Balazs Gerofi, Ph.D

Advanced Institute for Computational Science (AICS) – RIKEN
+81-80-2094-4286 • bgerofi@il.is.s.u-tokyo.ac.jp, bgerofi@riken.jp
http://www-sys-aics.riken.jp/Members/gerofi.html
Address: Toranomon Twin Bldg. East 16th Floor, 2-10-1 Toranomon, 105-0001, Tokyo, Japan

Education

The University Of Tokyo, Grad. School of Information Science and Technology — **Tokyo, JP**
Ph.D., Information Science and Technology — *Apr.2009–Mar.2012*

Vrije Universiteit, Faculty of Sciences, Dept. of Computer Science — **Amsterdam, NL**
M.Sc. (Cum Laude), Computer Science — *May.2005–Oct.2006*

Eotvos Lorand University, Faculty of Informatics, Dept. of Information Systems — **Budapest, HU**
M.Sc., Mathematician - Program Designer — *Sep.1999–Apr.2005*

Appointments and Working Experience

Intel Corporation — **Hillsboro, OR, USA**
Visiting Researcher — *2016.4–2016.7*
mOS Research Group

Intel Corporation — **Hillsboro, OR, USA**
Visiting Researcher — *2015.1–2015.3*
mOS Research Group

Advanced Institute for Computational Science (AICS), RIKEN — **Tokyo, JP**
Research Scientist — *2014.11–*
Exascale System Software Team

The University Of Tokyo — **Tokyo, JP**
Assistant Professor — *2013.9–2014.10*
Grad. School of Information Science and Technology

Advanced Institute for Computational Science (AICS), RIKEN — **Kobe, JP**
Postdoctoral Researcher — *2012.4–2013.8*
System Software Research Team

Argonne National Laboratory (ANL) — **Chicago, IL, USA**
Summer Intern — *2010*
Mathematics and Computer Science (MCS) Division

TomTom International BV — **Amsterdam, NL**
Linux Kernel Developer — *2007–2008*
System Integration Group

Awards & Scholarships

Best Paper Award at ACM ROSS'17 **Washington, D.C., USA**
Paper: Toward Full Specialization of the HPC Software Stack *Jun.2017*

Best Paper Award at ACM ROSS'15 **Portland, USA**
Paper: Exploring the Design Space of Combining Linux with Lightweight Ker... *Jun.2015*

Best Paper Award at ACM HPDC'14 **Vancouver, Canada**
Paper: CMCP: A Novel Page Replacement Policy for System Level Hierarchical... *Jun.2014*

Nomination for Best Paper Award at ACM/IEEE CCGRID'13 **Delft, Netherlands**
Paper: Partially Separated Page Tables for Efficient Operating System Assisted Hier... *May.2013*

Ph.D. Thesis Special Distinction (研究科賞授与) **Tokyo, JP**
by the Grad. School of Information Science and Technology, The Univeristy of Tokyo *Apr.2012*

Best Young Researcher's Paper Award at the IPSJ Computer System Symposium **Tokyo, JP**
Paper: Live Migration of Processes Maintaining Multiple Network Connections *Nov.2009*

MEXT Monbukagakusho Japanese Government Scholarship (文部科学省奨学金) **Tokyo, JP**
University of Tokyo *Apr.2008*

Dutch Government Computer Science Master's Scholarship **Amsterdam, NL**
Vrije Universiteit Amsterdam *May.2005*

Excellent Student of the Faculty of Informatics Award **Budapest, HU**
Eotvos Lorand University *Apr.2005*

Third prize at the Hungarian Scientific Competition for Students (TDK) **Budapest, HU**
Paper: Hidden Markov Model finds Behavioral Patterns of User Working with a Voicemouse *Mar.2005*

Research Interests

My research is mainly focused on system software and parallel / distributed computing. In particular, I am interested in operating system (kernel architectures for many-core CPUs, memory management, file systems), high-performance computing (parallel and distributed I/O, resiliency), virtualization (virtual machines, cloud-computing), and fault tolerant computing (replication, checkpoint-restart, message-logging).

