# Roberto Gioiosa

*Post-Doc Researcher*

*Barcelona Supercomputing Center*
*C/ Jordi Girona, 31*
*80034 Barcelona - Spain*
✉ *roberto.gioiosa@bsc.es*
*www.sprg.uniroma2.it/home/gioiosa*


## Research interests

Operating systems, high performance computing, programming models, reliability, heterogeneous systems, large parallel clusters, real-time systems, and embedded systems.

## Education

Jul 2008 **HiPEAC ACACES 2008 Summer school**, L'Aquila, Italy.

Jul 2007 **HiPEAC ACACES 2007 Summer school**, L'Aquila, Italy.

2002–2006 **PhD, Computer Science Engineering**, *University of Rome "Tor Vergata"*, Rome, Italy.

title *High performance computing clusters*

supervisors Prof. D.P. Bovet

description Modern clusters are built with COTS hardware components and run open-source software. These new configurations rise new problems in therms of performance, reliability and efficiency. This thesis addresses two problems of modern HPC systems: fault tolerance and OS noise. TICK is a new solution for transparent, low-overhead, kernel level fault tolerance. TICK is able to checkpoint parallel applications and automatically restart on a different node processes that have experienced a fault. The thesis also presents a study of the impact of OS noise on large HPC systems A first analysis shows how system performance are tightly sensible to uncoordinated events, such as interrupt handling. Even short activities may result in lack of performance at scale if uncoordinated. A global scheduler (coordinator) is proposed to reduce the effect of those activities that cannot be removed.

Oct 2002 **Certified Professional Engineer (PE) in Computer Engineering**, Rome, Italy.

1997–2002 **Master of Science, Computer Science Engineer**, *University of Rome "Tor Vergata"*, Rome, Italy, *Summa cum laude*.

title *Asymmetric kernels for multiprocessor systems*

supervisors Prof. D.P. Bovet, Dr. M. Cesati

description Multiprocessor systems, especially those based on multi-core or multi-threaded processors, and new operating system architectures can satisfy the ever increasing computational requirements of embedded systems. ASMP-Linux is a modified, high responsiveness, open-source, hard real-time OS for multiprocessor systems capable of providing high real-time performance while maintaining the code simple and not impacting on the performances of the rest of the system. Moreover, ASMP-Linux does not require code changing or application re-compiling/re-linking.

## Experience

### Research activities

2009–current **Post-Doc**, *Barcelona Supercomputing Center*, Barcelona, Spain.
Design and implementation of OS for exascale systems focusing on scheduling, power and thermal considerations; OS noise analysis and reduction; memory management; parallel programming model for future applications. Power, performance and thermal modeling for HPC systems. Performance analysis of HPC benchmarks and design of specific synthetic benchmarks for future generation processors. Virtual machine solutions for embedded real time systems.

Oct'10–Feb'11 **Visiting Scholar**, *University of Rochester*, Rochester, NY, US.
Accelerator-based computing for multimedia and real time systems. Resilient memory study.

2008–2009 **Post-Doc**, *IBM Research - Watson*, Yorktown Heights, NY, US.
Design, development and testing of the OS for new generation IBM Blue Gene systems. Performance analysis of HPC applications...

2006–2008 **Post-Doc**, *Barcelona Supercomputing Center*, Barcelona, Spain.
Research grant for the design and implementation of operating systems for multi-core/multi-threaded processors used in HPC systems, performance analysis of HPC applications; the research activity also focuses on real time systems. Involved in projects with IBM and SUN, MERASA FP7 European project; technical coordinator at BSC for HiPEAC2 European Network of Excellence.

2005–2006 **Researcher**, *University of Rome "Tor Vergata"*, Rome, Italy.
Research grant for the design and implementation of OS components and drivers for real time, military certified (DO178B) OSes for embedded systems (Linux, VxWorks). The research activity also focused on system overhead analysis (both Hardware and OS); global and local job scheduling across a computer cluster; transparent, low-overhead fault tolerant mechanism at kernel level and multi-processors/multi-core/multi-threaded systems analysis and scheduler design.

2004–2005 **Graduate student**, *Los Alamos National Laboratory (LANL)*, Los Alamos, NM, USA.
Research grant for the design and implementation of new features for operating system for HPC clusters. The research activity was focused on system overhead analysis (both hardware and operating system), global and local job scheduling across a computer cluster and transparent, low-overhead fault tolerant mechanism at kernel level (TICK).

### Teaching activities

May-Jul 2006 **Advanced Operating systems**, *Teacher Assistant*, University of Rome "Tor Vergata", Rome, Italy.

