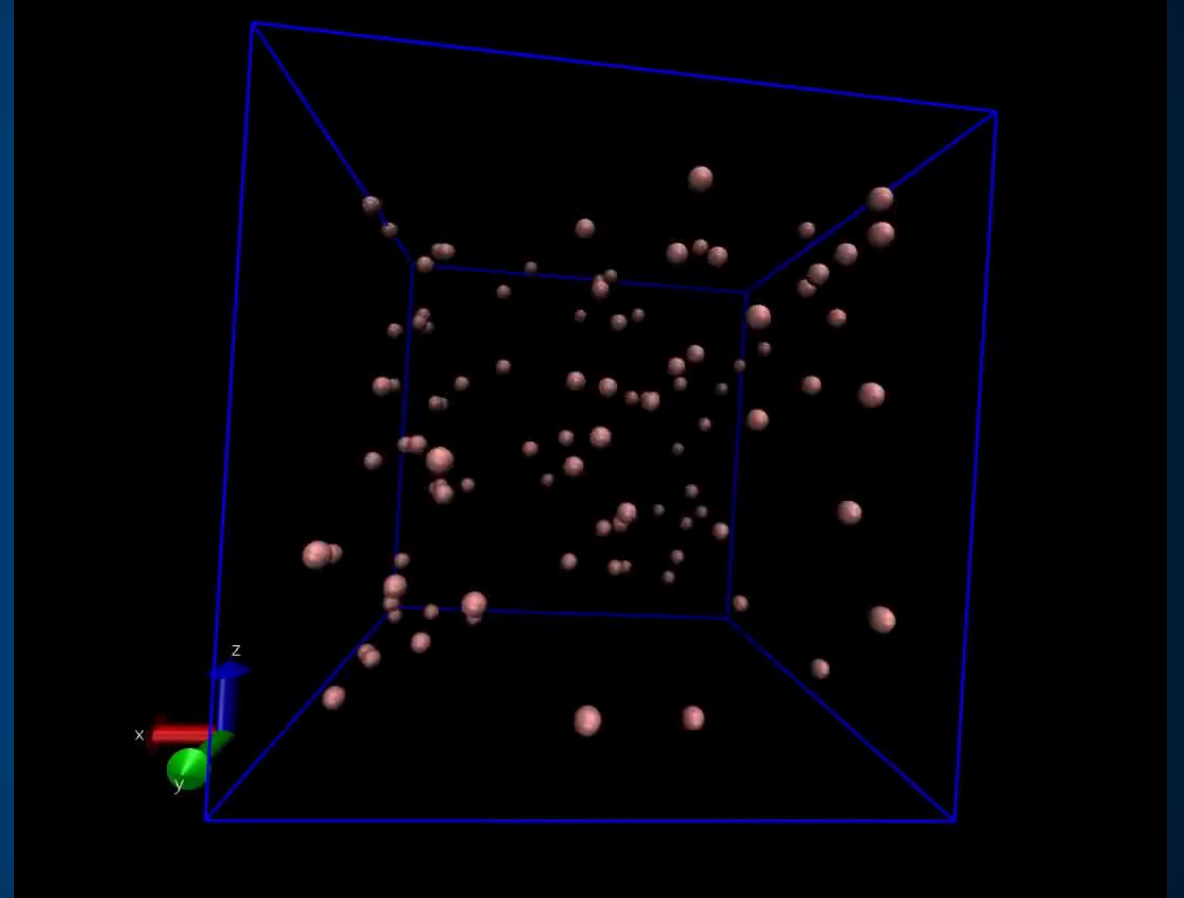
Development tools

- Debugger done

- VTune and Advisor next

It is working already

Compiled the latest version of GROMACS with SYCL acceleration enabled

Run on our ARM + PVC development platform

Outputs streamed and post-processed with VMD

Enormous "Thank you" to François Dugast for his contribution

Conclusion

oneAPI is the best way to take advantage of SiPearl + Intel solution

- Full native performance on SiPearl ARM CPU

- Full native performance on Intel X^e GPU

- Seamless integration

oneAPI is cross-platform portable

- No vendor ties

- Already available on many different architectures

Don't waste time to make the move

- We can help you

We will be ready for Exascale, will you?

Thank You

Questions ?