Satoshi Matsuoka

Research Activities

My research is principally in system architecture, both hardware and software, for large scale supercomputers and similar infrastructures such as Clouds for HPC, and more recently, convergence of Big Data / AI with HPC, as well as investigating the Post-Moore technologies towards 2030 and beyond.

Over the years I have been involved and lead a number of large collaborative projects that worked on basic elements that are now significant for the current and more importantly future exascale systems, such as fault tolerance, green/low power, strong scalability, programmability, as well as large-scale I/O. Rather than the research projects of the subject matters for pure academic endeavor, the topics always lead to innovative, leadership-class production supercomputers that I had designed and made responsible for.

The principal methodology is to precisely capture the trends of computing, hardware and software systems and other related technologies such as datacenter power/cooling, and the societal needs especially in emerging areas, and design and create a production supercomputer that is novel and cutting edge, and at the same time, useful in actual production. System design for production systems often becomes a fine balancing act of novelty vs. useability, capability vs. cost, speed vs. power efficiency, etc. and designers would normally design within a current design envelop in a rather conservative fashion to for operational stability. Instead, my approach has been to identify the potential areas where design breakthroughs can be made, and then launch academic research projects that would be an amalgamation of software and hardware, in order to develop novel algorithms, software systems, and system designs incorporating new but useable technologies, accompanied by experimental production systems with application partners on board such that the design envelop can be expanded to satisfy both ends of the design spectrum, such as power vs. performance. The results of the research project would feed into the production system design that would often trailing just behind concurrently, resulting in a leadership system that would often demonstratively become #1 in Japan or even the world in some key metrics that the research projects enabled. Also the system and the technologies would also be a design template for systems that follow globally often becoming mainstream. This was best demonstrated with application of GPUs for HPC in TSUBAME supercomputers for which I received the IEEE-CS Sidney Fernbach Award in 2014, but there are other research topics as well.

The first series of major projects, as a junior faculty at the University of Tokyo starting in 1989, I lead projects regarding scalable and efficient implementations of parallel object-oriented languages such as ABCL (Actor-Based Concurrent Language) and other novel highly efficient task parallel languages on 1000 CPU-scale machines in the early 90's (Fujitsu AP1000 and ETL EM-4), and produced a series of novel system techniques for such purpose, such as extremely efficient parallel runtime on both conventional and specialized processors; although similar efforts were happening at the time in the world, such as coarse-grain dataflow and parallel object machines at MIT (J-Machine), task parallel runtime at UC-Berkeley, committed choice parallel logic programming languages

at Japanese national project ICOT as well as similar projects in UK/US, in retrospect the projects, we were leaders (somewhat unknowingly) in terms on implementation at largest scale, utilizing the 512 node AP1000 system at Fujitsu.

After having moved to Tokyo Tech. as an associate professor in 1996 and starting a new laboratory, I started two new mainstream projects that were tightly related, in terms of scaling the limits of parallel computing. One was to pursue the emerging HPC over wider distributed computing setting such as the Internet, later termed as the 'Grid', in collaboration with a team at AIST (former ETL), lead by Satoshi Sekiguchi as the 'Ninf' project; another was to pursue commodity cluster computing to achieve scalability at lower cost & exploiting the commodity CPU ecosystem, in collaboration with RWCP (Real World Computing Project), which was also being run predominantly under AIST as a national project, lead by Yutaka Ishikawa, Mitsuhisa Sato, etc. Collectively, I designed and built a series of experimental commodity clusters called the 'PRESTO' series in my laboratory at Tokyo Tech., starting in 1997, for the purpose of a) collaborating and running the cluster software developed at RWCP to exploit scalability of commodity clusters, b) to continue the system software / runtime research for task parallel computing, carried over from AP1000, and c) to develop system software and applications amenable to Grid computing on the clusters. Also and more importantly, many system hardware-level experiments and insights such as power/cooling, packaging, interconnect, I/O, resilience, etc., were acquired and made into research topics in their own right during the research phase, and more importantly, for a series of production systems leading up to Fugaku and beyond. Although an amateur at first, engineering skills were quickly attained to the extent that in 2000, the PRESTO-II+ became the first AMD CPU-based cluster in history to make the Top500 list, and subsequently, a redesigned PRESTO-III became as large as 256 nodes / 512 CPUs in 2002, and became the second fastest commodity cluster in the world, and #47 overall, on the Top500. Also world records were set on some novel applications in various numerical optimizations on the Grid, employing multiple PRESTO clusters and the Ninf GridRPC software, using hundreds of nodes continuously over many hours, proving the viability of distributed computing using commodity processors, contributing in part to cloud computing in the modern times.