Mar-Jun 2006 **Linux Kernel Hacking 2006**, *Lecturer*, University of Rome "Tor Vergata", Rome, Italy.

Oct-Dec 2005 **Operating systems**, *Teacher*, University of Rome "Tor Vergata", Rome, Italy.

Oct-Dec 2003 **Operating systems**, *Teacher Assistant*, University of Rome "Tor Vergata", Rome, Italy.

Mar-Jun 2003 **Linux Kernel Hacking 2003**, *Lecturer*, University of Rome "Tor Vergata", Rome, Italy.

Oct-Dec 2002 **Operating systems**, *Teacher Assistant*, University of Rome "Tor Vergata", Rome, Italy.

### Professional activities

Jan-Jun 2006 **Consultant**, *Quadrics Italy*, Rome, Italy.
Development of a device driver for an ATA PCI (PMC) disk flash for VxWorks

Feb-Jun 2003 **Developer**, *Urmet TLC*, Rome, Italy.
Development of software applications for Banking Transactions on embedded system based on Hitachi H8/300 processor series

### Services

**Workshop organization** HPPAC'12, CAOS'12, HPPAC'11, CAOS'11, MULTIPROG'11, CAOS'10, MULTIPROG'10, MULTIPROG'09, MULTIPROG'08

**Program track chair** VLSI-SoC'10

**Program committee** CF'12, IA$^3$'12, ICS'11, CF'11, VLSI-SoC'11

**Publication Chair** CF'10

**Reviewer** Spinger JCC, CAL, VLSI-SoC'11, SYSTOR'11, ISCA'10, HPCA'10, IPDPS'10, HiPEAC'10, MTAAP'10, TCAD'10, ITNG'10, ICCAD'09, ISCA'08, MTAAP'08, IPDPS'07, ICS'07, SP&E

## Honor and Awards

2008 **HiPEAC best paper awards**, *HiPEAC*.

1997 **Honored graduate**, *I.T.C "G. Di Vittorio"*.

## Languages

Italian **Native speaker**

English **Excellent** *Read, written, spoken*

Spanish **Excellent** *Read, written, spoken*

## Computer skills

| | | | |
|---|---|---|---|
| **Operating systems** | Linux, CNK, WR VxWorks, Solaris, Windows, MS-DOS | **Scripting languages** | Pyton, bash, PHP |
| **Programming languages** | C/C++, x86/POWER/MIPS assembler, Java | **Miscellaneous** | Matlab, Latex |
| **Parallel paradigms** | MPI, OpenMP, pThreads, Transactional Memory | | |

## Invited talks

Sep 2011 **System software challanges in the exascale era**, *University of Siena*, Siena, Italy.

May 2011 **Improving the efficiency of HPC applications**, *Pacific Northwest National Laboratory*, Richland, WA, US.

Nov 2010 **Improving efficiency of HPC systems through dynamic HW resource allocation**, *University of Rochester*, Rochester, NY, US.

Apr 2010 **Balancing HPC applications through HW/SW co-design**, *Massachusetts Institute of Technology (MIT)*, Cambridge, MA, US.

Jan 2010 **The Many-Core Era: a look at the future**, *University of Rome "Tor Vergata"*, Rome, Italy.

Mar 2009 **A new Linux scheduler for balancing HPC applications on a IBM POWER5 processor**, *University of Illinois at Urbana-Champaign (UIUC)*, Urbana, Illinois, US.

May 2008 **Hardware/Software co-design for future HPC systems**, *IBM Research TJ Waston*, Yorktown Heights, NY, US.

Oct 2007 **Hard Real-Time Performance in Multiprocessor Systems using ASMP-Linux**, *University of Illinois at Urbana-Champaign (UIUC)*, Urbana, Illinois, US.

Apr 2006 **High Performance Computing Clusters**, *Barcelona Supercomputing Center (BSC)*, Barcelona, Spain.

Jun 2005 **Analysis of System Overhead on Parallel Computers**, *University of Salerno*, Fondo Per gli Investimenti della Ricerca di Base 2005 meeting (FIRB05), Salerno, Italy.

Jun 2003 **Linux Real-Time: an Asymmetric Approach**, *University of Naples "Federico II"*, Fondo Per gli Investimenti della Ricerca di Base 2003 meeting (FIRB03), Naples, Italy.

Nov 2002 **Asymmetric Multiprocessor Kernel**, *University of Rome "Tor Vergata"*, The Linux Day 2002, Rome, Italy.

## Publications

A. Morari, R. Gioiosa, R. W. Wisniewski, B. S. Rosenburg, T. A. Inglett, and M. Valero. Evaluating the impact of tlb misses on future hpc systems. In *26th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, May 2012. Best paper award.

