Fifth International Workshop on High–level Parallel Programming Models and Supportive Environments HIPS 2000

Chair: Martin Schulz
schulzm@in.tum.de

Lehrstuhl für Rechnertechnik und Rechnerorganisation, LRR–TUM
Institut für Informatik, Technische Universität München
80290 München, Germany

May 1st, 2000
Cancun/Mexico

Preface

Following the long tradition of this well established event, the 5th International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS 2000) provides a forum for researchers and developers from both academia and industry to meet and discuss the newest approaches and results in this active research area. It is again held in conjunction with IPDPS (formerly known as IPPS/SPDP), one of the premier events in the area of parallel and distributed processing.

Despite this long tradition of the HIPS workshop series, the topic — how to efficiently exploit parallel and distributed architectures with respect to performance and ease-of-use — has not lost any of its importance. On the contrary, with the rise of cluster architectures built from commodity–off–the–shelf components and possibly even interconnected with high–performance networking technologies such as SCI, Myrinet, or GigaNet, high-performance parallel computing has become available to a larger base of users than ever before.

Nevertheless, programming is still mostly restricted to the message passing paradigm, which with the advent of the de facto standards Parallel Virtual Machine (PVM) and Message Passing Interface (MPI) has reached a certain level of maturity. However, in terms of convenience and productivity, this parallel programming style is often considered to correspond to assembler-level programming of sequential computers.

One of the keys for a (commercial) breakthrough of parallel processing are therefore easy-to-use high-level programming models that allow to produce truly efficient code. In this respect, languages and packages have been established that are more convenient than explicit message passing and allow higher productivity in software development; examples are High Performance Fortran (HPF), OpenMP, thread packages for shared memory-based programming, and Distributed Shared Memory (DSM) environments.

Yet, current implementations of high-level programming models often suffer from low performance of the generated code, from the lack of corresponding high-level development tools, e.g. for performance analysis, and from restricted applicability, e.g. to the data parallel programming style. This situation requires strong research efforts in the design of parallel programming models and languages that are both at a high conceptual level and implemented efficiently. In addition, it is necessary to design and implement supportive tools and integrated programming environments to assist in the development of parallel and distributed applications. Hardware and operating system support for high-level programming, e.g. distributed shared memory and monitoring interfaces, are further areas of interest contributing to this goal.

The HIPS workshop series provides a forum to discuss these issues and present new approaches providing users with high-level models and tools that allow the easy and efficient exploitation of parallel and distributed architectures. It addresses a wide range of topics form different areas including, but not limited to:

- Concepts and languages for high-level parallel programming.
- Concurrent object-oriented programming.
- Distributed objects and components.
- Structured parallel programming (skeletons, patterns, ...).
- Software engineering principles.
- Automatic parallelization and optimization.
- High-level programming environments.
- Performance analysis, debugging techniques and tools.
- Distributed shared memory.
- Implementation techniques for high-level programming models.
- Operating system support for runtime systems.
- Architectural and high-speed communication support for high-level programming models.

These topics attracted numerous submissions to this year's HIPS 2000. Each submitted paper was assigned to at least 3 program committee members and underwent a rigorous review process. The reviewers spent a lot of time and gave very detailed comments and recommendations, which lead to an initial list of accepted papers. The program committee members were then given the opportunity to resolve differences in opinion. At the end, after a short second reviewing process for some disputed papers, 10 high-quality papers (listed further below) have been selected for publication. These are presented, together with one invited talk, during the workshop on May 1st, 2000.

Workshop Chair

- Martin Schulz, Technische Universität München, Germany

Steering Committee

- Michael Gerndt, Forschungszentrum Jülich, Germany
- Hermann Hellwagner, Universität Klagenfurt, Austria
- Frank Müller, Humboldt Universität Berlin, Germany

Program Committee

- Arndt Bode, Technische Universität München, Germany
- Helmar Burkhart, Universität Basel, Switzerland
- John Carter, University of Utah, USA
- Karsten Decker, Swiss Center for Scientific Compting, Switzerland
- Michael Gerndt, Forschungszentrum Jülich, Germany
- Hermann Hellwagner, Universität Klagenfurt, Austria
- Francois Irigoin, Ecole des Mines de Paris, France
- Vijay Karamcheti, New York University, USA
- Peter Keleher, University of Maryland, USA
- Gabriele Keller, University of Technology, Sydney, Australia
- Piyush Mehrotra, ICASE / NASA Lanley Research Center, USA
- Frank Müller, Humboldt Universität Berlin, Germany
- Susanna Pelagatti, Universita di Pisa, Italy
- Thierry Priol, IRISA, France
- Martin Schulz, Technische Universität München, Germany
- Xian-He Sun, Louisianna State University, USA
- Domenico Talia, Universita della Calabria, Italy
- George Thiruvathukal, DePaul University, USA

Acknowledgments

I would like to thank the following people for their significant contribution to the success of this workshop: the steering committee members listed above for their assistance and support, Jose Rolim and Danuta Sosnowska from the IPDPS team for their organizational support, the many programmers of the START software package used for the electronic submission process, especially Rich Gerber, and last but not least the members of the program committee and the anonymous external reviewers for their excellent work during the review process.

May 2000

Martin Schulz
Chair of HIPS 2000

Accepted Papers for HIPS 2000

Session 1

- Pipelining Wavefront Computations: Experiences and Performance
  E Christopher Lewis and Lawrence Snyder, University of Washington
- Specification Techniques for Automatic Performance Analysis Tools
  Michael Gerndt and Hans-Georg Eßer, Research Centre Juelich
- PDRS: A Performance Data Representation System
  Xian-He Sun, Louisiana State University and Illinois Institute of Technology
  Xingfu Wu, Northwestern University and Louisiana State University

Session 2

- Clix - A Hybrid Programming Environment for Distributed Objects and Distributed Shared Memory
  Frank Müller, Humboldt University Berlin
  Jörg Nolte, GMD FIRST
  Alexander Schlaefer, University of Washington
- Controlling Distributed Shared Memory Consistency from High Level Programming Languages
  Yvon Jegou, IRISA/INRIA
- Online Computation of Critical Paths for Multithreaded Languages
  Yoshihiro Oyama, Kenjiro Taura, and Akinori Yonezawa, University of Tokyo

Session 3

- Problem Solving Environment Infrastructure for High Performance Computer Systems
  Daniel C. Stanzione, Jr. and Walter B. Ligon III, Clemson University

Session 4

- Combining Fusion Optimizations and Piecewise Execution of Nested Data-Parallel Programs
  Wolf Pfannenstiel, Technische Universität Berlin
- Declarative concurrency in Java
  Rafael Ramirez and Andrew E. Santosa, National University of Singapore
- Scalable Monitoring Technique for Detecting Races in Parallel Programs
  Yong-Kee Jun, Gyeongsang National University
  Charles E. McDowell, University of California at Santa Cruz