During that time, in Spring 2000, the computing center at Tokyo Tech. was being completely restructured to become GSIC (Global Scientific Information and Computing) Center, and I was recruited to leading the supercomputing division thereof; although Tokyo Tech. had comparable budget to other leadership supercomputing centers in Japan, its, presence was almost non-existent globally, hosting custom-designed supercomputing whose design metrics were was becoming seriously outdated, such as the NEC SX machine. The results obtained in research from PRESTO clusters at the lab was quickly applied, to create a production testbed called the "Titech Campus Grid" (Tsubame0) project (2002-2006) which deployed a 650 node / 1300 CPU grid-of-clusters in pseudo-production mode within the Tokyo Tech campus, deploying the grid and cluster software. This was one of the first and largest such infrastructure to be deployed for daily use, side-by-side to the existing classic supercomputers at GSIC that were more than an order of magnitude expensive but similar in performance, not only contributing back to research, but also allowing GSIC to gain experience for designing and running an extremely large scale

cluster-based supercomputer in production leading to the TSUBAME series of supercomputers. Overall the project was a great success, hundreds of users prepping for the transition to a large-scale cluster-based supercomputing environment.

On the research side, in addition to research continued on grid and cluster middleware, especially system issues such as resource management, resilience, I/O, etc., it became clear that power efficiency would become the central metric to machine scalability. A series of sizeable inter-institutional projects were launched, the first being the 'Megascale' project which essentially ran in almost in parallel to IBM BlueGene L/P projects and attempted to establish the same low power / high scalability technology using low power processors but using commodity cluster technology that IBM USA refuted as being viable for system scalability. The project was lead by Hiroshi Nakashima @ Kyoto university and I served as a Co-PI. A prototype "MegaProto" high density commodity cluster system was built with assistance from IBM Japan using Transmeta processors. Although the project was successful, it also revealed limitations in merely seeking frequency+voltage reductions for power saving. In the NSF CTWatch journal, I wrote an article in the August 2005 edition, with the following observation"

"There are direct uses of low power embedded processors, augmented with HPC features such as vector processing, low power high performance networking, etc. Examples are BlueGene/L, Green Destiny,[4] and MegaProto.[5] Fundamental power savings are realized with lower voltage, smaller number of transistors, intricate DVS features, etc. In fact, BlueGene/L has demonstrated that the use of low power processors is one of the most promising methodologies. There are still issues, however, since the power/performance ratio of embedded processors applied to HPC are not overwhelmingly advantageous, especially with the development of the power efficient processors that will be arriving in 2006-2007, where similar implementation techniques are being used. Moreover, although one Petaflop would be quite feasible with today's technologies, to reach the next plateau of performance, i.e. ten Petaflops and beyond, we will need a ten-fold increase in power/performance efficiency. In light of the limits in voltage reduction and other constraints, the question of where to harvest such efficiency is a significant research issue."

In the mean time, the success of commodity clusters in RWCP/AIST and at Tokyo Tech., along with some early successes of the AIST-Tokyo Tech. Grid projects, a joint project proposal was made by AIST and Tokyo Tech. and its collaborators to launch a national grid project, with underlying facilitation of sets of modern commodity cluster computers for modernization of supercomputing centers in Japan; as a result of this, a national project of developing a Grid middleware common to all supercomputing centers in Japan was launched as the NAREGI project (2003-2007) as a large national project, headed by Ken Miura formerly of Fujitsu, and hosted by the NII (National Institute of Informatics), Japan, and I become the co-PI for developing the core grid resource management system. Also, although the infrastructure proposal part became somewhat cut down, still a set of divergent, moderate sized clusters of 32~128 nodes, and dedicated supercomputers were facilitated at NII and IMS (Institute for Molecular Science) Japan as testbeds, and molecular science became the primary target application. The lack of large scale clusters were compensated by the project participants from respective supercomputing centers facilitate a production cluster on their own, such as the Tokyo Tech. Tsubame 1.0. Note that this timeline is in complete parallel to the NFS TeraGrid which also started in 2003.

Although NAREGI was a mixed success, some of the middleware and the operational procedures became the platform for HPCI, a supercomputer consortium in Japan involving flagship K/Fugaku as Tier-1 and other major centers hosting machines such as TSUBAME as Tier-2, and became a national foundation for supercomputing for research to this current day.