G. Kestor, R. Gioiosa, O. Unsal, A. Cristal, and M. Valero. Hardware/software techniques for assisted execution runtime systems. In *The 2nd Workshop on Runtime Environments, Systems, Layering and Virtualized Environments (RESoLVE)*, March 2012.

A. Morari, C. Boneti, F. Cazorla, R. Gioiosa, C. Cher, A. Buyuktosunoglu, P. Bose, and M. Valero. SMT malleability in IBM POWER5 and POWER6 processors. *IEEE Transaction on Computers*, To appear 2012.

C. Luque, M. Moreto, F. J. Cazorla, R. Gioiosa, A. Buyuktosunoglu, and M. Valero. Cpu accounting for multicore processors. *To appear in IEEE Transaction on Computers*, February 2012.

G. Kestor, R. Gioiosa, T. Harris, A. Cristal, O. Unsal, M. Valero, and I. Hur. STM2: A parallel stm for high performance simultaneous multi-threading systems. In *The 20th IEEE International Conference on Parallel Architectures and Compilation Techniques (PACT)*, October 2011.

V. Jimenez, F. Cazorla, R. Gioiosa, M. Valero, C. Boneti, E. Kursun, C. Cher, C. Isci, A. Buyuktosunoglu, and P. Bose. Characterizing power and temperature behavior of power6-based system. (invited paper). *IEEE Journal of Emerging and Selected Topics in Circuits and Systems*, September 2011.

V. Jimenez, F. Cazorla, R. Gioiosa, E. Kursun, C. Isci, C. A. Buyuktosunoglu, P. Bose, and M. Valero. A case for energy-aware accounting and billing in large-scale computing facilities cost metrics and design implications. *IEEE Micro*, May/June 2011.

A. Morari, R. Gioiosa, R. W. Wisniewski, F. J. Cazorla, and M. Valero. A quantitative analysis of os noise. In *25th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, May 2011.

R. Gioiosa. Towards sustainable exascale computing. In *The 18th IEEE/IFIP VLSI System on Chip Conference (VLSI-SoC)*, September 2010.

V. Jimenez, R. Gioiosa, E. Kursun, F. Cazorla, C. Cher, A. Buyuktosunoglu, P. Bose, and M. Valero. Trends and techniques for energy efficient architectures. In *The 18th IEEE/IFIP VLSI System on Chip Conference (VLSI-SoC),*, September 2010.

R. Gioiosa, S. A. McKee, and M. Valero. Designing os for hpc applications: Scheduling. In *The 2010 IEEE International Conference on Cluster Computing (CLUSTER)*, September 2010.

V. Jimenez, C. Boneti, F. Cazorla, R. Gioiosa, E. Kursun, C. Cher, C. Isci, A. Buyuktosunoglu, P. Bose, and M. Valero. Power and thermal characterization of power6 system. In *The 19th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, September 2010.

B. Goel, S. A. McKee, R. Gioiosa, K. Singh, M. Bhadauria, and M. Cesati. Portable, scalable, per-core power estimation for intelligent resource management. In *The 1st Green Computing Conference, 2010 International (IGCC)*, August 2010.

V. Jimenez, F. Cazorla, R. Gioiosa, E. Kursun, C. Isci, A. Buyuktosunoglu, and M. Valero. A case for energy aware accounting in large scale computing facilities: Cost metrics and implications for processor design. In *Workshop on Architectural Concerns in Large Datacenters (ACLD), in conjunction with ISCA*, June 2010.

K. Kedzierski, F. Cazorla, R. Gioiosa, A. Buyuktosunoglu, and M. Valero. Power and performance aware reconfigurable cache for cmps. In *Workshop on Next Generation Multicore/Manycore Technologies (IFMT), in conjunction with ISCA*, June 2010.

M. Paolieri, I. Bonesana, R. Gioiosa, and M. Valero. J-dse: Joint software and hardware design space exploration for application specific processors. In *Third Workshop on Programmability Issues for Multi-Core Computers (MULTIPROG)*, January 2010.

C. Luque, M. Moreto, F.J. Cazorla, R. Gioiosa, A. Buyuktosunoglu, and M. Valero. Itca: Inter-thread conflict-aware cpu accounting for cmps. In *The International Symposium on Parallel Architectures and Compilation Techniques (PACT09)*, September 2009.

E. Betti, M. Cesati, R. Gioiosa, and F. Piermaria. Global Operating System for HPC Clusters. *The 2009 IEEE International Conference on Cluster Computing (Cluster 2009)*, August 2009.

