Supreme-K:
A Korean Initiative for Pioneering the AB21 HPC accelerator and Compute Node

* Supreme-K : National Supercomputer Development Project

2025. 8. 20

Yoomi Park (parkym@etri.re.kr) and Woojong Han(woojong.han@etri.re.kr)
ETRI

Contents

Contents

❖ (Survey '19, Korea) Not enough HPC resources vs. Growing demand for HPC resources

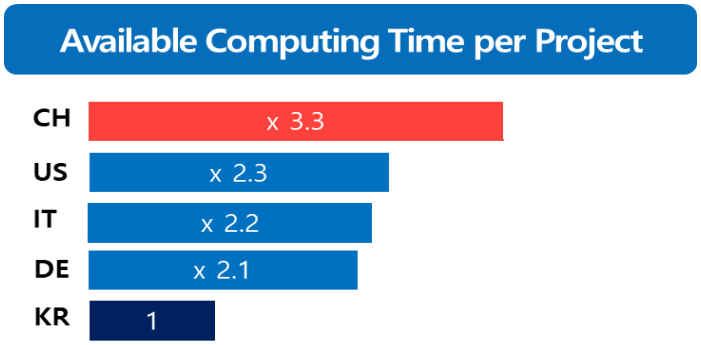

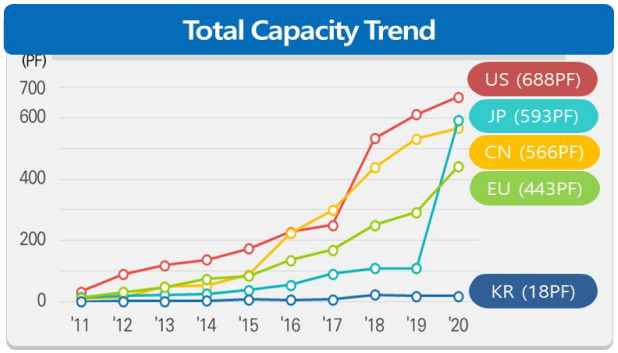

❖ (Challenges) supercomputers from global vendors (made in the US or China)

TOP500 rank	Name	Site	Vendor
25	Sejong	Naver	Nvidia
32	SSC-21	Samsung	HPE
44	kakaocloud	kakao	Supermicro
58	Guru	Korea Meteorological Administration	Lenovo
59	Maru		Lenovo
70	kakaocloud	kakao	Supermicro
73	Titan	SK Telecom	HPE
75	Nurion	KISTI	HPE
90	KT DGX SuperPOD	KT	Nvidia
148	DAIDC		HPE
178		CUBOX	HPE
275	Dream-AI	GIST	Inspur
468	SSC-21 Scalable Module	Samsung	HPE

❖ (Risks) Critical national resource → supply chains disruption → national risk

- **2nd Master Plan for National Supercomputer Development**
- **Government's Execution Strategy : 3-Pillars of undertaking**
  - to build Sustainable Top-tier HPC ecosystem
  - being realized through R&D projects by government

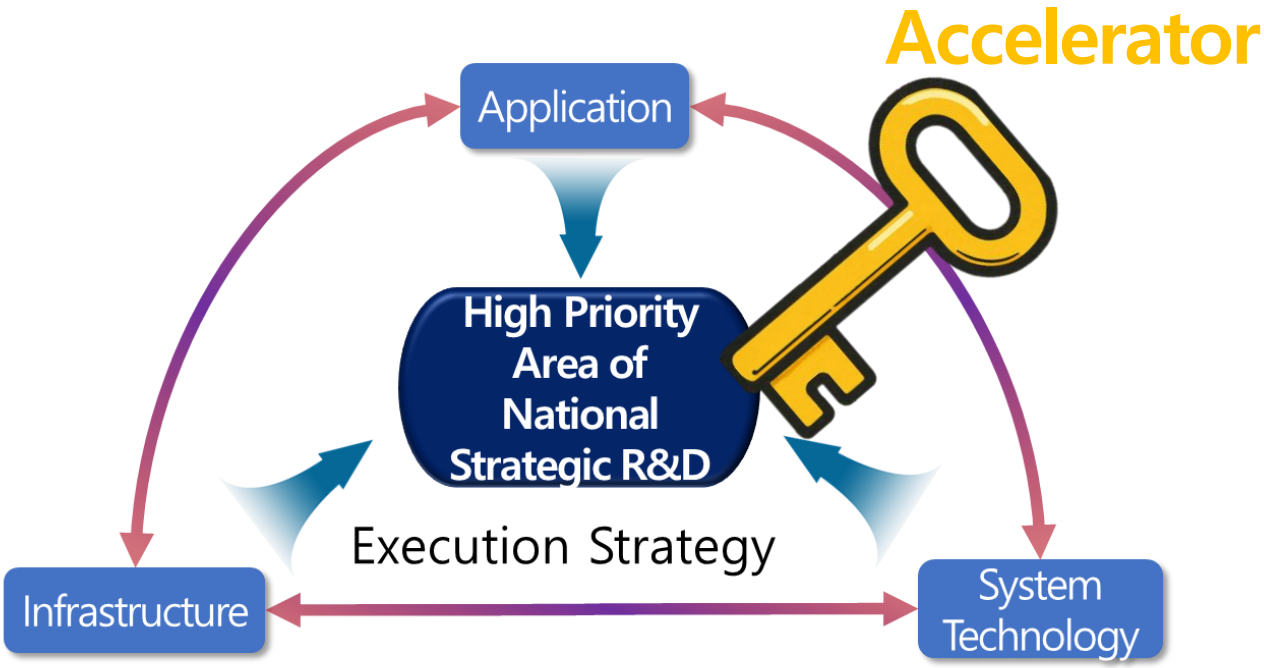

Contents

❖ **Goal** of Project 'Supreme-K'

- Development of a PoC of HPC system based on a domestic **HPC accelerator ('AB21'), system SWs, and compute nodes('Silver')**
- supported by Ministry of Science & ICT of Korea