Parallel to NAREGI, design of the TSUBAME-1 started in 2004, based on the research results from PRESTO clusters and experimental production results from Titech Campus Grid and the NAREGI testbed. The objective was to create a machine whose capabilities would match or even succeed the Earth Simulator, while being based on commodity technologies, and be within GSIC budget which is a small fraction of the Earth Simulator project, at around $30mil vs. $700mil. Numerous technical and administrative challenges had to be overcome, but in the end TSUBAME-1.0 was built, jointly by NEC, Sun Microsystem and Tokyo Tech, and chiefly co-designed chiefly and jointly by Andy Bechtolsheim then at Sun, and myself. TSUBAME-1.0 was installed late March 2006, and successfully superseded the Earth Simulator on the Top500 in its first entry in June 2006, until then the incumbent No.1 flagship supercomputer in Japan and Asia (#7 in the world). Tsubame-1.0 was an epitome of the most advanced cluster technologies of the time, using multi-core AMD x86 processors in fat node (8 socket) configuration, nodes interconnected via multi-rail Infiniband, that not only served as parallel computing interconnect but also as I/O interconnect with Lustre parallel filesystem being run on a cluster of JBOD disk nodes. A total of 5000 nodes / 10,000 CPU (cores) were available to the user, with comprehensive scheduling software that allowed coexistence of large-scale MPI jobs with significant numbers of small, ensemble jobs in a multi-tenancy environment, realized in modern clouds today. Tsubame-1 was a great success of not only in terms of the machine itself, but also motivated the classic supercomputing users to modernize their applications to massively parallel computing, eventually leading to K and Fugaku. Most importantly, it had put Tokyo Tech. GSIC as one of the leadership supercomputing centers in Japan and of the world, from being nowhere for decades, by immediate application of research results to system design.

Despite the success of Tsubame-1, it was already apparent that, novel technological leap had to be made to achieve more than an order of magnitude performance leap for the next generation Tsubame-2 slated to be in operation in 2010, as mentioned earlier. In fact strawman architecture of Tsubame-2 was being considered in 2004, and concluded that the use of multimedia/graphics processors such as the upcoming Sony Cell Processor, or emerging GPU with programmable visualization pipeline repurposed as general compute engine would be viable, as fundamentally there are modern incarnation of classic vector processor in commodity space, with superior power efficiency c.f. CPUs. After a few years of unfunded research, a major research project “Ultra Low Power (ULP)-HPC” (2007-2013) was funded, and studies were done to achieve low power and programmability of GPUs on supercomputers. The first prototype GPU cluster was build in 2007, using 128 NVIDIA GF92 GPU cards, and was the first sizeable GPU-powered HPC cluster in the world; there, various system software and algorithms as well as selected applications were tested, in particular molecular docking application with Yutaka Akiyama @ Tokyo Tech, and in particular optimized implementation of 3-D FFT which was an adaptation of classic

algorithm of vector supercomputer was applied to great success, allowing us to estimate that we get 4 times better performance/Watt compared to BlueGene/P that ran the same algorithm. We then teamed up directly with NVIDIA to retrofit TSUBAME-1.0 to facilitate 648 NVIDIA G200 GPU, the first GPU designed to be installed in a datacenter environment, in experimental production mode, the machine redesignated to be TSUBAME-1.2. We conducted application porting and performance & operational studies as well as encourage users to try to use the GPUs in their applications, as well as algorithmic studies to study the hybrid use of GPUs and CPUs, in high performance setting such as Linpack; as a result, TSUBAME-1.2. became the first GPU-based supercomputer to be ranked on the Top500, retaking the #1 spot among the Japanese supercomputers (#29 overall).

Based on the ULP-HPC project and Tsubame-1.2, TSUBAME-2.0 was designed to make extensive use of GPUs as the main computing engine, as well as to inherit the operational cluster software stack from Tsubame-1. After extensive co-design with the vendors, NEC-HPE alliance won the contract, and Tsubame-2.0 was installed in Nov. 2010 as the first petascale supercomputer in Japan, facilitating 4224 NVIDIA M2050 GPUs hosted in 1408 newly designed HPE SL390s blade nodes, interconnected with dual rail QDR full fattree Infiniband. Linpack performance was 1.192 Petaflops, ranked #1 in Japan and #4 globally on the Top500, increase from Tsubame-1.0 being 31.2x in 4.5 years. Moreover, on the Green 500, it scored 958 Megaflops/W, being ranked #3 overall and #1 as a production supercomputer. Although it was superseded by the newly deployed K Computer in the June 2021 ranking, nonetheless it retained #5 ranking, and moreover, won the ACM Gordon Bell Prize along with the K Computer in November 2011, realizing 2 Petaflops on a production application (dendric metal solidification simulation). Another novel aspect of TSUBAME-2 was the incorporation of flash-based drives on every compute node, which is now quite common including burst buffers, but at the time was met with significant reservation from the vendors. Trace-based records of I/O activities on TSUBAME-1.0 had revealed that, significant overhead was being incurred in scratchpad I/O in the parallel filesystem, which could be alleviated with local flash drives, and I calculated that the I/O duty cycle would last the lifetime of the machine. Local Flash drives also facilitated extremely fast checkpointing as will be described later.