Teaching

The University Of Tokyo, Graduate School of Information Science and Technology **Tokyo, JP**
Computer Architecture and System Software *Fall 2016*

The University Of Tokyo, Graduate School of Information Science and Technology **Tokyo, JP**
Introduction to PowerPC Assembly *Fall 2013*

The University Of Tokyo, Graduate School of Information Science and Technology **Tokyo, JP**
Practical System Programming *Fall 2013*

The University Of Tokyo, Graduate School of Information Science and Technology **Tokyo, JP**
Practical System Programming *Fall 2011*

The University Of Tokyo, Graduate School of Information Science and Technology **Tokyo, JP**
Practical System Programming *Fall 2009*

Languages

English: Fluent

Japanese: Conversational

German: Basic

Hungarian: Native

Thesis

Degree: Ph.D Thesis at The University of Tokyo

Title: ***Efficient Replication Mechanisms for Highly Available Virtual Machines***

Supervisor: Prof. Yutaka Ishikawa

URL: http://www.il.is.s.u-tokyo.ac.jp/∼bgerofi/bgerofi-phdthesis.pdf

Degree: Master's Thesis at the Vrije Universiteit Amsterdam

Title: ***Design and Implementation of the MINIX Virtual File System***

Supervisor: Prof. Andrew S. Tanenbaum

URL: http://www.minix3.org/theses/gerofi-minix-vfs.pdf

Selected Publications

B. Gerofi, R. Riesen, M. Takagi, T. Boku, R. W. Wisniewski, and Y. Ishikawa, "**Performance and Scalability of Lightweight Multi-Kernel based Operating Systems**," in *Processing of IEEE International Parallel and Distributed Processing Symposium*, ser. IPDPS '18, May 2018.

B. Gerofi, R. Riesen, R. W. Wisniewski, and Y. Ishikawa, "**Toward Full Specialization of the HPC Software Stack: Reconciling Application Containers and Lightweight Multi-kernels**," in *Proceedings of the 7th International Workshop on Runtime and Operating Systems for Supercomputers ROSS 2017*, ser. ROSS '17. New York, NY, USA: ACM, 2017, pp. 7:1–7:8. [Online]. Available: http://doi.acm.org/10.1145/3095770.3095777

T. Miyoshi, G. Y. Lien, S. Satoh, T. Ushio, K. Bessho, H. Tomita, S. Nishizawa, R. Yoshida, S. A. Adachi, J. Liao, B. Gerofi, Y. Ishikawa, M. Kunii, J. Ruiz, Y. Maejima, S. Otsuka, M. Otsuka, K. Okamoto, and H. Seko, "**Big Data Assimilation: Toward Post-Petascale Severe Weather Prediction: An Overview and Progress**," *Proceedings of the IEEE*, vol. PP, no. 99, pp. 1–25, 2016.

S. Perarnau, J. A. Zounmevo, B. Gerofi, K. Iskra, and P. Beckman, "**Exploring Data Migration for Future Deep-Memory Many-Core Systems**," in *Proceedings of IEEE International Conference on Cluster Computing*, ser. CLUSTER '16, Sept. 2016.

J. Liao, B. Gerofi, G.-Y. Lien, S. Nishizawa, T. Miyoshi, H. Tomita, and Y. Ishikawa, "**Toward a General I/O Arbitration Framework for netCDF based Big Data Processing**," in *Proceedings of the Euro-Par International European Conference on Parallel Processing*, ser. EuroPar '16, Aug. 2016.

B. Gerofi, M. Takagi, G. Nakamura, T. Shirasawa, A. Hori, and Y. Ishikawa, "**On the Scalability, Performance Isolation and Device Driver Transparency of the IHK/McKernel Hybrid Lightweight Kernel**," in *Processing of IEEE International Parallel and Distributed Processing Symposium*, ser. IPDPS '16, May 2016, pp. 1041–1050.

B. Gerofi, M. Takagi, and Y. Ishikawa, "**Revisiting RDMA Registration in the Context of Lightweight Multi-kernels**," in *Proceedings of the European MPI Users' Group Meeting*, ser. EuroMPI '16, 2016 (short paper).

