


Co-Design Story of Fugaku

Yuichiro Ajima
Principal Architect
System Development Division
Advanced Technology Development Unit
Fujitsu Limited



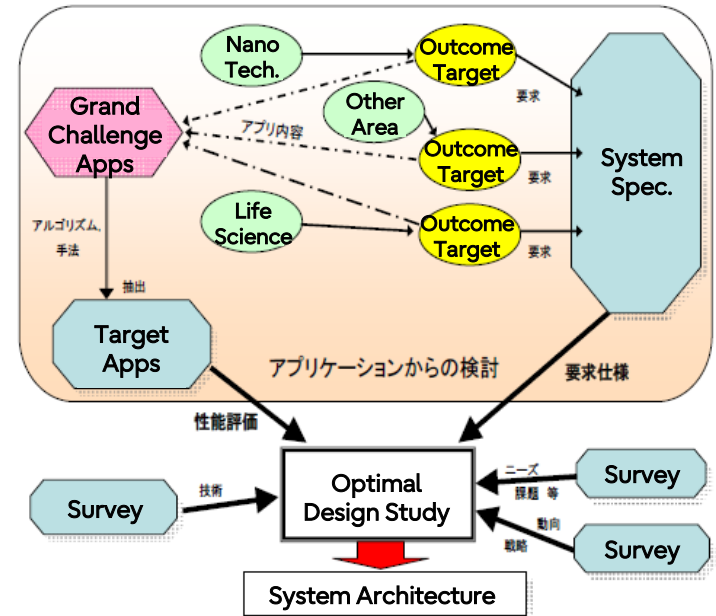


- Current Role
 - Principal Architect for system and interconnect architecture of supercomputer
- Past Projects
 - Development of the K computer
 - Development of the supercomputer Fugaku
 - Architecture lead of Tofu interconnect series
- Current Projects
 - FugakuNEXT Basic Design
 - Feasibility Study 3.0 for Future HPCI

- A system development methodology that integrates actual application developers into the evaluation and improvement feedback loop
- Core principle
 - Benefit from detailed feedback based on realistic user scenarios
 - Avoid over-reliance on synthetic benchmarks
- This approach evolved from lessons learned during the development of the K computer, Fugaku's predecessor

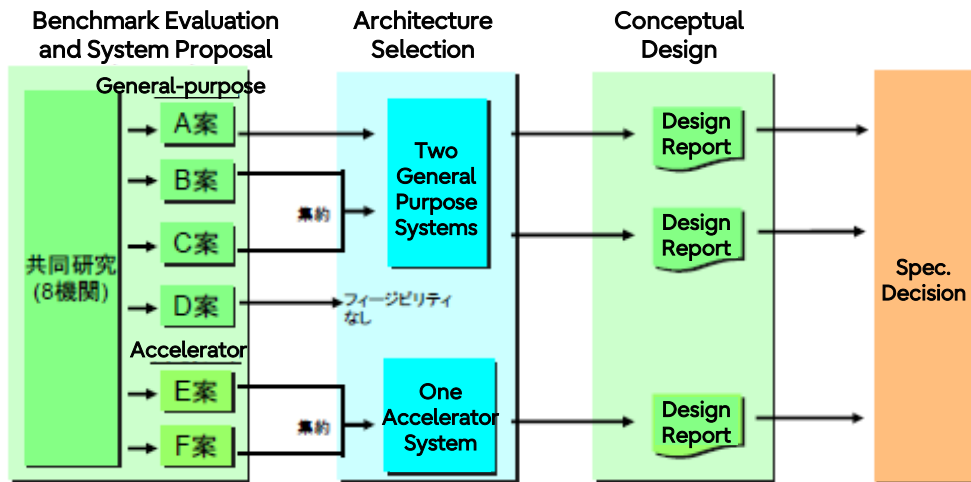
Looking Back the K Computer

- Conceptual Design: FY2006–07
- Detailed Design: FY2007–09
- Workflow concept in 2006 (right)
 - Key concept: “Grand Challenge Apps”
 - Yet, the overall process followed a one-way, procurement-like model
 - However, timeline was much longer
 - It was a development project