Tsubame-2 was further upgraded to Tsubame-2.5 in the fall of 2013 by replacing all the M2050 GPUs with the Kepler K20 GPUs, boosting the performance to 5.4 Petaflops in double precision floating point, and 17.1 Petaflops in single precision floating point, the latter superseding the K-Computer performance. In order to exploit the latter, various mixed precision research was conducted, as well as the resulting applications such as molecular dynamics apps deployed in production.

To further improve the power efficiency as well as improving the cooling efficiency in terms of PUE, prototypes for the succeeding TSUBAME-3.0 was being constructed in the ULP-HPC project. In 2014, in collaboration with Green Revolution Cooling, NEC and NVIDIA, a prototype cluster TSUBAME-KFC was constructed, which hosted 40 nodes facilitating 160 Kepler K20 GPUs, immersed in liquid (mineral oil) for cooling. This facilitated both extremely efficient cooling as well as low temperature operation for the

semiconductors, collectively realizing high efficiency. Along with systematic tuning effort, TSUBAM-KFC became #1 on the Nov. 2013 Green 500 at 4.503 Gigaflops/W, nearly 5x improvement over Tsubame-2.0 in 3 years, signifying the importance of the modern liquid-cooling infrastructure to achieve high green efficiency. TSUBAME-KFC was later upgraded to TSUBAME-KFC/DL by replacing the Kepler K20 GPUs with K80 GPUs in 2015, in order to study the emerging deep learning to prepare for the upcoming TSUBAME-3.0, as well as to slightly improve power efficiency.

As research results from ULP-HPC was being transferred to TSUBAME-2.0 and KFC for production, three new research topics were started: 1) one was the "resilience in billion way exascale supercomputer" project (2011-2016), to seek utmost reliability in extreme scaling supercomputer with billion-way parallelism (Fugaku in practice facilitates billion way parallelism with 8 million cores, each capable of computing 64x FP16 FMA operations per clock cycle, resulting in 2 Exaflops of performance at 2GHz clockcycle). Early results from the project were already groundbreaking in that the hierarchical checkpoint system developed utilizing erasure codes and fast local flash drives on TSUBAME-2.0 was awarded the "perfect score paper award" at SC11 (it could not be awarded the 'Best Paper' Award due to COI, as one of the authors, Franck Cappello, was the technical papers co-chair). The work evolved into FTI (Fault Tolerant Interface), a widely used fault tolerance API for supercomputers, and recently in turn evolved into VeloC, now de-facto standard for checkpoint restart in large-scale supercomputers. Overall, the project showed that, resilience in extremely large systems would be possible while being compatible with commodity computing ecosystem.

The second project was the "Extreme Big Data" project (2013-2018) which seeked to investigate application of supercomputing system technologies towards extreme scaling of big data on modern supercomputers e.g., TSUBAME, and in turn, devise evolutionary roadmap of supercomputers to handle exabytes to zetabytes of data. One of the key objectives of the research was the development of efficient data-centric algorithms on massively scale supercomputers, including those with GPUs, such as sorting, word counting, as well as graph processing. In particular, the breadth-fast graph search algorithm developed for the K-Computer became the incumbent #1 on the Graph500 rankings, outclassing machines with much higher compute capabilities. Strawman study results of supercomputers for extreme big data was reflected to the design of TSUBAME-3.0 and its sister machine, AIST ABCI (AI-Based Computing Infrastructure), both of which I served as the chief designer / project leader.