J. Liao, F. Trahay, B. Gerofi, and Y. Ishikawa, "**Prefetching on Storage Servers through Mining Access Patterns on Blocks**," *IEEE Transactions on Parallel and Distributed Systems*, vol. PP, no. 99, pp. 1–1, 2015.

B. Gerofi, M. Takagi, and Y. Ishikawa, "**Toward Operating System Support for Scalable Multithreaded Message Passing**," in *Proceedings of the 22Nd European MPI Users' Group Meeting*, ser. EuroMPI '15. Bordeaux, France: ACM, 2015, pp. 1:1–1:10. [Online]. Available: http://doi.acm.org/10.1145/2802658.2802661

R. Riesen, A. B. Maccabe, B. Gerofi, D. N. Lombard, J. J. Lange, K. Pedretti, K. Ferreira, M. Lang, P. Keppel, R. W. Wisniewski, R. Brightwell, T. Inglett, Y. Park, and Y. Ishikawa, "**What is a Lightweight Kernel?**" in *Proceedings of the 5th International Workshop on Runtime and Operating Systems for Supercomputers*, ser. ROSS '15. New York, NY, USA: ACM, 2015, pp. 9:1–9:8. [Online]. Available: http://doi.acm.org/10.1145/2768405.2768414

B. Gerofi, M. Takagi, Y. Ishikawa, R. Riesen, E. Powers, and R. W. Wisniewski, "**Exploring the Design Space of Combining Linux with Lightweight Kernels for Extreme Scale Computing**," in *Proceedings of the 5th International Workshop on Runtime and Operating Systems for Supercomputers*, ser. ROSS '15. New York, NY, USA: ACM, 2015, pp. 5:1–5:8. [Online]. Available: http://doi.acm.org/10.1145/2768405.2768410

T. Shimosawa, B. Gerofi, M. Takagi, G. Nakamura, T. Shirasawa, Y. Saeki, M. Shimizu, A. Hori, and Y. Ishikawa, "**Interface for Heterogeneous Kernels: A Framework to Enable Hybrid OS Designs targeting High Performance Computing on Manycore Architectures**," in *Proceedings of the IEEE International Conference on High Performance Computing*, ser. HiPC '14, Dec. 2014.

B. Gerofi, T. Masamichi, and Y. Ishikawa, "**Exploiting Hidden Non-uniformity of Uniform Memory Access on Manycore CPUs**," in *Proceedings of the International Workshop on Multi/Many-Core Computing Systems, held in conjunction with Euro-Par International European Conference on Parallel Processing*, ser. MuCoCoS '14, Aug. 2014.

Y. Soma, B. Gerofi, and Y. Ishikawa, "**Revisiting Virtual Memory for High Performance Computing on Manycore Architectures: A Hybrid Segmentation Kernel Approach**," in *Proceedings of the International Workshop on Runtime and Operating Systems for Supercomputers (ROSS), held in conjunction with ACM/SIGARCH International Conference on Supercomputing (ICS)*, ser. ROSS '14, Jun. 2014.

B. Gerofi, A. Shimada, A. Hori, T. Masamichi, and Y. Ishikawa, "**CMCP: A Novel Page Replacement Policy for System Level Hierarchical Memory Management on Many-cores**," in *Proceedings of the ACM International Symposium on High-Performance Parallel and Distributed Computing*, ser. HPDC '14, Jun. 2014.

M. Takagi, Y. Nakamura, A. Hori, B. Gerofi, and Y. Ishikawa, "**Revisiting Rendezvous Protocols in the Context of RDMA-capable Host Channel Adapters and Many-Core Processors**," in *Proceedings of 20th European MPI Users' Group Meeting*, ser. EuroMPI '13, 2013.

A. Tokuhisa, J. Arai, Y. Joti, Y. Ohno, T. Kameyama, K. Yamamoto, M. Hatanaka, B. Gerofi, A. Shimada, M. Kurokawa, F. Shoji, K. Okada, T. Sugimoto, M. Yamaga, R. Tanaka, M. Yokokawa, A. Hori, Y. Ishikawa, T. Hatsui, and N. Go, "**High-speed classification of coherent X-ray diffraction patterns on the K computer for high-resolution single biomolecule imaging**," in *International Union of Crystallography (IUCr) Journal of Synchrotron Radiation*, ser. IUCrJSR '13, 2013.

