

Satoshi Matsuoka

Professor, Dr. Sci. / High Performance Computing Systems Group
Global Scientific Information and Computing Center
& Department of Mathematical and Computing Sciences
Tokyo Institute of Technology
2-12-1 Oo-okayama, Meguro-ku
Tokyo, Japan 205-0021

a. Professional Preparation:

The University of Tokyo, Tokyo, Japan, Information Science, B.S., 1986
The University of Tokyo, Tokyo, Japan, Information Science, M.S., 1988
The University of Tokyo, Tokyo, Japan, Information Science, Ph.D., 1993

b. Appointments:

Visiting Professor, National Institute of Informatics, Tokyo, Japan 2003-Present
Professor, Global Scientific Information and Computing Center, Tokyo Institute of
Technology, Tokyo, Japan 2001-Present
Associate Professor, Mathematical and Computing Sciences, the School of Information
Science and Engineering, Tokyo Institute of Technology, Tokyo, Japan 1996-2001
Assistant Professor, Information Engineering, the School of Engineering, the University
of Tokyo, Tokyo, Japan, 1993-1996
Research Associate, Information Science, the School of Science, the University of Tokyo,
Tokyo, Japan, 1989-1993

c. Five recent publications

Akira Nukada, Hiroyuki Takizawa, Satoshi Matsuoka. “NVCR: A Transparent
Checkpoint-Restart Library for NVIDIA CUDA”, Proc. 20th International
Heterogeneity in Computing Workshop (HCW 2011), Anchorage, Alaska, The IEEE
Press, May, 2011, pp.1-10.

Leonardo Bautista, Akira Nukada, Naoya Maruyama, Franck Cappello, Satoshi Matsuoka.
“Low-overhead checkpoint for large-scale GPU-accelerated systems”, Proc. High
Performance Computing (HiPC 2010), The IEEE Press, Goa, India, Dec. 2010, DOI:
10.1109/HIPC.2010.5713163

Koichi Shirahata, Hitoshi Sato, Satoshi Matsuoka. “Hybrid Map Task Scheduling for
GPU-based Heterogeneous Clusters”, First International Workshop on Theory and
Practice of MapReduce (MAPRED'2010), The IEEE Press, Nov. 2010. Indianapolis,
IN, pp. 733-740, DOI: 10.1109/CloudCom.2010.55.

T. Shimokawabe, T. Aoki, C. Muroi, J. Ishida, K. Kawano, T. Endo, A. Nukada, N.
Maruyama, S. Matsuoka. “An 80-Fold Speedup, 15.0 TFlops Full GPU Acceleration
of Non-Hydrostatic Weather Model ASUCA Production Code”, Proc. IEEE/ACM
Supercomputing (SC2010), New Orleans, LA, Nov. 2010, DOI: 10.1109/SC.2010.9

Hitoshi Nagasaka, Naoya Maruyama, Akira Nukada, Toshio Endo, Satoshi Matsuoka.
“Statistical Power Modeling of GPU Kernels Using Performance Counters”, Proc.
Int'l Green Computing Conference (IGCC'10), Chicago, IL, Aug. 2010, DOI:
10.1109/GREENCOMP.2010.5598315

Five other significant publications

- Leonardo Bautista Gomez, Naoya Maruyama, Franck Cappello, Satoshi Matsuoka. "Distributed Diskless Checkpoint for Large Scale Systems", Proc. 10th IEEE/ACM Cluster, Cloud and Grid Computing (CCGrid 2010), The IEEE Press, Melbourne, Australia, May, 2010, pp.63-72, DOI: 10.1109/CCGRID.2010.40
- Satoshi Matsuoka, Ikuhei Yamagata, Hideyuki Jitsumoto. "Speculative Checkpointing: Exploiting Temporal Affinity of Memory Operations", Proc. HPC Asia 2009, NCHC, Kaohsiung, Taiwan, Mar. 2009, pp.390-396, ISBN: 978-986-85228-0-0
- Hideyuki Jitsumoto, Toshio Endo, Satoshi Matsuoka. "Environmental-Aware Optimization of MPI Checkpointing Intervals", Proc. IEEE Int'l Conf. Cluster Computing (Cluster 2008), The IEEE Press, Tsukuba, Japan, Sep. 2008, pp. 326-329, DOI: 10.1109/CLUSTR.2008.4663790
- Naoya Maruyama and Satoshi Matsuoka. "Model-Based Fault Localization in Large-Scale Computing Systems", In Proc. 22nd IEEE International Parallel & Distributed Processing Symposium (IPDPS 2008), The IEEE Press, Miami, FL, April 2008, pp. 1-12, DOI: 10.1109/IPDPS.2008.4536251
- Hideyuki Jitsumoto, Toshio Endo, Satoshi Matsuoka. "ABARIS: An Adaptable Fault Detection/Recovery Component Framework for MPIs", 12th IEEE Workshop on Dependable Parallel, Distributed and Network-Centric Systems (DPDNS'07), in conj. with IPDPS2007, The IEEE Press, Long Beach, CA, MARCH 2007, pp.1-8, DOI: 10.1109/IPDPS.2007.370603