❖ **Period** : Jul. 2020~Sep. 2025

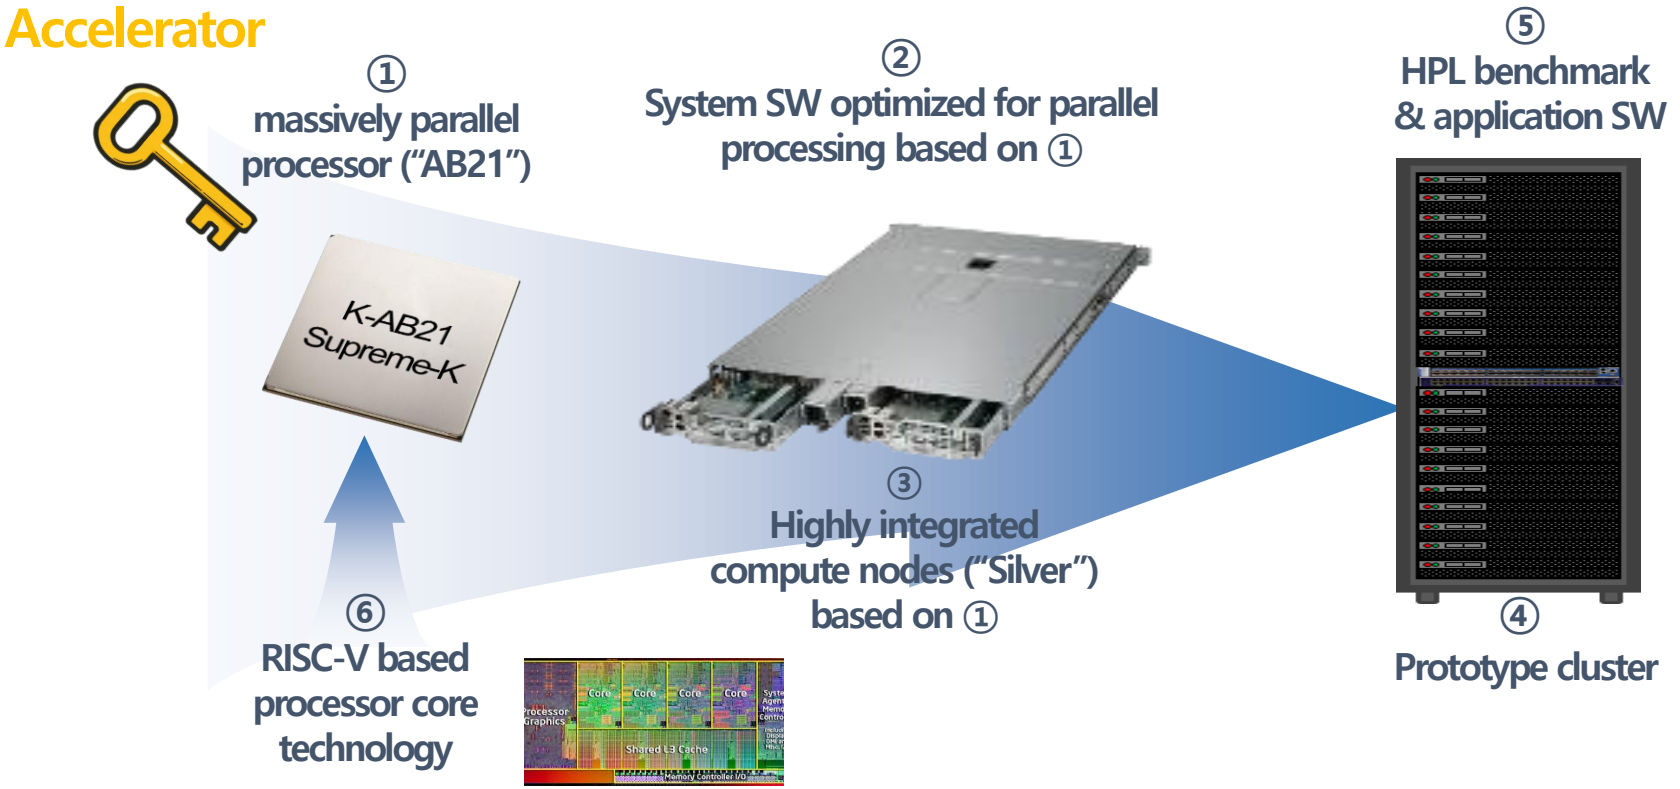

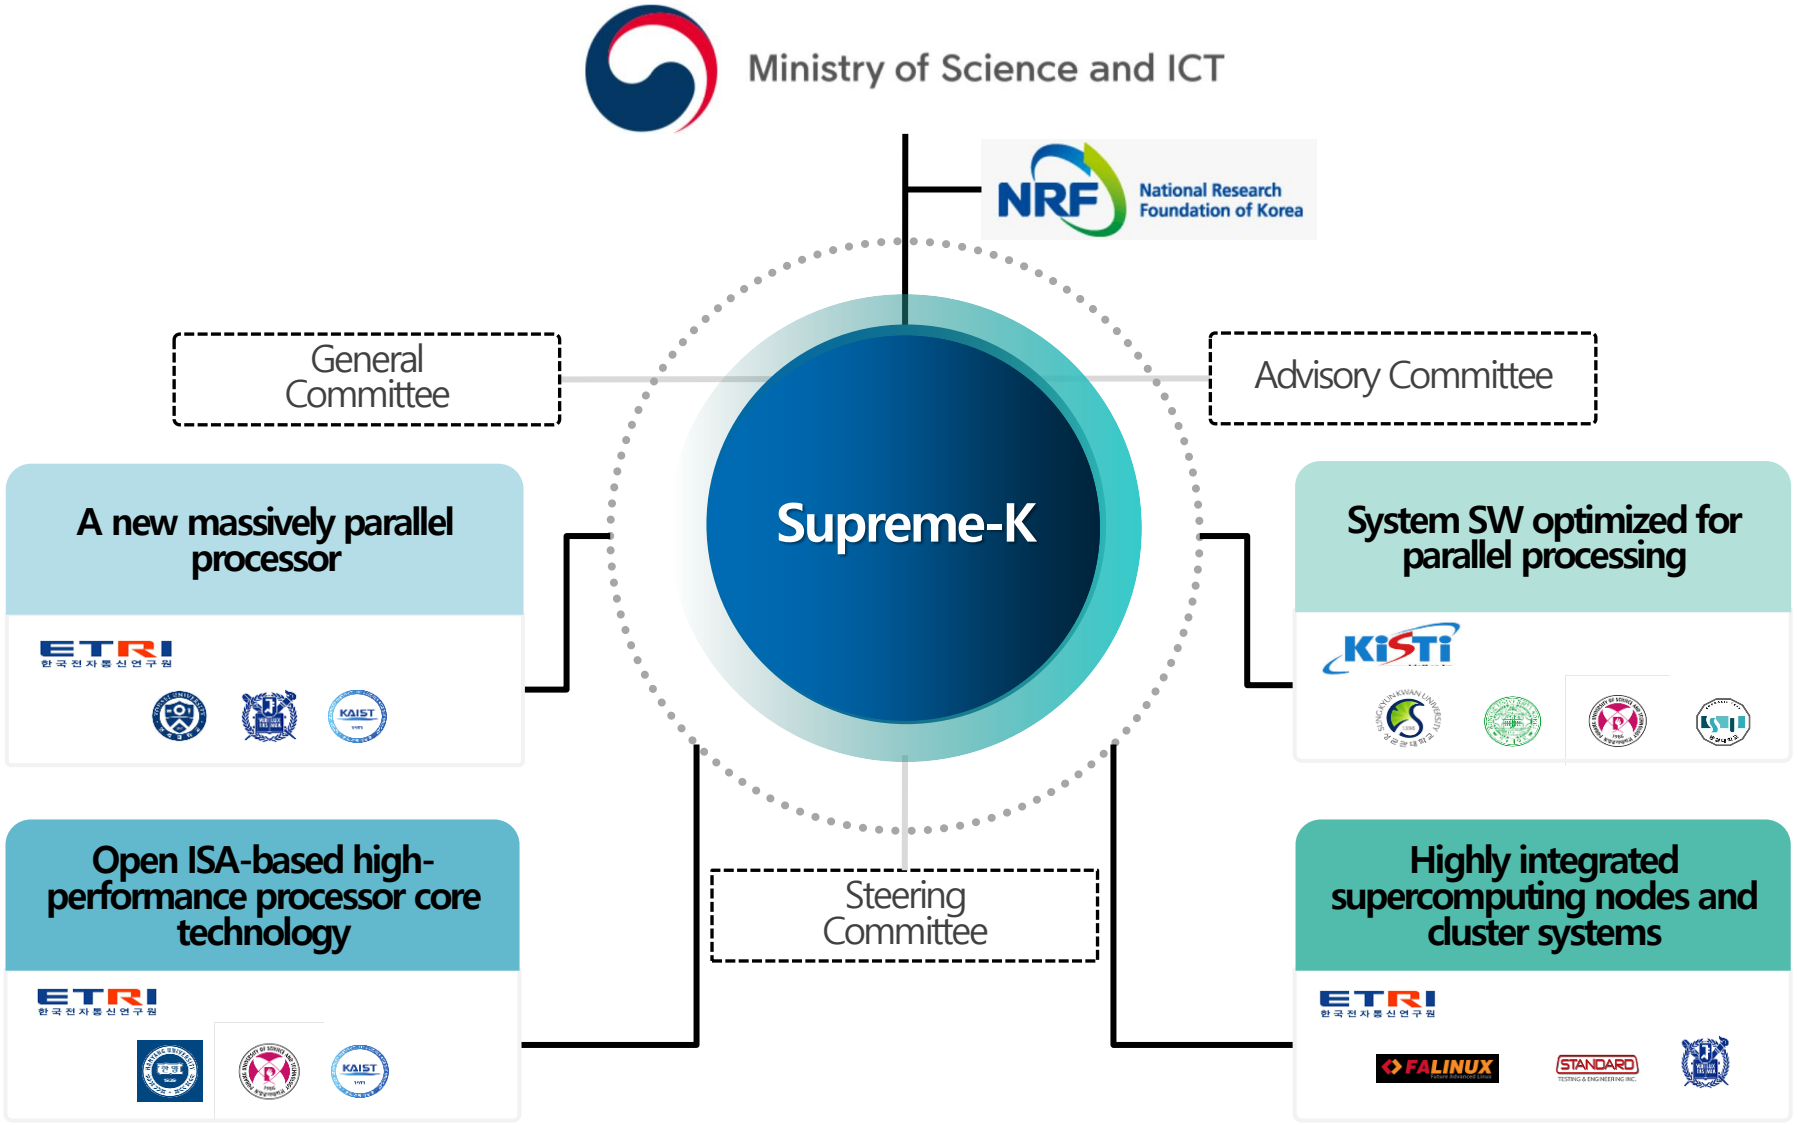
Ministry of Science and ICT
NRF National Research Foundation of Korea
General Committee
Advisory Committee
Supreme-K
A new massively parallel processor
ETRI
KAIST
System SW optimized for parallel processing
KiSTi
Open ISA-based high-performance processor core technology
ETRI
KAIST
Steering Committee
Highly integrated supercomputing nodes and cluster systems
ETRI
FALINUX
STANDARD

- HW/SW Co-design for HPC accelerator "AB21"