A. Shimada, B. Gerofi, A. Hori, and Y. Ishikawa, "**Proposing a new Task Model towards Many-core Architecture**," in *International Workshop on Many-core Embedded Systems*, ser. MES '13, 2013.

B. Gerofi, A. Shimada, A. Hori, and Y. Ishikawa, "**Partially Separated Page Tables for Efficient Operating System Assisted Hierarchical Memory Management on Heterogeneous Architectures**," in *Proceedings of the ACM/IEEE International Symposium on Cluster, Cloud and Grid Computing*, ser. CCGRID '13, May. 2013.

B. Gerofi, A. Shimada, A. Hori, and Y. Ishikawa, "**Towards Operating System Assisted Hierarchical Memory Management for Heterogeneous Architectures**," in *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '12 (poster paper), Nov. 2012.

B. Gerofi, A. Hori, and Y. Ishikawa, "**clone_n(): Parallel Thread Creation for Upcoming Many-Core Architectures**," in *Proceedings of IEEE International Conference on Cluster Computing*, ser. CLUSTER '12 (poster paper), Sept. 2012, pp. 592 –596.

A. Shimada, B. Gerofi, A. Hori, and Y. Ishikawa, "**PGAS Intra-node Communication towards Many-Core Architecture**," in *Proceedings of the International Conference on Partitioned Global Address Space Programming Models*, ser. PGAS '12, 2012.

B. Gerofi, Z. Vass, and Y. Ishikawa, "**Utilizing Memory Content Similarity for Improving the Performance of Highly Available Virtual Machines**," in *ELSEVIER International Journal of Future Generation Computer Systems*, ser. FGCS '12, 2012.

B. Gerofi and Y. Ishikawa, "**Enhancing TCP Throughput of Highly Available Virtual Machines via Speculative Communication**," in *Proceedings of the ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments*, ser. VEE '12, 2012.

B. Gerofi and Y. Ishikawa, "**Workload Adaptive Checkpoint Scheduling of Virtual Machine Replication**," in *Proceedings of the IEEE Pacific Rim International Symposium on Dependable Computing*, ser. PRDC '11, 2011.

B. Gerofi and Y. Ishikawa, "**Utilizing Memory Content Similarity for Improving the Performance of Replicated Virtual Machines**," in *Proceedings of the ACM/IEEE International Conference on Utility and Cloud Computing*, ser. UCC '11, 2011.

B. Gerofi and Y. Ishikawa, "**RDMA based Replication of Multiprocessor Virtual Machines over High-Performance Interconnects**," in *Proceedings of the 2011 IEEE International Conference on Cluster Computing*, ser. CLUSTER '11, 2011, pp. 35–44.

B. Gerofi and Y. Ishikawa, "**A Multi-core Approach to Providing Fault Tolerance for Non-deterministic Services**," in *Network Computing and Applications (NCA), 2010 9th IEEE International Symposium on*, july 2010, pp. 233 –238.

B. Gerofi, H. Fujita, and Y. Ishikawa, "**An Efficient Process Live Migration Mechanism for Load Balanced Distributed Virtual Environments**," in *Cluster Computing (CLUSTER), 2010 IEEE International Conference on*, sept. 2010, pp. 197 –206.

B. Gerofi, H. Fujita, and Y. Ishikawa, "**Live Migration of Processes Maintaining Multiple Network Connections**," *IPSJ Transactions on Advanced Computing Systems (ACS 29)*, vol. 3, no. 1, p. 8, 2010.

G. Hévízi, B. Gerofi, B. Szendrő, and A. Lorincz, "**Assisting robotic personal agent and cooperating alternative input devices for severely disabled children**," in *Proceedings of the 4th International Central and Eastern European conference on Multi-Agent Systems and Applications*, ser. CEEMAS'05. Berlin, Heidelberg: Springer-Verlag, 2005, pp. 591–594. [Online]. Available: http://dx.doi.org/10.1007/11559221_68

Professional Society Memberships

I am a member of the IEEE Computer Society, the Association for Computing Machinery (ACM) and the Information Processing Society of Japan (IPSJ).