C. Luque, M. Moreto, F.J. Cazorla, R. Gioiosa, A. Buyuktosunoglu, and M. Valero. Cpu accounting in cmp processors. *Computer Architecture Letters*, 2009.

C. Boneti, R. Gioiosa, F. J. Cazorla, and M. Valero. A dynamic scheduler for balancing HPC applications. In *SC '08: Proc. of the 2008 ACM/IEEE conference on Supercomputing*, November 2008.

P. Radojkovic, V. Cakarevic, F. Cazorla, R. Gioiosa, A. Pajuelo, and J. Verdu. Measuring Operating System Overhead on CMT Processors. In *The 20th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD08)*, Campo Grande, Mato Grosso do Soul, Brazil, October 2008.

C. Boneti, F. Cazorla, R. Gioiosa, C-Y. Cher, A. Buyuktosunoglu, and M. Valero. Software-Controlled Priority Characterization of POWER5 Processor. In *The 35th International Symposium on Computer Architecture (ISCA).*, Beijing, China, June 2008.

V. Cakarevic, P. Radojkovic, R. Gioiosa, F. Cazorla, J. Verdu, A. Pajuelo, and M. Valero. Understanding the overhead of the spin-lock loop in CMT architectures . In *"Workshop on the Interaction between Operating Systems and Computer Architecture" (WIOSCA), in conjunction with ISCA'08*, Beijing, China, June 2008.

C. Boneti, R. Gioiosa, F. Cazorla, J. Corbalan, J. Labarta, and Mateo Valero. Balancing HPC applications through smart allocation of resources in MT processors. In *to appear in the 22nd IEEE Int. Parallel and Distributed Processing Symp. (IPDPS08)*, Miami, FL, April 2008.

V. Cakarevic, P. Radojkovic, R. Gioiosa, F. Cazorla, J. Verdu, A. Pajuelo, and M. Valero. Overhead of the spin-lock loop in UltraSPARC T2. *the 5th HiPEAC Industrial Workshop*, June 2008.

C. Boneti, F. Cazorla, R. Gioiosa, and Mateo Valero. Scheduling real-time systems with explicit resource allocation processors. In *the 21st International Conference on Architecture of Computing Systems (ARCS08)*, Dresden, Germany, February 2008.

E. Betti, D.P. Bovet, M. Cesati, and R. Gioiosa. Hard real-time performances in multiprocessor embedded systems using ASMP-Linux. *Operating System Support for Embedded Real-Time Applications special issue on Eurasip Journal on Embedded Systems*, October 2007.

E. Betti, D.P. Bovet, M. Cesati, and R. Gioiosa. complete data set for the article "Hard real-time performances in multiprocessor embedded systems using ASMP-Linux". *Technical Report SPRGTR002*, September 2007.

R. Gioiosa, D.P. Bovet, and M. Cesati. vxWorks driver for Asine ASPMC660 flash card. *Technical Report SPRGTV-Q001*, July 2006.

M. Cesati, R. Gioiosa, and D.P. Bovet. On the Execution Time of a FFT procedure with and without caches. *Technical Report SPRGTV-Q002*, July 2006.

R. Gioiosa, D.P. Bovet, and M. Cesati. Real-time tests on vxWorks. *Technical Report SPRGTV-Q004*, July 2006.

R. Gioiosa. High Performance Computing Clusters. *Ph.D. Thesis*, May 2006.

R. Gioiosa, J.C. Sancho, S. Jiang, F. Petrini, and K. Devis. Transparent Incremental Checkpoint at Kernel level: A Foundation for Fault Tolerance for Parallel Computers. *SC|05 (Supercomputing): Int. Conf. for High Performance Computing, Networking, and Storage*, November 2005. http://bravo.ce.uniroma2.it/home/gioiosa/pub/sc05.pdf.

José Carlos Sancho, Fabrizio Petrini, Kei Davis, Roberto Gioiosa, and Song Jiang. Current Practice and a Direction Forward in Checkpoint/Restart Implementations for Fault Tolerance. In *First Workshop on System Management Tools for Large-Scale Parallel Systems*, Denver, CO, April 2005. Available from `http://www.c3.lanl.gov/~fabrizio/papers/survey-checkpoint.pdf`.

R. Gioiosa, F. Petrini, K. Davis, and F. Lebaillif-Delamare. Analysis of System Overhead on Parallel Computers. In *The 4th IEEE Int. Symp. on Signal Processing and Information Technology (ISSPIT 2004)*, Rome, Italy, December 2004. http://bravo.ce.uniroma2.it/home/gioiosa/pub/isspit04.pdf.