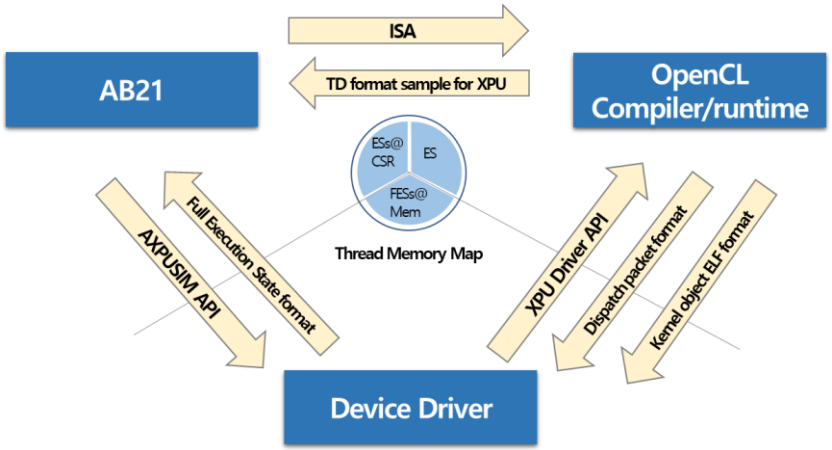

- Optimization of Data movement
  - Data sharing in a shared virtual memory between Host and AB21
- Comprehensive Systematic Verification for System integration
  - SW simulator, FPGA-based HW emulator, SVP(Silicon Virtual Prototyping)
- Project Management
  - Issue/Risk management
  - System/Document Configuration management
- Preliminary research on expanding memory capacity to tackle the memory wall issue

Contents

01 Background

- Why
- For what

02 Overview

- Goal
- Consortium
- R&D Strategy

03 Development

- Accelerator
- Compute node
- Software stacks
- Cluster systems

04 Conclusions

- Lesson learned
- Future direction

Architecture Design

- [NoC] **ARM-CMN**(Coherent Mesh Network)-700
- [PE] **dual ARM(zeus) cores**
- [**Parallel Processing Unit**]
  - [XEMIS] 64 clusters
  - [XECM] 4 x Parallel processing unit/XEMIS
  - {16 FP array+64-b RISC-V (4,096 FP units/chip)}/XECM
  - D-cache & SPM (2x128KB)/XEMIS
- **DDR5 / PCIe GEN5** interface on ARM-CMN

Implementation

(w supergate)

- TSMC 12nm process node (2Q, 2024)
- Rpeak : **8TFLOPS@FP64**, **32TFLOPS@FP32**
- 1.4 billion gates and approximately 4,900 balls