- Conceptual design began in 2006, with participation from multiple vendors and institutions

- Workflow concept (right)



- Concluded in 2007

- A hybrid system of Fujitsu's Scalar and NEC's Vector architectures

- Iterative benchmark submissions by both vendors
 - Resulted in architectures with similar final characteristics
 - A possible factor in NEC's project exit in 2009
- Design progressed steadily and ahead of plan
 - Architecture remained stable during detailed design
 - Contributed to half-year early production completion
- However, application feedback opportunities were limited

Co-Design Process Pioneered in Fugaku Development

2012 H24	2013 H25	2014 H26	2015 H27	2016 H28	2017 H29	2018 H30	2019 H31/R1	2020 R2
Feasibility Study 2012.7-2014.3		Basic Design		Detailed Design			Production and Delivery	

- “Feasibility Study on Future HPC Infrastructure” Project
 - FY2012-2013: Feasibility Study Phase
- “Flagship 2020” Project
 - FY2014-2018: Design Phase
 - FY2019-2020: Production and Delivery Phase

- Comprehensive survey of problems, social impacts, issues, and computational requirements conducted first
 - Moving from direct selection of focus areas and applications
- Computational Science Roadmap
 - Led by RIKEN
 - Completed March 2014, 344 pages
 - Social problems: 5 areas, 33 issues
 - Scientific problems: 20 areas, 112 issues

<https://cs-forum.github.io/hpci-aplfs/roadmap-2014/>




- Conducted in parallel with the Science Roadmap survey
- Three feasibility study teams
 - U-Tokyo team (our team): Latency-optimized cores
 - U-Tsukuba team: Accelerators
 - Tohoku-U team: Memory-intensive applications
- Each team selected appropriate applications for evaluation based on inputs from the Science Roadmap team
 - This was co-design in action from the very first step

- First half-year of the Flagship 2020 project
- Objective: Maximize social impact
 - Based on Science Roadmap outcomes and architecture feasibility study results
- Key decisions: Monolithic vs. heterogeneous system
- U-Tokyo team's architecture was selected
 - Fujitsu then joined as basic design contractor

- Focused on 9 priority area and corresponding applications
- Direct collaboration with application developers

- Co-Design report published
 - March 2022 (after project completion)
 - 424 pages
 - Documented detailed co-design outcomes
 - Evolution of architecture and design decisions
 - Application optimization methods and results
 - Also included summaries of feasibility studies

<https://www.r-ccs.riken.jp/fugaku/history/codesign-report/>



Examples from My Co-Design Experiences in Fugaku Development

- Three WGs: Architecture, System Software, Application

Work Group	U-Tokyo	Kyushu-U	Riken	Fujitsu	Hitachi	NEC
Architecture	✓	✓		✓		
System Software	✓			✓		✓
Application	✓		✓		✓	