The third and most important was kickstarting and organizing a set of smaller projects leading to towards a project to build Fugaku as the successor to the K Computer in the exascale era, with 2020 as the target goal of deployment. Endeavor towards exascale actually commenced internationally back in Supercomputing 2008 in which the IESP (International Exascale Software Project) exploratory meeting was held, and then meeting one was held in Santa Fe in April 2009, with the objective of drafting a common, international document to kickstart exascale projects in respective countries/regions. I became a member of the executive committee as the only member from Japan/Asia, along with members from the US such as Prof. Jack Dongarra and Paul Messina, and from Europe

such as Thomas Lippert and Mateo Valero. I helped to organize the meetings as well as became one of the contributors as well as one of the editors for the IESP Roadmap document. Parallel to the IESP efforts, domestic efforts to consider an exascale machine as a successor to the K-Computer was launched in the fall of 2010, which eventually turned into the SDHPC effort and document which had more of s comprehensive system focus covering issues from hardware to analysis of potential applications; again I served as one of the main and contributor to the effort. These two efforts collectively helped to launch the 'feasibility study' project for Fugaku/Post-K (2012-2013), and there I co-lead the applications investigation group to facilitated commonization of requirements as well as benchmarking, so that broad sets of applications can be quantitatively assessed together to devise the design requirements. After two years a 600 page report was crafted which identified the requirements of the next machine to achieve groundbreaking scientific results among multiple disciplines. I also was a member of the HPCI committee to assess the necessity of the next generation machine, and contributed to the report to justify the launch of the 'Flagship 2020 project' at Riken, from 2014.

During 2014-2017, after the project was launched, I served on the architecture investigation group which was the authoritative committee to officially examine the hardware architecture being developed day-to-day by the Flagship 2020 office at Riken AICS (predecessor to Riken R-CCS) and Fujitsu, and determine the overall architectural development directions. I also served on the advisory and governance committees for R&D efforts accompanying the Flagship 2020 hardware development project, namely the JST CREST "Post PetaScale" program that funded various system software efforts, as well as the Strategic Applications program that covered the co-design R&D of 9 target application areas being developed for Fugaku/Post-K. Overall, my comprehensive knowledge and contribution role at high (and sometimes very technically detailed) level helped the project to overcome some of the most difficult target goals and various crisis during that occurred during the years, and likely lead to my directorship at R-CCS.

As the research projects were going along productively, TSUBAME-2.5 / KFC production as well as various unique features I came up in research and quickly put into production such as multi-tenancy on 'fat' GPU nodes for resource utilization improvement, dynamic software power capping algorithm for hot summer daytime in light of the electricity shortage caused by the 3/11 Tohoku Earthquake, etc., were being deployed to make TSUBAME-2.5 to make it one of the most advanced operated supercomputer in the world, not just being GPU enabled and high performance comparable to the K, as well as contributing to the Fugaku/Post-K projects, I started designing TSUBAME-3.0 as the successor of TSUBAME-2.5. The objective was to make it not only a viable successor to the highly successful TSUBAME-2.5, the major objective was essentially make it a near perfect 'template' architecture for future GPU-based supercomputer, pursuing world-leading performance (both flops and bandwidth), (power)-efficiency, programmability, as well as incorporating the research results in resilience, as well as make it amenable to modern applications involving big data and emerging AI. As opposed to TSUBAME-2.5 which was only water cooled at the rack level, experience and data from TSUBAME-KFC was re-analyzed to achieve high density with direct warm-water cooling but with high efficiency, reliability, and serviceability, coexisting with high system density. Significant

shortcoming of TSUBAME-1.0 which was improved for TSUBAME-2.0 but still insufficient was the lack of bandwidth in the system, in that any data residing anywhere in the entire system (irrespective of CPU/GPU memory or NVM storage) can be moved freely without CPU intervention, only limited by device or intra- node or inter-node interconnects, using RDMA. Modern cloud / big data / AI features needed to be readily incorporated and made available to the user while not compromising performance; for example, TSUBAME-2.0 multi-tenancy was statically determined, and moreover, overhead was significant for interconnect and I/O for virtualized partition (users that required high performance was encouraged to use bare metal partition), but these limitations had to be eliminated, such that number of CPUs/GPUs were essentially invisible to all users, in that they all were allocated virtually partitioned nodes to fit their application needs.

After significant design efforts in collaboration with multiple vendors, HPE/SGI won the contract, and all of the design requirements were incorporated in their brand new Apollo 8600 system which was designed for TSUBAME-3.0 and then commercialized. Altogether, TSUBAME-3.0 was realized as a 540 node, 2160 NVIDIA P100 GPU and 1080 Intel Broadwell CPUs configured into a direct warm water cooled blade, with extensive intra- and inter-node full bandwidth connections using combinations of available chips and technologies. Applying the results from TSUBAME-KFC, TSUBAME-3.0 became world #1 on the Green 500 in June 2017 at 14.110 GigaFlops/W, besting the previous #1 in Nov. by 50% and 14.7x improvement over TSUBAME-2.0 in 6.5 years, or nearly 500x improvement over TSUBAME-1.0 in 11 years, allowing prepping for exascale.