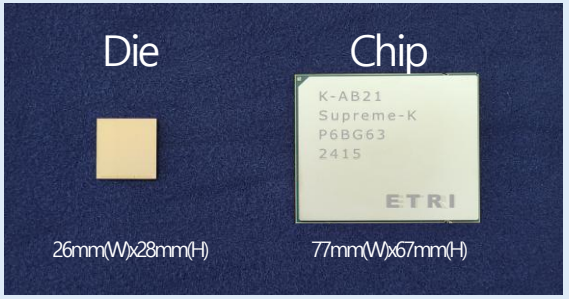

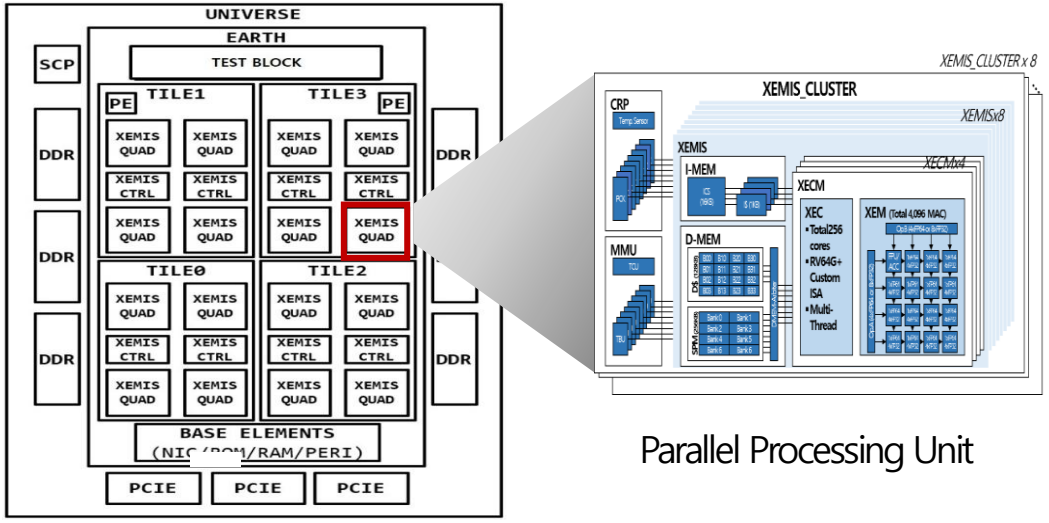

Parallel Processing Unit

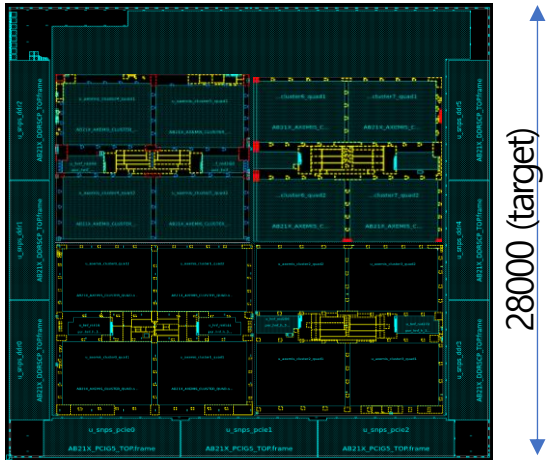

Floorplan after BE

❖ Parallel Processing Unit in AB21

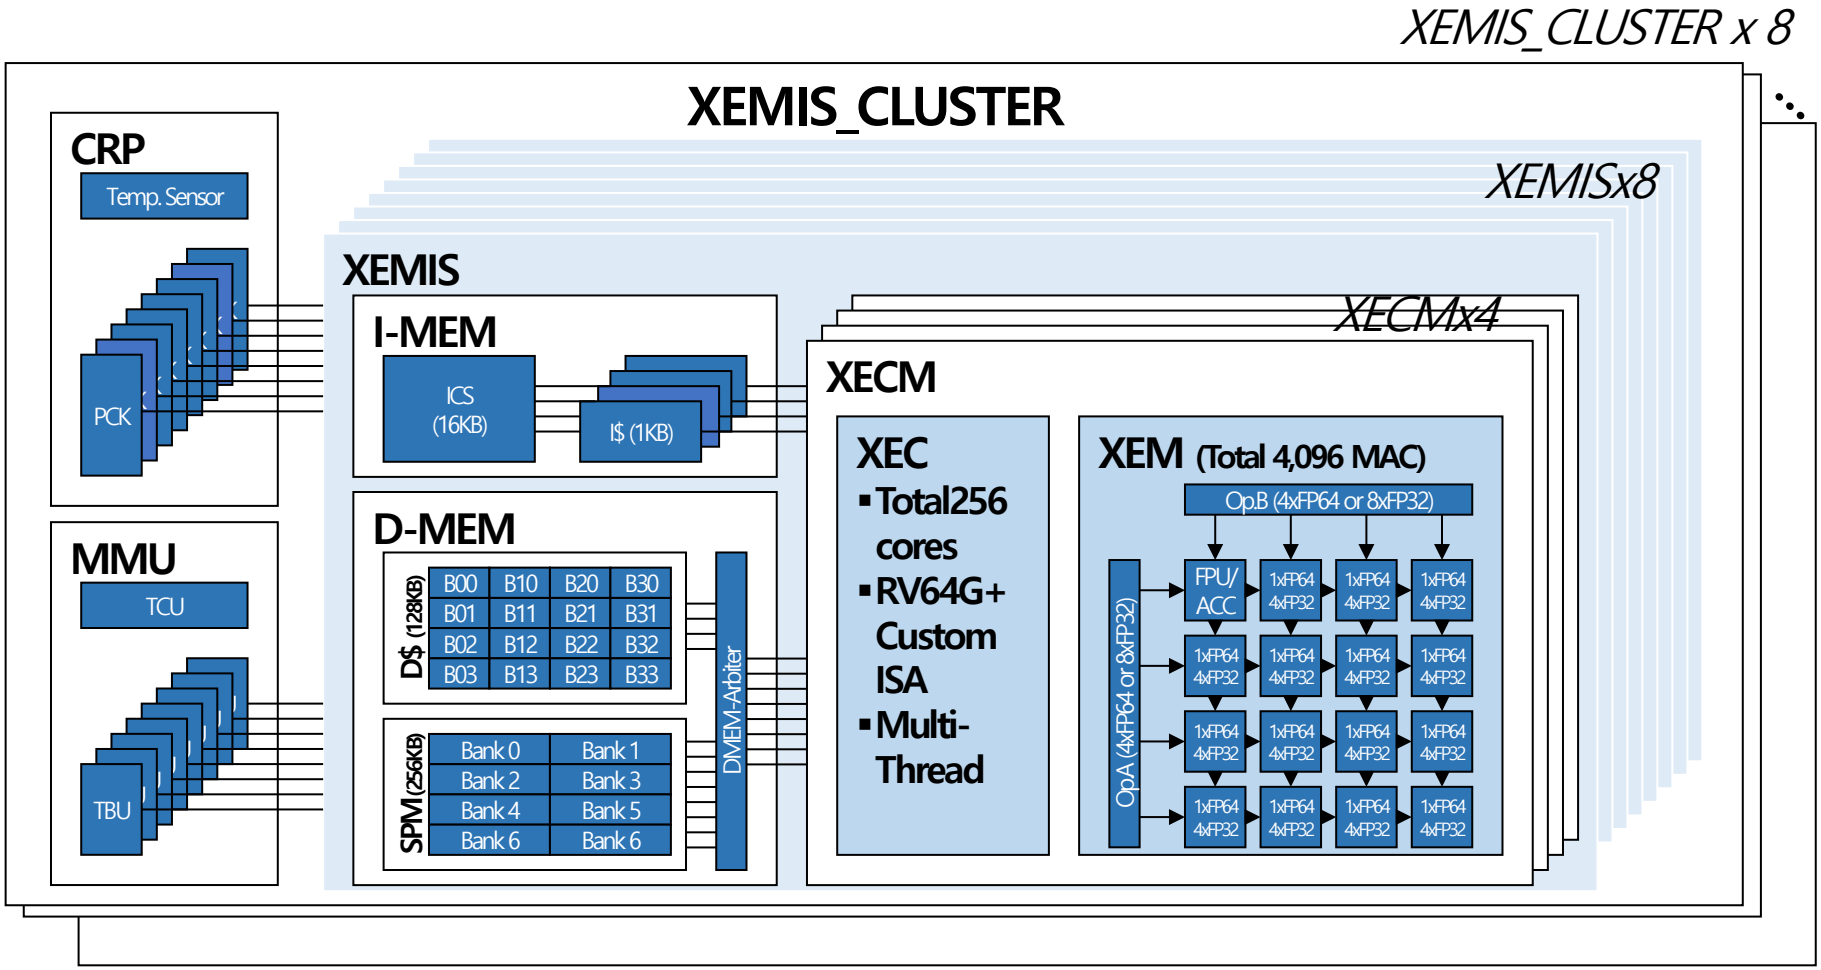

* Chun-gi Lyuh et. al., "Development of SoC to Accelerate 64/32 bit floating point matrix calculation for Supercomputers," JCCI 24, Apr. 2024.

Comparison

NVDIA GPU

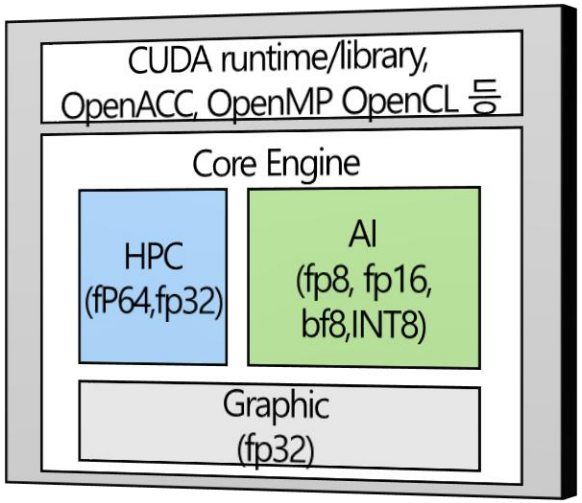

AB21

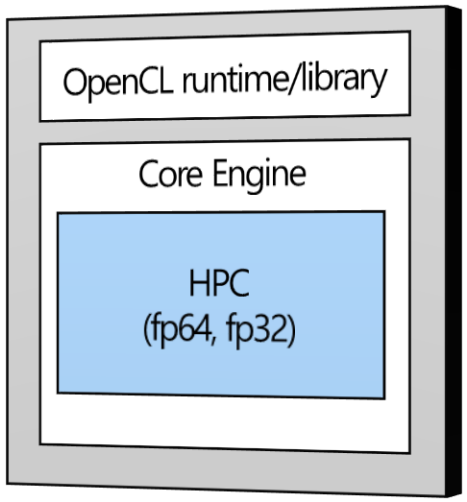

NPU

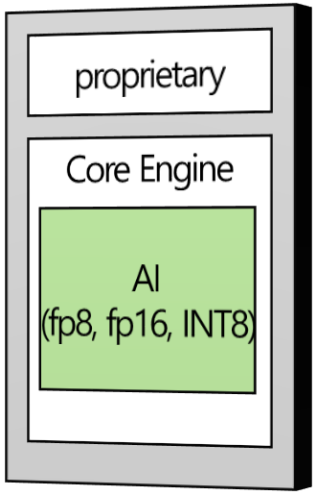

Product	AB21	NVIDIA V100	NVIDIA A100	NVIDIA H100	AMD MI250X
Release Year	'24.5	'17.6	'20.5	'22.3	'21.11
Rpeak (fp64)	8TF	7TF	9.7TF	26TF	47TF
Fab	TSMC 12nm	TSMC 12nm	TSMC 7nm	TSMC 4nm	TSMC 6nm

14% ↑

HW/SW Codesign

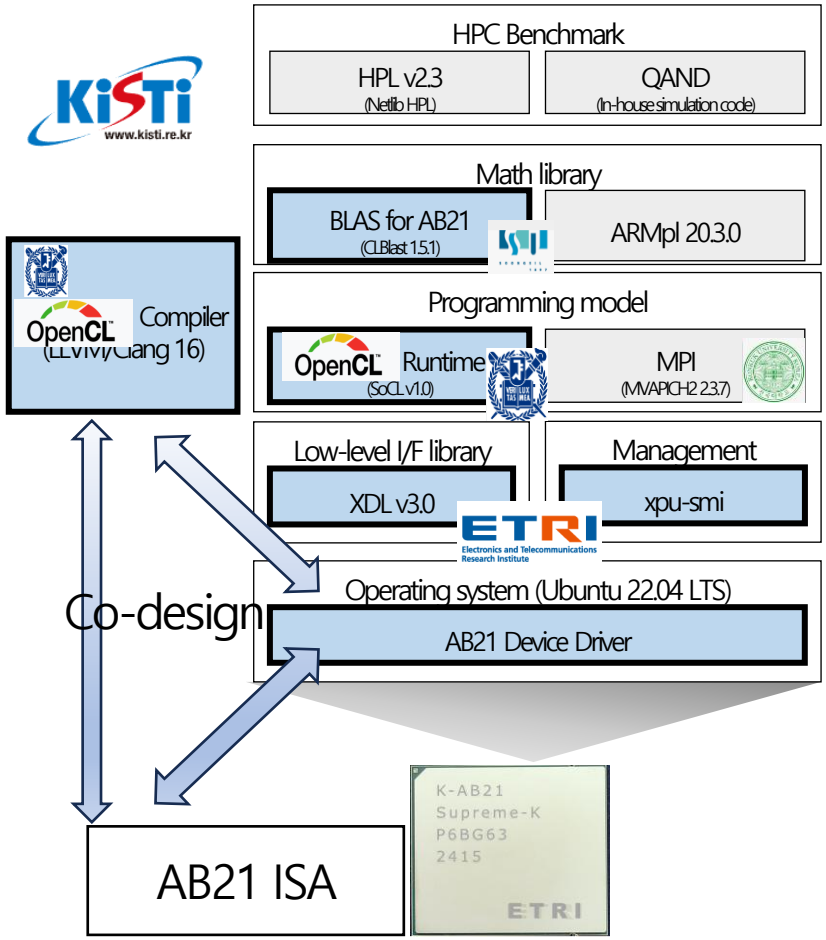

Implementation

- Device driver : controls AB21's behavior, parallel task scheduling, task execution, memory access, and device communication
  - Providing low level I/F library(XDL) and management tool(xpu-smi)
- OpenCL back-end compiler and runtime for AB21
  - OpenCL (Open Computing Language : an open, royalty-free standard for cross-platform parallel programming of heterogeneous systems
  - Compier : extended based on LLVM/clang
  - Runtime : extended based on POCL
- Customized BLAS Library for AB21
  - Supporting level 1, 2, 3
  - BLAS : customized based on CLBlast