d. Synergistic Activities

My research is principally in system software for large scale supercomputers and similar infrastructures such as Clouds for HPC. 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, low power, strong scalability, programmability, as well as large-scale I/O. Some of the major projects I have lead are scalable implementations of parallel object-oriented languages on 1000CPU-scale machines in the early 90's (Fujitsu AP1000 and ETL EM-4), GridRPC/Ninf project in collaboration with ETL/AIST (1995-2005), the "Titech Campus Grid" project (2002-2006) which deployed a 1300 CPU grid-of-clusters in production mode within the Tokyo Tech campus, the NAREGI project (2003-2007) which is a national project to develop a deployable advanced grid middleware for national supercomputing centers, and Tsubame1 (2006-2010), a production supercomputer which was the #1 machine in Asia for 1.5 years. Some of the recent projects include Ultra Low Power HPC (2007-2013) which aims to achieve 1000-fold power/performance improvement in supercomputers, Info-Plosion (2005-2011) which aims to develop basic system software technologies for large-scale data and information processing, Tsubame2.0 (2010-2014) which became the 4th fastest supercomputer in the world and the "Greenest Production Supercomputer in the World" on the Green 500 two consecutive times, heavily utilizing GPUs and other power efficient technologies. The most recent award is the Billion-way fault tolerancy in supercomputers (2011-2015) which will investigate the key technologies for reliable exascale in supercomputing.

e. Collaborators

H. Nakada, A. Yonezawa, S. Sekiguchi, T. Endo, M. Sato, N. Maruyama, T. Suzumura, O. Tatebe, S. Takahashi, K. Taura, U. Nagashima, H. Casanova, Y. Morita, T. Boku, T. Igarashi, H. Jitsumoto, H. Sato, H. Harada, H. Tanaka, H. Masuhara, K. Miyashita, A. Takefusa, D. Takahashi, Y. Soda, Y. Sakae, A. Nukada, T. Kamada, Y. Tanaka, H. Takagi, K. Seymore, K. Shimura, Y. Hotta, L. Baduel, H. Hosobe, J. Rekimoto, T. Watanabe, K. Asai, M. Yasugi, S. Kawachiya, S. Itou, S. Takizawa, Y. Kimura, Y. Takamiya, H. Nakashima, Y. Watase, Y. Ichisugi, Y. Ayatsuka, J. Dongarra, S. Kawai, K. Aida, F. Maruyama, T. Chiba, Y. Machida, K. Wakita, H. Nakamura, F. Cappello, , T. Aoki, P. Arzberger, T. Nakagawa, S. Ogura, K. Tanaka, K. Saga, K. Sato, Y. Hosogaya, A. Cevahir, T. Hamano, T. Kamada, H. Hasegawa, M. Aoyagi,, C. Lee, Y. Ishikawa, D. Katz, Y. Kodama, F. Berman (200 others).

Graduate students

Current: Sumeth Lerthirunwong, Mohamed Amin Jabri, Leonardo A. Bautista Gomez, Irina Demeshko, Aleksandr Drozd, Jin Guang-hao, Kento Sato Amer Abdelhalim, Koichi Shirahata, Nguyen Toan, Keisuke Fukuda, Takafumi Saito, Zhang Jiayue. (Have advised over 30 grad. students and 10 postdoc scientists while at Tokyo Tech.)

Graduate advisor

Satoru Kawai, The University of Tokyo
Akinori Yonezawa, The University of Tokyo

Satoshi Matsuoka received his Bachelor's, Master's, and Ph.D. respectively in Information Science from the University of Tokyo. After being a Research Associate and Lecturer at Information Science and Information Engineering Departments of University of Tokyo respectively, he became an Associate Professor at the Dept. of Mathematical and Computing Sciences, and five years later a Full professor at the Global Scientific Information and Computing Center (GSIC) of Tokyo Institute of Technology (Tokyo Tech), one of the best universities in Japan in the area of science and technology.

He had worked principally on software systems for high-performance computing on advanced infrastructural platforms such as macro dataflow machines, large-scale MPP supercomputers, large clusters, computational and data grids, and heterogeneous GPU/CPU supercomputers such as the latest petascale Tsubame2.0

Early on at University of Tokyo starting in 1989 he worked as an Assistant Professor in Professor Akinori Yonezawa's group, leading the efforts to implement concurrent and parallel object-oriented languages such as ABCL (Actor-Based Concurrent Language, designed by Prof. Yonezawa's group) efficiently on massive parallel machines of the time, such as the 64 to 1024 processor AP1000, as well as macro dataflow machines such as the ETL-EM/4. These were actually one of the very few actual parallel languages of the day that ran on supercomputers at those scales. He also worked on language features to make the resources explicit as first-class objects using reflective techniques. Also, he conducted definitive work identifying the source of anomaly occurring in inheritance when concurrency is introduced to object-oriented languages, and coined the term "inheritance anomaly".