As TSUBAME-3.0 was being put immediately into production, two joint projects with AIST emerged, centered around their new AIRC (Artificial Intelligence Research Center) which was founded in 2015. One was to jointly establish a new joint research center between AIST and Tokyo Tech. on big data using supercomputing (Real World Big data Computing – Open Innovation Laboratory, or RWBC-OIL for short), and I became the foundational director by becoming a joint appointee, in 2017. Concurrently, an AIST leadership supercomputer focused on Big Data & AI to serve AIRC & RWBC-OIL, especially their collaborations with the Japanese industry that seriously lacked restructures c.f. US GAFA, was proposed and funded by the ministry METI, thereupon I had set forth to design and build such a supercomputer based on TSUBAME-3.0, but AIST had totally deprecated its supercomputing efforts after RWCP. I turned that into an advantage by designing a datacenter and the machine collectively, with two goals 1) increased commoditization of the machine c.f. TSUBAME-3.0, eliminating expensive, custom build components, 2) commoditized construction of a datacenter, quick and inexpensive to build but retain the efficiency and environmental sturdiness resistant to disasters e.g., typhoons & earthquakes, and 3) despite the commoditization, retain performance & efficiency of TSUBAME-3.0. Another requirement was that, the procurement specifications of the RFP had to be entirely Big Data and AI centric, not only system specifications, but also the procurement benchmarks. So I set up a collaboration team consisting of researchers from AIST AIRC & RWBC, Tokyo Tech, to come up with a set of detailed benchmarks for various big data & AI supercomputing workloads, some of which likely to have affected some of such benchmarks we see today.

Fujitsu/NVIDIA/Intel was awarded the contract, and the resulting supercomputer was ABCI (AI-Based Cloud Infrastructure), facilitating 4352 NVIDIA V100 GPUs as the AI engine, as well as total of nearly 2PBytes of nearline NVMe's for high bandwidth data processing for big data and AI training, all of which hot water cooled at 42 degrees Celsius, completely chiller free even during hot summer. It became operational merely two years after being funded, including the datacenter which was built in six months including all the electrical and cooling facilities, on a former Rugby field. ABCI became the fastest supercomputer in Japan on the Top500, the spot it held until Fugaku came online, or #5 overall in June 2018 with 19.88 Petaflops, and also became the fastest machine in the world in May 2019 momentarily in the unofficial global Resnet-50 training competition that directly preceded MLPerf benchmarks, by being immediately put to research & AI industry use, including various research at RWBC-OIL. ABCI has recently been upgraded to ABCI2.0 in 2021 by the colleagues at AIST, adding 960 NVIDIA A100s to the existing platform, which was easy as I designed the whole machine and the datacenter infrastructure to readily accommodate such extensions.

Based on TSUBAME-3.0 and ABCI as well as the upcoming Fugaku, another research project was launched directly to tackle scaling issues of AI Training we were experiencing; Professor Koichi Shinoda became the PI and I became the Co-PI for the DEEP-CREST project (2017-2022), in which the objective of my team was to upscale the scalability of training to 100,000 nodes on large-scale supercomputers such as the upcoming Fugaku; this is difficult, not only due to systems issues but also algorithmic issues regarding training convergence. The project funded the creation of a dedicated High Performance AI Systems team I headed directly at Riken R-CCS upon my moving there in 2018 as the director, where we addressed and devised various novel algorithms and systems to attain such performance and scalability on ABCI and Fugaku, and was actually demonstrated as will be described below.

In the mean time, in the Summer of 2017 I was recruited to become the director of Riken R-CCS, which was to be restructured from AICS; it was created to parallel to the commencement of facilitation and operations of the K-Computer in 2011 for such purpose, but also accompanied research teams, being one of Riken's research centers. The tasks confronted with the new directorship was quite challenging in all fronts. Fugaku R&D was the mainstream project in the center, but was facing multiple, difficult challenges, not only technical, but also financial, administrative, political, even before COVID-19. For example, although the development was progressing, it was not clear whether the various project goals and facility constraints would be met. As a national research center for HPC, although outputs were being made in respective application areas due to the use of K, its visibility in terms of academic outputs in the mainstream HPC journals and top conferences such as ACM/IEEE Supercomputing was trailing that of Tokyo Tech. and its partners, e.g. AIST, with significantly smaller budgets. The most lacking were big data, cloud, and AI, and their convergence strategies on Fugaku. There were zero considerations for supporting scalable deep learning, neither in hardware initially, and as a result, in software/frameworks.