Step-by-step verification through several stages

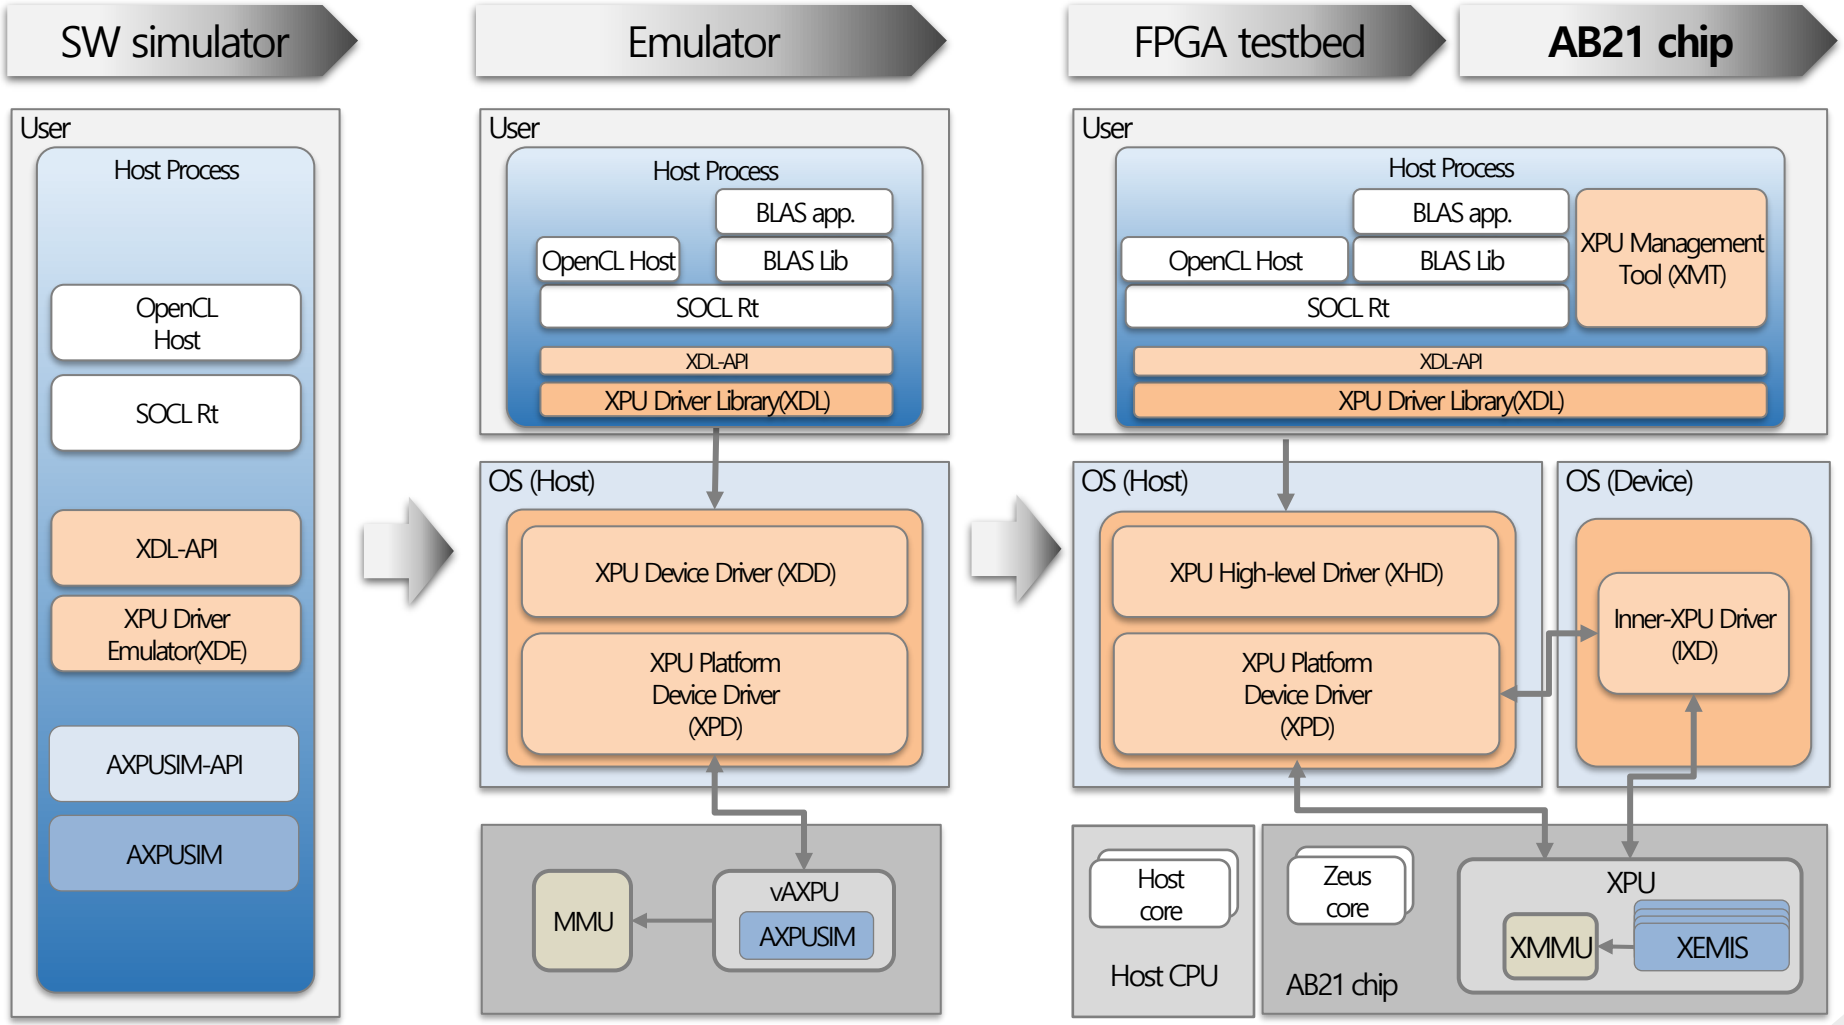

Architecture Design

- Customized Compute Node for dual AB21

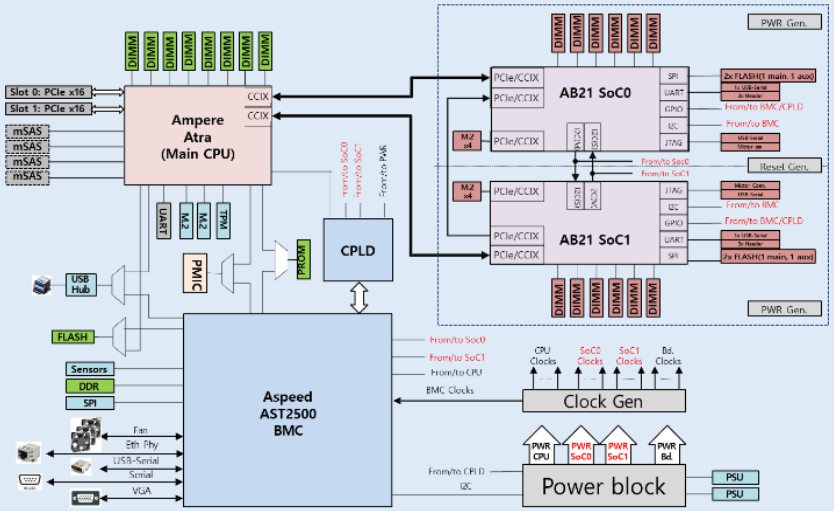

Implementation

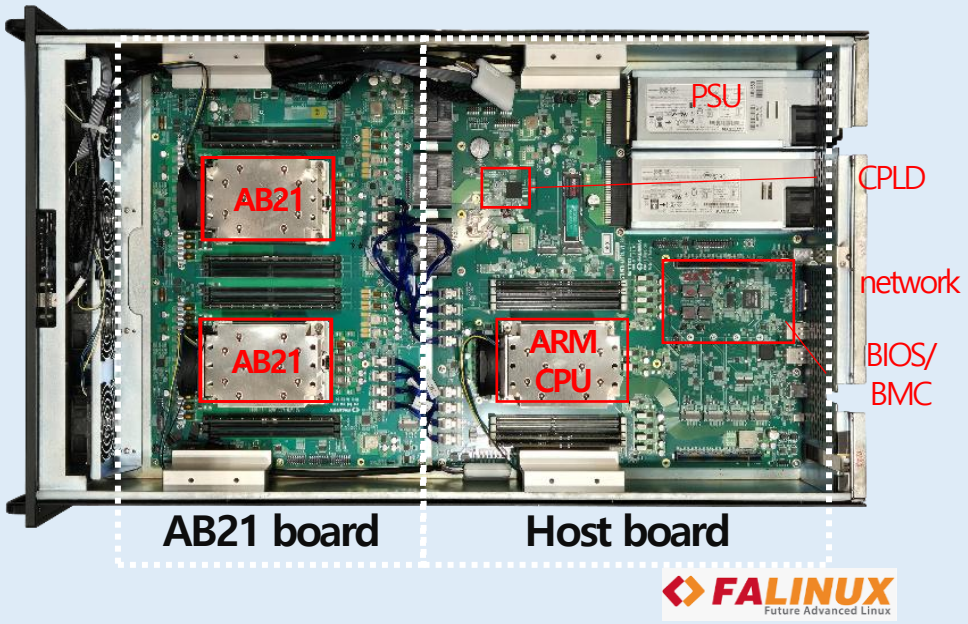

Main board (Host+AB21)	Host board	▪ Host CPU : ARM CPU (Ampere) ▪ DDR4 64GB * 8 ch. / PCIe Gen4 * 2ch. *16 ▪ Aspeed AST2500 BMC FW (OpenBMC), BIOS FW (UEFI)
	AB21 board	▪ K-AB21 * 2EA ▪ Refer to AB21 spec.
Network		▪ Ethetnet, IB, IPMI
Power		▪ 2400W * 2
Enclosure		▪ 19" standard chassis/3U ▪ Hybrid(Air + Liquid) cooling