After having transferred internally to the Information Engineering department in 1993, and continuing with the efforts of utilization and adaptation of object-oriented languages such as Java and C++ on real parallel code, he moved to Tokyo Tech in 1996 as an Associate Professor, and started work on what was to be called grid computing, as well as developing commodity clusters so that they would scale to large configurations matching that of dedicated MPPs, as well as in a heterogeneous environment. He co-lead many collaborative projects on the emerging grid computing, as one of the pioneers in Japan, such as the Ninf project as well as various grid scheduling and simulations work. Also, some of the prototype clusters that were built in his lab were ranked on the Top500, as high as 47th in the world, much faster than the supercomputers that were in production at the computing center of Tokyo Tech. at the time.

Such efforts lead to his being promoted to full professor in 2000 to head the Research Infrastructure Division of GSIC, being responsible for Tokyo Tech's supercomputing infrastructure. During 2002-2006 he led the Titech Campus Grid Project that hosted up to 1,300-processor "Grid of Clusters" in actual production mode throughout the Tokyo Tech. campus, alongside the dedicated production supercomputers. Also, during 2003-2008 he served as a sub-leader of the Japanese National Research Grid Initiative (NAREGI) project, which was a \$100 million national project in creating middleware and applications for next-generation Japanese Cyber-Science (Grid) infrastructure, and made collaborations with various international grid initiatives such as the TeraGrid in the U.S., and EGEE, and DEISA in the EU.

At the same time he led the Tokyo Tech's TSUBAME supercomputer project, which resulted in the fastest supercomputer in Asia-Pacific starting from June 2006 for three consecutive Top 500s. Based on the experiences, he then lead the design and deployment of TSUBAME2.0, which became the first Petaflops supercomputer in Japan in November 2010 at 2.4 Petaflops peak and 1.192 Petaflop Linpack benchmark performance, being ranked 4th fastest supercomputer in the world. TSUBAME2.0 was also ranked 2nd in the world for power efficiency on the Green 500 in November, and was awarded the "Greenest Production Supercomputer in the World" award. TSUBAME2.0 became 5th fastest in the world and re-awarded the "Greenest Production" in June, 2011. TSUBAME makes heavy use of GPUs to achieve its performance excellence, and is a forerunner in heterogeneous computing architecture that combines many-core "throughput" processors with smaller number of

multi-core “latency processors” to cope with the scalability and power challenges leading to exascale and beyond. Not only is TSUBAME2.0 exhibiting excellence in benchmarks, but also in real application code with scientific significance---petascale applications running on TSUBAME2.0 were selected to be two of the five ACM Gordon Bell Prize finalists for 2011.

Currently, he is also the head of a research lab of nearly 25 researches and graduate students at the Dept. of Mathematical and Computing Sciences at Tokyo Tech. and plays a key role in the Japanese HPCI (High-Performance Computing Infrastructure), whose aim is to actually deploy the vision of NAREGI centered around the K supercomputer and a set of multi-petaflop storage pool.

His main recent research interests are software issues in Peta/Exaflop supercomputers including scalability and resilience, green computing for HPC, programming on millions of cores, and dealing with I/O bottlenecks in peta and exascale storage, and the use of Clouds for e-Science. His recent major research projects include being a deputy leader for Info-Plosion (2006-2011), a \$40mil inter-university project to deal with explosion of information and data. He is a project PI for ULPHPC (Ultra Low-Power HPC) that aims to drastically reduce the power consumption of HPC systems by a factor of 1000 over 10 years, and its successor Green Supercomputing that strives for intricate energy control and to almost nullify the need for cooling in supercomputers even in hotter climates, Billion Core Fault Tolerance project whose goal is to come up with a set of tools and methodologies for exascale resilience. He is also collaborating internationally in the French-Japan JST-ANR FP3C project for post petascale computing, as well as the G8 project for exascale climate research. The fruits of such research should pave the way for TSUBAME3.0 in 2014/15 timeframe, which will sport a design scalable to 100 Petaflops and beyond.

He has published over 200 refereed publications, many in top-tier conferences and journals, as well as has won several awards including the prestigious JSPS Prize from the Japan Society for Promotion of Science in 2006 from His Royal Highness Prince Akishinomiya. He was also selected as one of the HPC Wire’s “People to Watch 2010”, the only Japanese ever to be chosen so.

Not only that he has served on a few hundred program committee roles in prestigious conferences, he has played key chairmanship roles in a number of international conferences, such as the Program Chairs of ACM OOPSLA 2002 as well as IEEE CCGrid 2003, and the Technical Paper Chair of IEEE/ACM Supercomputing 2009 (SC09), the latter gathering over 10,000 participants worldwide from academia to industry. In 2011 he is serving as the Community Chair for SC11, and in 2013 the Technical Program Chair for SC13. He has served various other board and leadership roles internationally including the secretariat of ACM Japan Chapter, steering group member of the Open Grid Forum, as well as the steering group member for Supercomputing during 2009-2012, and a Fellow of the International Supercomputing Conference in Europe. He is also the main leader for GSIC to be recognized as the Microsoft’s Technical Center of Excellence, as well as NVIDIA’s CUDA Center of Excellence, both for the first time for a Japanese university.