Since the design of the Fugaku at the hardware level was in its final stages, hardware tweaks that could be made were minor, and most of the new efforts to accommodate the

deficiencies would have to be made in terms of algorithms and software; but this was in principle in line with the system development approach I had undertaken for the last 30 years, and as a director I formulated several research teams within R-CCS in partnership with our research collaborators, one in big data, one on high performance AI systems that I transferred from Tokyo Tech. initially under the DEEP-CREST budget and a cloud effort lead by the operations teams.

As an example, early tests on the K computer revealed that significant implementation effort had to be conducted to realize deep learning performance that would be competitive to GPUs and latest many-core commercial CPUs. Even before my arrival at Riken, through the architecture committee Riken and I vouched heavily for adding the wide SIMD FP16 vector operations on the A64FX, and at the very last minute in the architecture finalization the feature was added by Fujitsu. Just before my arrival as the R-CCS Director, I initiated a joint development project among multiple development partners including Fujitsu and Arm to exploit the feature, and various libraries such as batched small matrix BLAS for FP16 as well as Arm version of OneDNN was developed, along with other research results to scale deep learning training to large number of nodes, up to 100,000 or more. As a result, Fugaku was able to demonstrate excellent performance in deep learning, achieving 2 Exaflops and #1 in the world in the newly established HPL-AI benchmark, as well as in the new MLPerfHPC benchmark, Fugaku was able to scale the training of the CosmoFlow application to 89,000 nodes, or half of Fugaku, and achieve training speed of about 2x of GPU-based supercomputers hosting 5000 NVIDIA A100 GPUs.

These efforts outlined above are just the tip of the iceberg; there were many research, administrative, political, financial problems that Fugaku faced which had to be resolved to make it not only actually come into being, but also immediately demonstrate its excellence to the global HPC community. One such exemplar was the our reaction to COVID-19, which became a serious pandemic at the beginning of 2021 just when Fugaku installation started and Fujitsu was at the height of its manufacturing. Together with Fujitsu, we devised mitigation measures for supply-chain problems versus installation, early production operation, benchmarking, etc., as well as sustainability strategies even if the engineers are infected and the facility required sanitization quarantine. Such crisis lead to the realization that, Fugaku ramping up quickly despite these difficulties, could in turn be utilized to fight COVID-19. In February I proposed to the MEXT Ministry to utilize early production Fugaku to start this fight, in particular, to conduct translational research of applications that have been developed and optimized for Fugaku, applied such that we could conduct simulated / digital twin studies of the pandemic with compute capacity which would be impossible even for most leadership supercomputers in the world, as they will be totally dominated. A program was quickly set up and a call was made, and in the end 6 projects were selected, from micromolecular studies of bonding strength of different COVID strains, detailed molecular docking studies to investigate transpositional usage of 2,000 pharma substances, to more epidemiological studies of transmissions of virus through aerosols in various societal situations replicated as digital twins, all the way to societal simulations regarding the economic effect of lockdowns. Despite the differences in the field, I organized regular research meetings involving all the teams so that there would be mutual collaboration on the use of Fugaku as well as research input from the

Fugaku side including myself. The program was a great success, especially the aerosol simulation which for the first time provided details of exactly how the viruses which are essentially airborne with aerosols could spread, and quantitative assessment of mitigation measures could be accurately analyzed, which refuted some of the early statements that were made by responsible organization such as WHO that initially stated that masks are not very effective, whereas Fugaku simulation demonstrated that they are very effective in various societal situations. More than 1,000 digital twins were created, simulated and analyzed, and the result was directly reflected in government and industrial anti-COVID countermeasure policies. Also the simulations were visualized and broadcast to the public via media, which educated the public to adopt proper anti-COVID measures. The technical details of the effort was submitted as a ACM Gordon Bell Prize entry, including myself as one of the contributing authors, to the special prize category for COVID-19; among the excellent six finalists chosen, we were able to win the prestigious prize, which was the second Gordon Bell Prize for me since 2011, and proved the worthiness of the research program.

