Workshop on Large-Scale Parallel Processing – LSPP

Workshop Theme

The workshop on Large-Scale Parallel Processing is a forum that focuses on computer systems that utilize thousands of processors and beyond. Large-scale systems, referred to by some as extreme-scale and Ultra-scale, have many important research aspects that need detailed examination in order for their effective design, deployment, and utilization to take place. These include handling the substantial increase in multi-core on a chip, the ensuing interconnection hierarchy, communication, and synchronization mechanisms. Increasingly this is becoming an issue of co-design involving performance, power and reliability aspects. The workshop aims to bring together researchers from different communities working on challenging problems in this area for a dynamic exchange of ideas. Work at early stages of development as well as work that has been demonstrated in practice is equally welcome. Of particular interest are papers that identify and analyze novel ideas rather than providing incremental advances in the following areas:

- **Large-scale systems**: exploiting parallelism at large-scale, the coordination of large numbers of processing elements, synchronization and communication at large-scale, programming models and productivity
- **Novel architectures and experimental systems**: the design of novel systems, the use emerging technologies such as Non-Volatile Memory, Silicon Photonics, application-specific accelerators and future trends.
- **Monitoring, Analysis, and Modeling**: tools and techniques for gathering performance, power, thermal, reliability, and other data from existing large scale systems, analyzing such data offline or in real time for system tuning, and modeling of similar factors in projected system installations.
- **Multi-core**: utilization of increased parallelism on a single chip, the possible integration of these into large-scale systems, and dealing with the resulting hierarchical connectivity.
- **Energy Management**: Techniques, strategies, and experiences relating to the energy management and optimization of large-scale systems.
- **Applications**: novel algorithmic and application methods, experiences in the design and use of applications that scale to large-scales, overcoming of limitations, performance analysis and insights gained.
- **Warehouse Computing**: dealing with the issues in advanced datacenters that are increasingly moving from co-locating many servers to having a large number of servers working cohesively, impact of both software and hardware designs and optimizations to achieve best cost-performance efficiency.

Workshop Chairs

Darren J. Kerbyson, Pacific Northwest National Laboratory
Ram Rajamony, IBM Austin Research Lab
Charles Weems, University of Massachusetts

Additional Steering Committee Members

Johnnie Baker, Kent State University
Alex Jones, University of Pittsburgh
H.J. Siegel, Colorado State University
Lixin Zhang, Institute of Computing Techology, Chinese Academy of Sciences
Guangming Tan, Institute of Computing Techology, Chinese Academy of Sciences

Provisional Program Committee

- Pavan Balaji, Argonne National Laboratory, USA
- Kevin J. Barker, Pacific Northwest National Laboratory, USA
- Laura Carrington, San Diego Supercomputer Center, USA
- I-Hsin Chung, IBM T.J. Watson Research Lab, USA
- Tim German, Los Alamos National Laboratory, USA
- Georg Hager, University of Erlangen, Germany
- Simon Hammond, Sandia National Laboratory, USA
- Martin Herbordt, Boston University, USA
- Daniel Katz, University of Chicago, USA
- Celso Mendes, University of Illinois Urbana-Champagne, USA
- Bernd Mohr, Forschungszentrum Juelich, Germany
- Phil Roth, Oak Ridge National Laboratory, USA
- Jose Sancho, Barcelona Supercomputer Center, Spain
- Gerhard Wellein, University of Erlangen, Germany
- Pat Worley, Oak Ridge National Laboratory, USA
- Ulrike Yang, Lawrence Livermore National Laboratory, USA

IEEE computer society