

5th International Workshop on Embedded/Distributed HPC Systems and Applications (EHPC 2000)

Workshop Co-Chairs

Devesh Bhatt

Honeywell Technology Center
3660 Technology Drive
Minneapolis, MN 55418, USA
devesh.bhatt@honeywell.com

Lonnie R. Welch

Ohio University
School of Engineering and Computer Science
Athens, OH 45701-2979, USA
welch@ohio.edu

Preface

The International Workshop on Embedded/Distributed HPC Systems and Applications (EHPC) is a forum for the presentation and discussion of approaches, research findings, and experiences in the applications of High Performance Computing (HPC) technology for embedded/distributed systems. Of interest are both the development of relevant technology (e.g.: hardware, middleware, tools) as well as the embedded HPC applications built using such technology.

We hope to bring together industry, academia, and government researchers/users to explore the special needs and issues in applying HPC technologies to defense and commercial applications.

Topics of Interest

- **Algorithms and Applications:** addressing parallel computing needs of embedded military and commercial applications areas such as signal/image processing, advanced vision/robotic systems, smart-sensor based systems, industrial automation/optimization, vehicle guidance.
- **Networking Multiple HPC Systems:** in-the-large application programming models/API's, partitioning/mapping, system integration, debugging and testing tools.
- **Programming Environments:** software design, programming, and parallelization methods/tools for DSP-based, reconfigurable, and mixed-computation-paradigm architectures.
- **Operating Systems and Middleware:** distributed middleware service needs (e.g. QoS, object distribution) of high-performance embedded applications, configurable/optimal OS features needs, static/dynamic resource management needs.

- **Architectures:** special-purpose processors, packaging, mixed-computation-paradigm architectures, size/weight/power modeling and management using hardware and software techniques.

EHPC 2000 Contents

The EHPC 2000 workshop will feature technical paper presentations, and an open discussion session. This year, we have papers covering several topic areas of interest. The following is a highlight of the papers.

In the algorithm and applications area, Yang et al. present a reconfigurable, dynamic load balancing parallel sorting algorithm applicable to information fusion. Hadden et al. present system health management application domain which would benefit from embedded HPC architectures.

In the programming environments area, Janka and Wills present a specification and design methodology for signal-processing systems using high-performance middleware and front-end tools. Patel et al. present performance comparison of high-performance real-time benchmarks using hand-crafted design versus automated glue-code generation from data-flow specification using their design tool.

In the operating systems and middleware area, we have several papers ranging from network load monitoring to communication scheduling for high-performance applications. Islam et al. present a technique for evaluating network load based upon dynamic paths using embedded application benchmarks. Pierce et al. present an architecture for mining of performance data for HPC systems, extending the capabilities of current instrumentation tools. Huh et al. present an approach for predicting the real-time QoS in dynamic heterogeneous resource management systems. VanVoorst and Seidel present the use of a real-time parallel communication benchmark to compare several MPI implementations. West and Antonio present an approach for optimizing the communication scheduling in parallel Space-Time Adaptive Processing (STAP) applications.

In the architecture area, we have papers on software and hardware perspectives on power management, as well as a new architecture for embedded applications. Osmulski et al. present a probabilistic power-prediction tool for Xilinx 4000-series reconfigurable computing devices. Unsal et al. present an energy consumption model addressing task assignment and network topology/routing, using replication of shared data structures. Schulman et al. present a system-on-chip architecture containing an array of VLIW processing elements, with reconfiguration times much smaller than FPGA-based architectures.

Program Committee

Ashok Agrawala, Univ. of Maryland, USA

Bonnie Bennett, Univ. of St. Thomas, USA

Bob Bernecky, NUWC, USA

Alberto Broggi, Universita' di Pavia, Italy

Hakon O. Bugge, Scali Computer, Norway

Richard Games, MITRE, USA

Farnam Jahanian, Univ. of Michigan, USA
Magnus Jonsson, Halmstad University, Sweden
Jeff Koller, USC/Information Sciences Institute, USA
Bruce Lewis, US Army AmCom, USA
Mark Linderman, USAF Rome Laboratory, USA
Craig Lund, Mercury Computer Systems, Inc., USA
Stephen Rhodes, Adadvanced Systems Architectures Ltd., UK
Samuel H. Russ, Mississippi State Univ., USA
Behrooz Shirazi, University of Texas at Arlington, USA
Anthony Skjellum, Mississippi State Univ., USA
Brian VanVoorst, Michigan Technological Univ., USA
Sudhakar Yalamanchili, Georgia Tech., USA

Advisory Committee

Keith Bromley, NRaD, USA
Dieter Hammer, Eindhoven Univ. of Technology, The Netherlands
David Martinez, MIT Lincoln Laboratory, USA
Viktor Prasanna, Univ. of Southern California, USA