Fugaku of course demonstrated excellent performance as the first supercomputer to exhibit exascale performance over wide range of applications. It was installed and was commissioned earlier than the planned schedule, and did not go over budget due to multitudes of tactics I played out. Immediately after all the racks were installed of the 160,000 node machine, the largest supercomputer ever created, I had Fujitsu and Riken take separate responsibilities for developing and/or tuning for the four major raking benchmarks, Fujitsu being responsible for the Top500 and HPCG as initially planned, while R-CCS research team were deployed to be responsible for HPL-AI and the Graph500 with their own codes. Although it was an enormous logistical challenge to have the installation and field testing of the machine ongoing while we conduct whole machine benchmarks in an unstable environment, the teams pulled through, and Fugaku successfully became world #1 in all the metrics in June 2020, besting the #2 machine by several factors. Fugaku has so far successfully 'defended' its #1 title across all four for the past four editions, namely June / Nov. 2020, and again June / Nov. 2021.

By all means Fugaku was not designed to excel in benchmarks, rather, the design objective was to improve the performance of applications that are important to the society and the humanity, predominantly SDGs applications, such that there would be nearly two orders of magnitude performance improvement over K. Fugaku successfully met the goal, and was accepted at the end of 2020, and for operation to commence early 2021 but still within the Japanese fiscal year 2020, such that the project name 'Flagship 2020' would not be tarnished. Since its launch, Fugaku has been an extremely busy machine despite its size, with utilization hitting maximum continuously.

As is, the entire center with our collaborators is actively conducting research to upgrade Fugaku as the core infrastructure for Society 5.0 initiative of Japan, i.e., achieve digital transformation of the society towards SDGs agenda. For this, along with the commencement of the Fugaku operations in March 2021, I also created the "Society 5.0" office in Tokyo, and at the same time, succeeded in allocating 5% of Fugaku's resources to applications that are collaborative efforts to create SDGs platform on Fugaku accessible

not only by research and academia but also by the industry and local governments. The office will coordinate the various activities, and the first proposals are starting to be submitted, which will make supercomputers to become increasingly indispensable infrastructure for the society as a whole. The widespread recognition of Fugaku and its possible contribution to the society is now ubiquitously recognized by the Japanese public.

Along with leading the current research using Fugaku as the infrastructure and helping to improve it, I have been leading groups of next generation mainstream researchers in HPC to seek the next generation Fugaku or Fugaku NEXT, as well as path forward to progress even beyond that towards the 2030s and beyond. I have launched or helped to launch several projects for this purpose, replicating the timeline evolution of K to Fugaku. One major project I proposed as a PI is the NEDO 100x processor project, in which we investigate various technologies to achieve processor performance that would be 100x that of the one in 2018 in 10 years in 2028, around the time of Fugaku NEXT deployment, and propose a strawman architecture with quantitative assessment to do so. We have mainly recognized the importance of system bandwidth and strong scaling capabilities will be the key to such achievement.

Another effort is the community-driven whitepaper compilation of the NGACI (Next Generation Accelerated Computing Infrastructure), which, much like the SDHPC report for Fugaku, attempts to conduct the initial analysis regarding the broader, more general technical assessments and possibilities towards Fugaku NEXT compared to the NEDO project, and the first draft has been published, effort lead by Masaaki Kondo, one of the team leaders at R-CCS and I am advising the content. Another research effort I have started is to step up the current state of the art in performance models and benchmarking, e.g., create a comprehensive benchmarking and simulation platform involving multitudes of applications, architectures, and analyzed parameters so that we have comprehensive data towards architectural and algorithmic directions we should take, as well as to create new framework for parameterized performance models such that applications can be easily modeled and their performance estimated without having to execute them repeatedly over huge multidimensional parameter space. I am creating a new team dedicated to this research, and in the selection of a new PI candidate to lead the effort.

Such efforts collectively will be further stepped up stating the latter half of 2022, as we have been successful in securing a national budget to conduct a feasibility study for Fugaku NEXT over fiscal year 2022 and 2023. I am already engaging multiple vendors from Japan and internationally to discuss the supercomputing system architectural visions towards 2030, such that they will participate in the study, hopefully leading to Fugaku NEXT which will again demonstrate leapfrogging performance over Fugaku, and will become one of the leadership templates for machines following that, not just large supercomputers but high performance IT infrastructure as a whole. Again, I will be following the successful pattern of identifying the key performance innovators and conducting system HW/SW research there, and rapidly transferring the results to the design of the real system. Hopefully Fugaku will be regarded as the first Post Exascale system, with qualitative differences from the exascale machines that are starting to be deployed today.