Execution BLAS level 1(DAXPY)

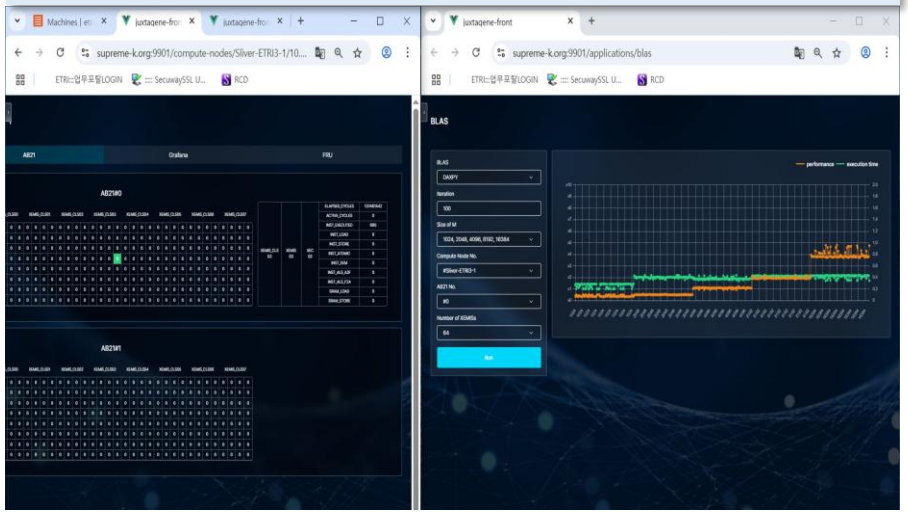

Execution BLAS level 2 (DGEMV)

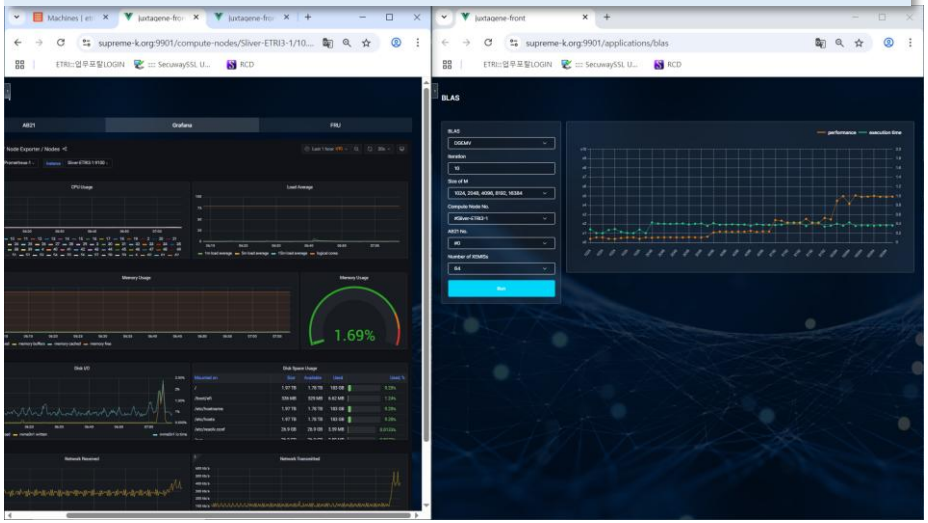

Execution BLAS level 3(DGEMM)