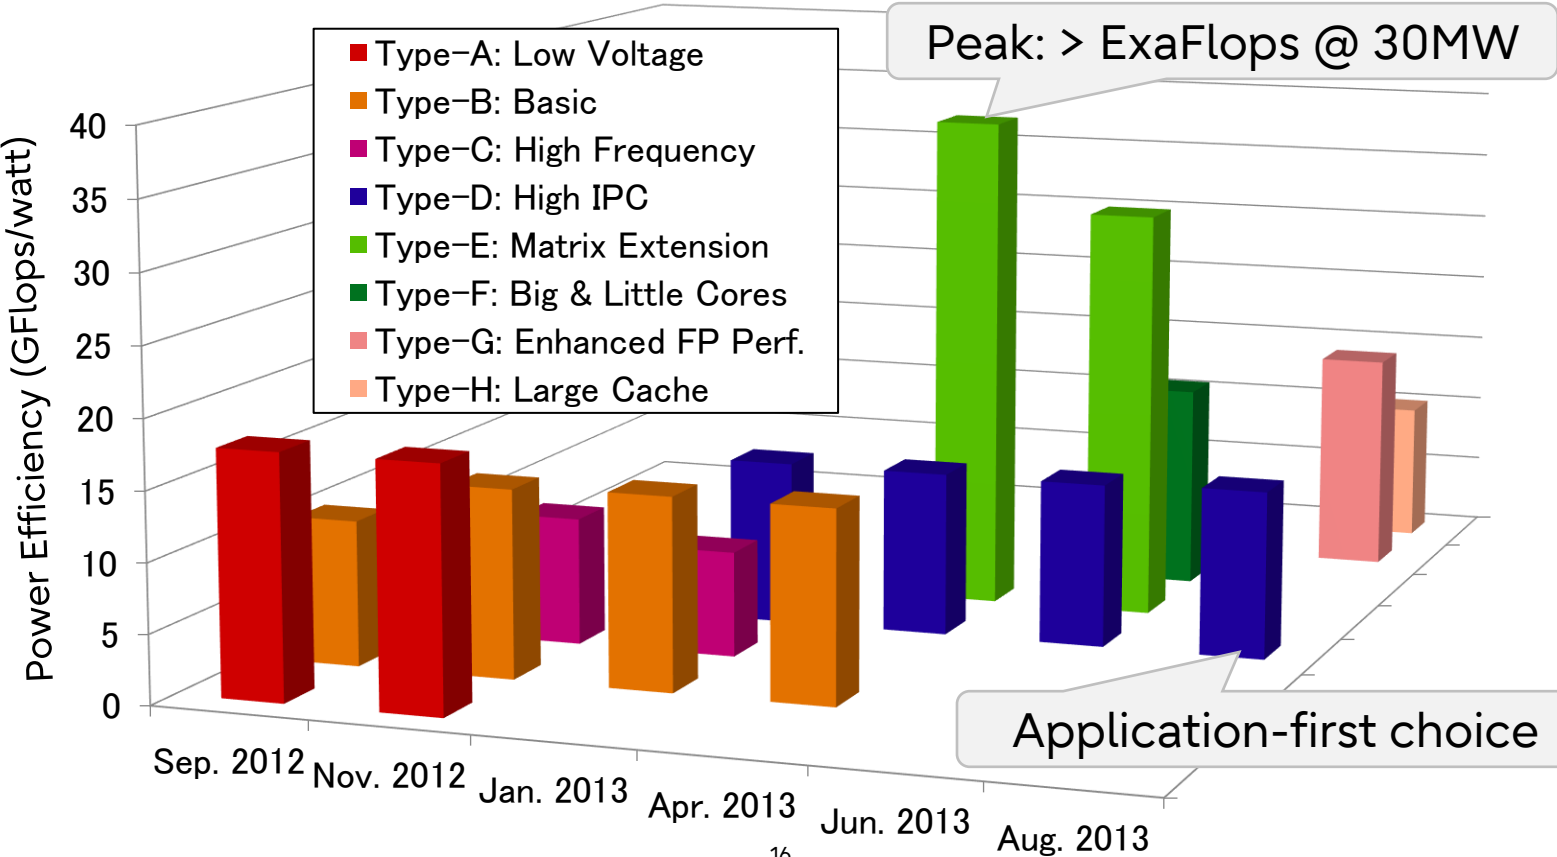
- Architecture Study: Iterative feedback loop

- Arch. WG presented three candidate architectures

- App. & Arch. WGs collaborated to predict application performance

- Arch. WG renewed candidates every other month

Architecture Study Progress



- Feature: Tofu Barrier – barrier synchronization hardware
 - Integrated all-reduce support
 - Originally supported single double-precision element per node
- Application request: Complex number support
 - Required parallel reduction of real and imaginary parts
- Architecture response: Extended element capacity
 - Leveraged available space in minimum packet size
 - Result: Capability for 3-element reduction

- Fugaku: Focused on simulation
 - The architecture was highly optimized through co-design
- Fugaku-NEXT: Expand computational science frontiers
 - AI for Science
 - Quantum-Hybrid Computing
- **Now in basic design phase**

Thank you