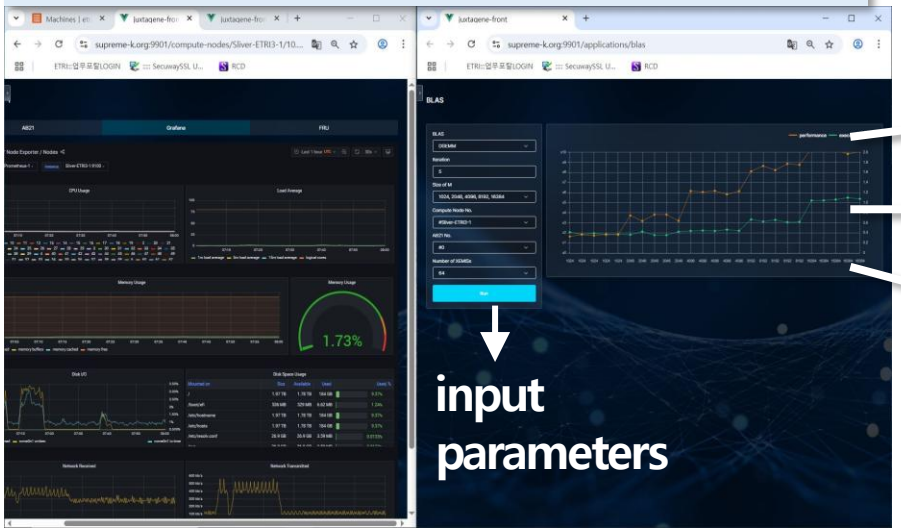

- performance
- execution time
- growing problem size

Logical view

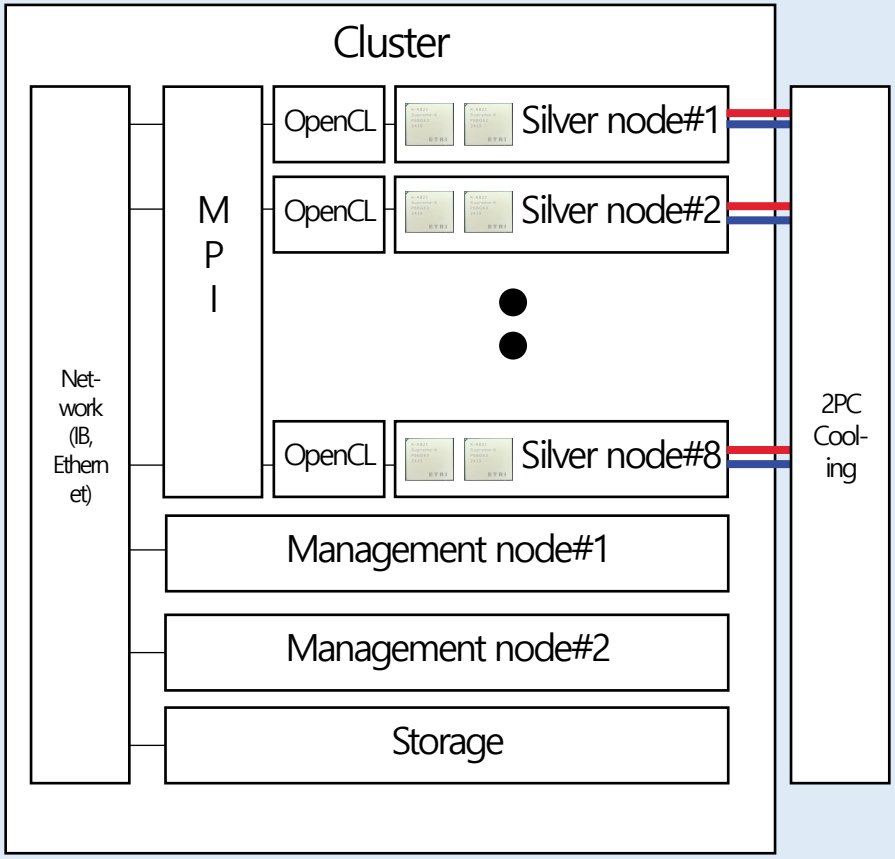

Physical view

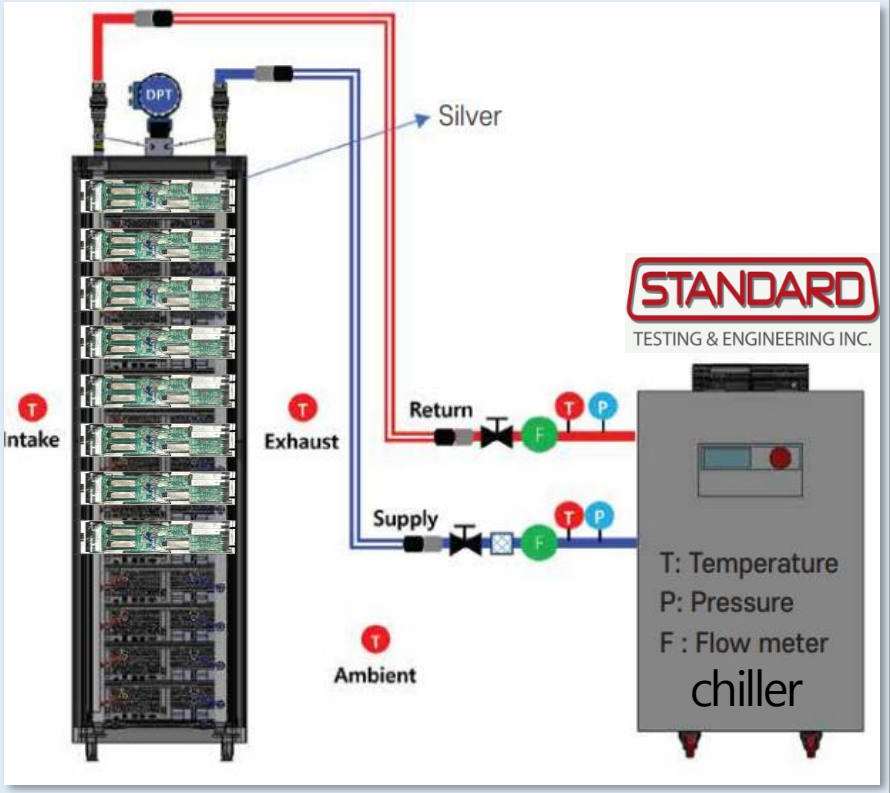

Single rack-scale cluster
with Partial 2PC(2 Phase liquid Cooling)*

@ISC25 ETRI booth, Hamburg

HPC cluster

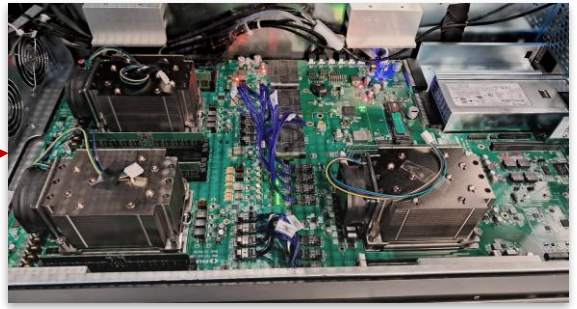

Compute node with dual AB21 chips

@ETRI Conference, Seoul

- What is the purpose of this chip?
- Which process node was used?
- What are the specs and performance?
- How is it different from NVIDIA GPUs?
- Is it commercially available?
- Does the chip support AI?
- What are the target applications?
- What are the software stacks for AB21
- What programming model does it use?
- What are the future plans based on current results?

Contents

01 Background

- Why
- For what

02 Overview

- Goal
- Consortium
- R&D Strategy

03 Development

- Accelerator
- Compute node
- Software stacks
- Cluster systems

04 Conclusions

- Lesson learned
- Future direction

❖ Nurion : Korea's fifth national supercomputer

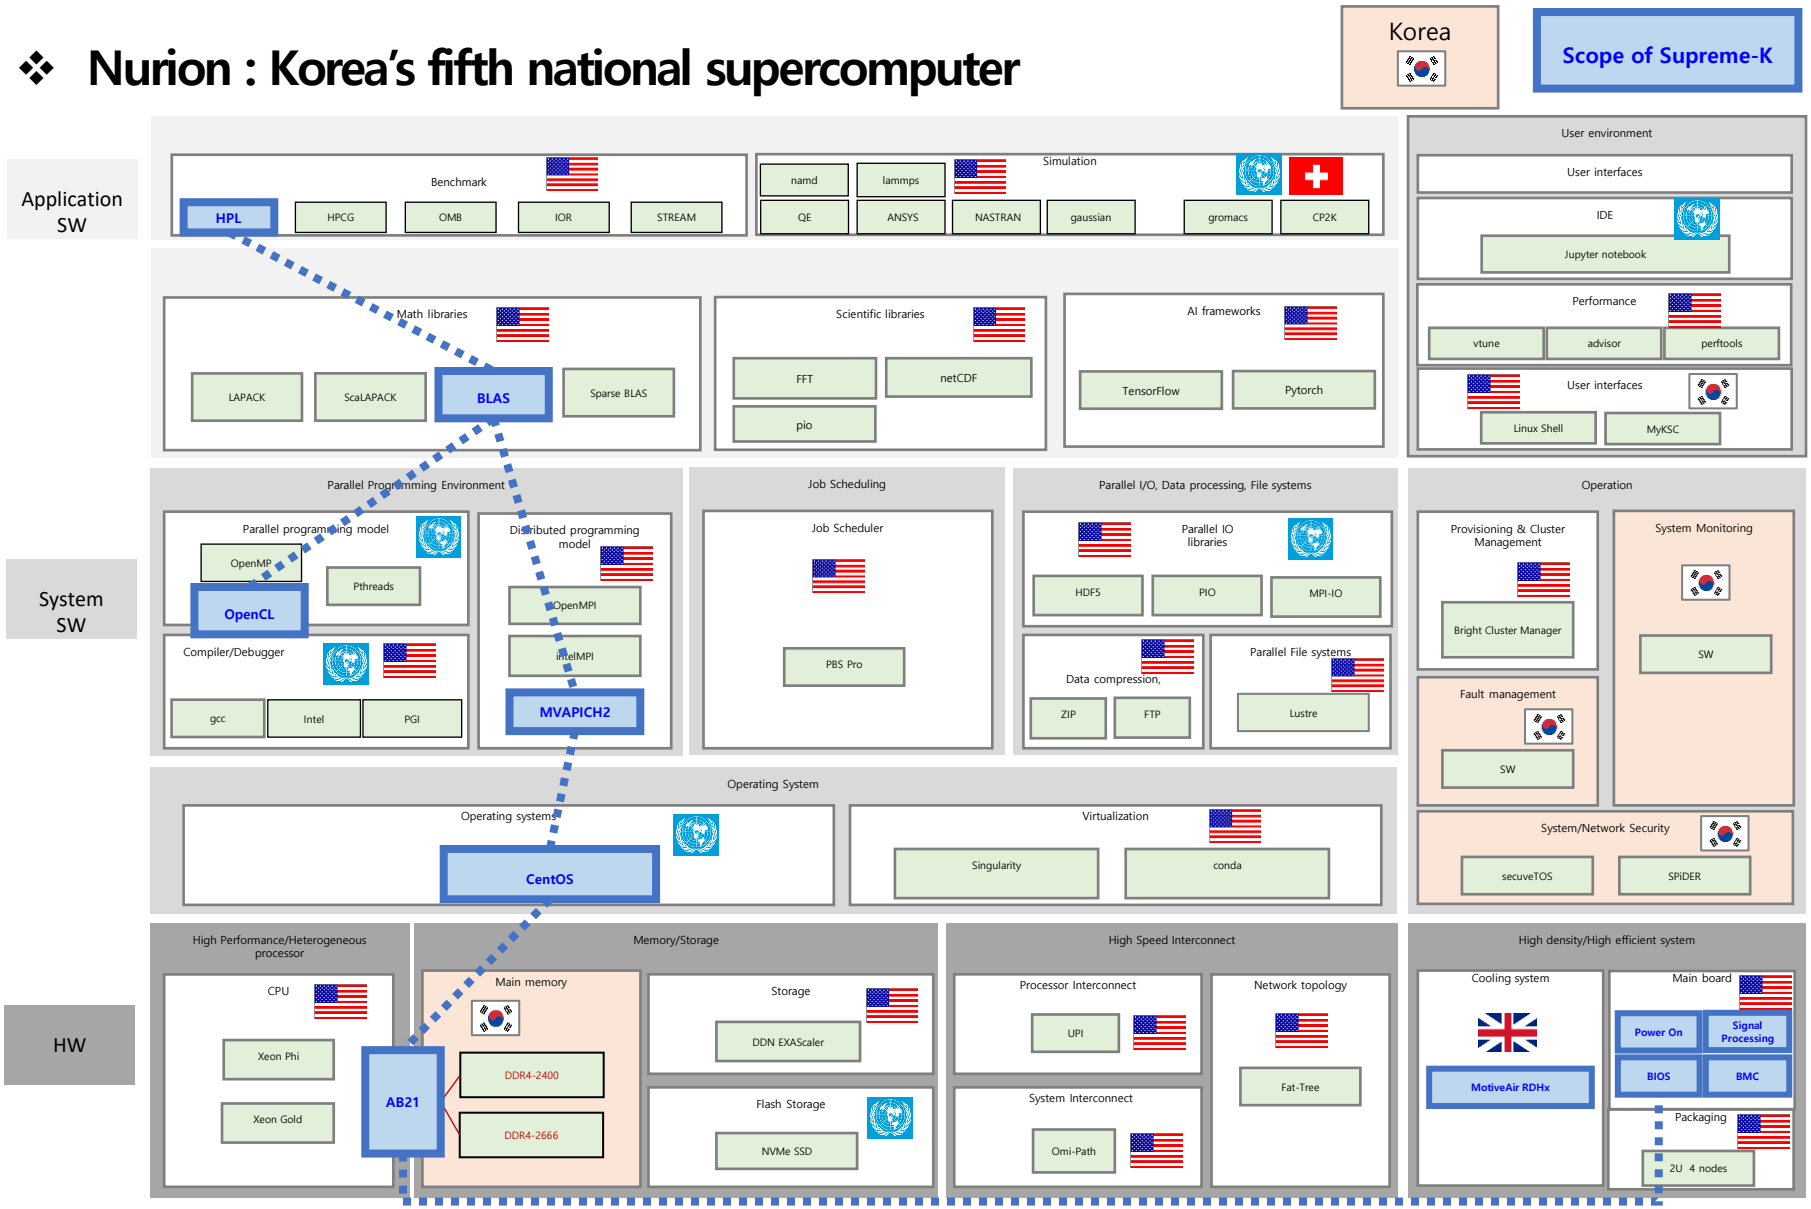

04
Conclusions :: Takeaway

SemiDynamics

ARM

Ampere Computing

RISC-V Core

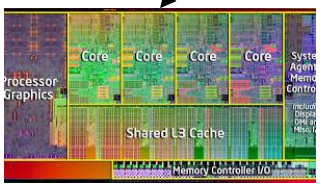

AB21

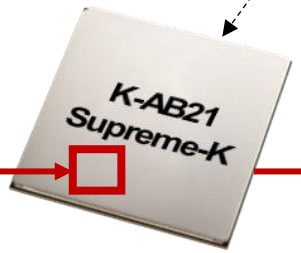

Compute Node

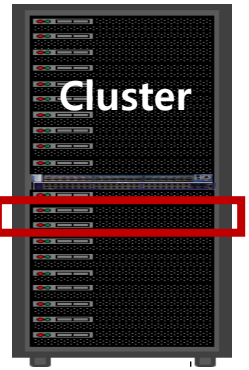

MVAPICH

OSU

ATOS/Eviden

MemVerge/Enfabrica

ORNL/LANL/TACC

04
Conclusions :: Future direction

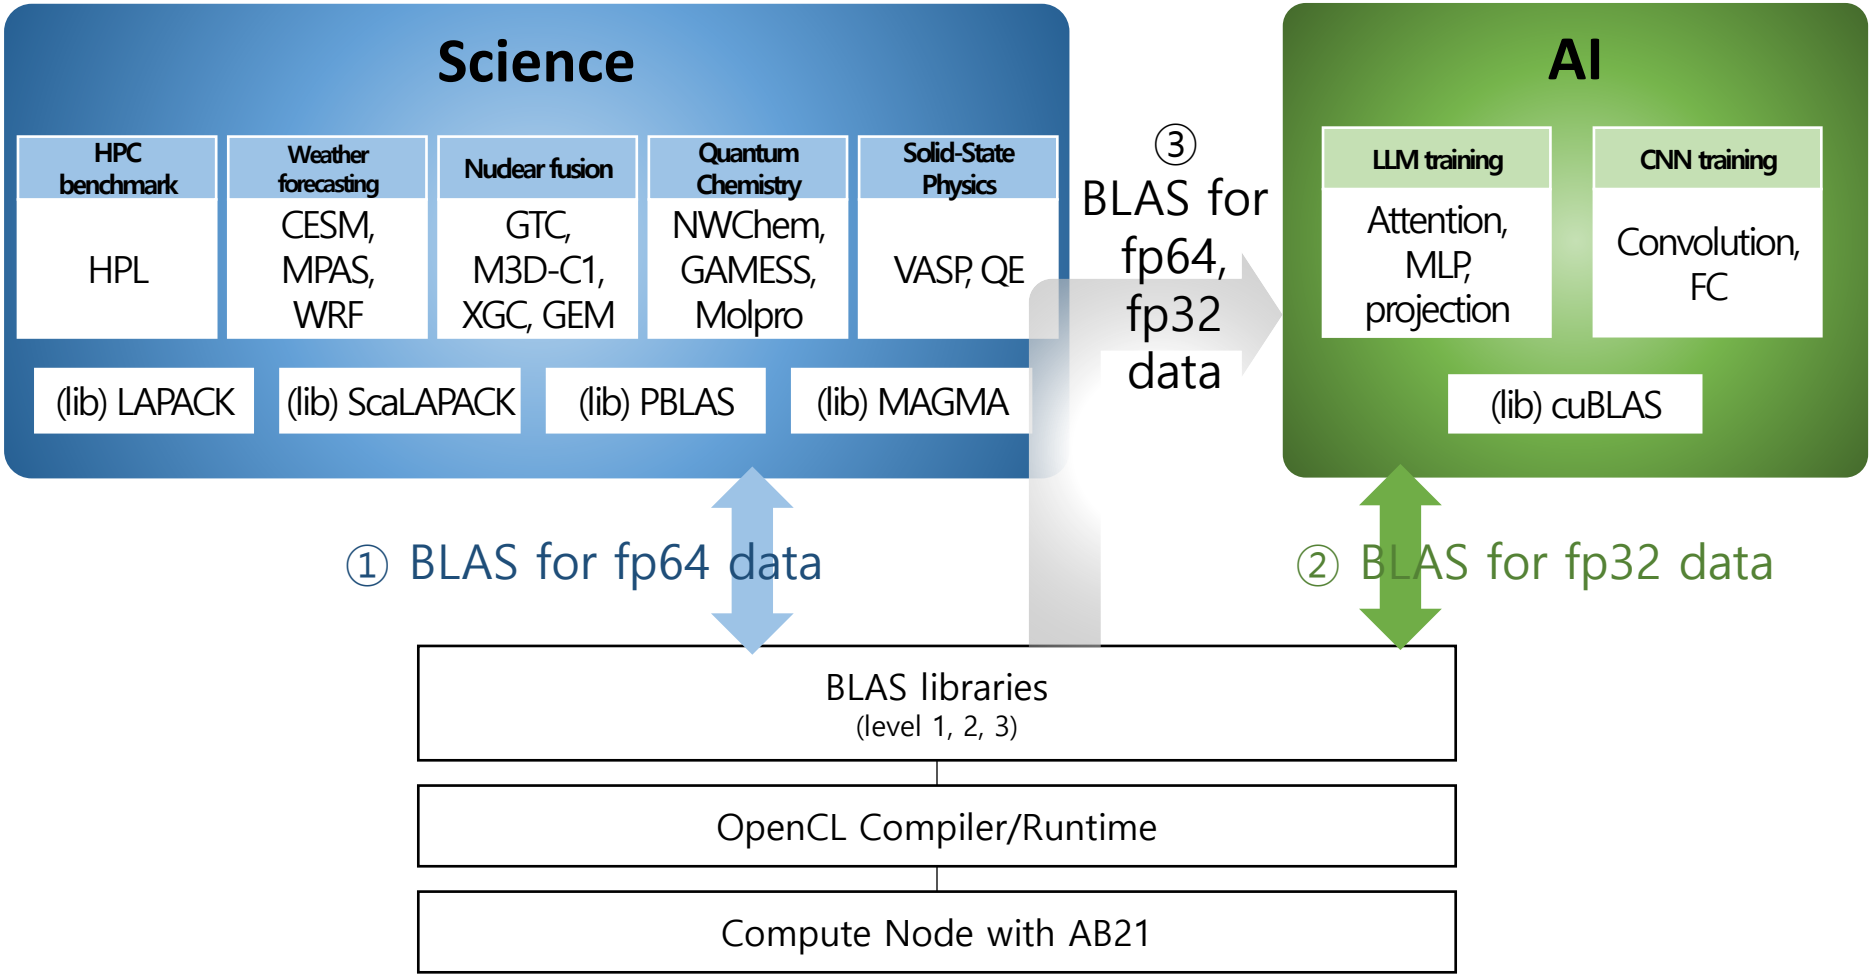

- ❖ GEMM(GEneral Matrix Matrix multiplication)
- ❖ GEMV(GEneral Matrix Vector multiplication)
- ❖ GEMM/GEMV operations account for 70%–90% of computations in both fundamental science and AI.

❖ Lesson Learned

- Taking a New path
- Walking it Together
- Proving That Korea can achieve this!

❖ Moving forward

- First stride toward HPC sovereignty
- Expanding our vision beyond HPC to AI
- Actively collaborating with global-leading partners
- Keep moving forward, step by step, guided by Korea's national HPC roadmap

Thank you!

ETRI
Electronics and Telecommunications
Research Institute