


I. Banicescu (Mississippi State Univ., USA), K. Pingali (Cornell Univ. USA), T. Rauber (Univ. Bayreuth, D), G. Rünger (TU Chemnitz, D)
(Editors)

Adaptivity in Parallel Scientific Computing

Dagstuhl Seminar 03211 – May 18 to May 23, 2003
Dagstuhl-Seminar-Report No. 380



SCHLOSS DAGSTUHL

INTERNATIONALES
BEGEGNUNGS-
UND FORSCHUNGSZENTRUM
FÜR INFORMATIK

ISSN 0940-1121

Herausgegeben von IBFI gem. GmbH, Schloss Dagstuhl, 66687 Wadern, Germany.

Das Internationale Begegnungs- und Forschungszentrum für Informatik (IBFI) Schloss Dagstuhl ist eine gemeinnützige GmbH. Sie veranstaltet regelmäßig wissenschaftliche Seminare, welche nach Antrag der Tagungsleiter und Begutachtung durch das wissenschaftliche Direktorium mit persönlich eingeladenen Gästen durchgeführt werden.

Gesellschafter:

- Gesellschaft für Informatik e.V. – Bonn
- TH Darmstadt
- Universität Frankfurt
- Universität Kaiserslautern
- Universität Karlsruhe
- Universität Stuttgart
- Universität Trier
- Universität des Saarlandes

Summary

Although progress in parallel and distributed methodologies for scientific computing have been quite remarkable during the past years, this area of computer science remains still active, especially in topics concerning the relationship between performance and aspects such as: irregularity of applications and algorithms, adaptive characteristics of software and hardware, heterogeneity of hardware platforms, and flexibility of programming environments. Recent research activities include development of complex hardware architectures, including storage hierarchies or heterogeneous (parallel and distributed) computing platforms with large numbers of processors, as well as irregular applications that involve complex domain decomposition and hierarchical, adaptive and multi-level organization of computation and data structures.

The corresponding irregular algorithms comprise applications with sparse, block-structured or adaptive data structures, as well as applications with irregular, runtime-dependent computation and control structures.

Over time, to improve scientific applications' performance on sequential machines, several techniques in hardware and algorithm design, such as storage hierarchies and hierarchical domain decomposition, have been introduced. However, the simulation of large irregular problems still requires the use of parallel and distributed environments. The irregular and dynamically changing runtime behavior makes an efficient parallel realization difficult, since the memory access patterns and the evolution of dynamic structures cannot be determined a priori, and therefore, cannot be planned statically. Consequently, an efficient parallel implementation of this class of problems necessitates the exploitation of flexible programming environments as well as techniques to improve scalability.

This seminar was a forum that brought together researchers working in different areas of parallel scientific computing and its applications, to solve scientific and industrially oriented problems. It provided a fertile environment for the participants to meet and exchange ideas, as well as to foster future research collaborations.

Of particular interest was the exchange of experiences in interdisciplinary research projects. Topics covered by this seminar included:

- parallel numerical algorithms
- parallel implementation of irregular applications
- algorithms for memory hierarchies with enhanced locality of memory access
- libraries for supporting parallel scientific computing
- mixed task and data parallel executions on large parallel machines
- performance analysis evaluation and prediction
- compiler transformations for increasing the locality of memory references
- dynamic load balancing techniques

- partitioning and scheduling strategies
- heterogeneous computing (cluster and grid computing)
- combination of different programming models for heterogeneous parallel machines

During the seminar, a number of presentations lead to formulation of interesting open questions followed by discussions on optimal integration of adaptivity at various levels of technology in application, algorithms and system development. In the following paragraphs, we summarize a few concepts and ideas for new approaches, methodologies, and future directions that spawned from various talks and discussions.

In recent years, research in modeling and simulation is becoming increasingly important for a wide variety of scientific and engineering disciplines. It addresses the need for developing a safe, dependable and effective information environment, as well as the one for expanding of basic research in revolutionary fields which are of vital importance to our society. As a result, the research community is now faced with new challenges, such as the ones to incorporate additional physics, length scales and time scales, into models for adaptivity, higher fidelity and resolution, or to process variable amount of data from distributed datasets, which in turn place significant demands on software design and hardware implementation. Therefore, there is a need to explore and devise the design of a flexible, robust, and effective *vertical integration strategy*, for advanced development of scientific applications. This integration is expected to facilitate an effective fusion of advances in application algorithms, with the ones in programming environments, system software and hardware capabilities, for the purpose to enable terascale modeling and simulation.

The "Tinker-toy Parallel Programming", an interesting approach to building scientific applications via an aggregation of multiple, light-weight toolkits has been introduced during the seminar presentations. A solution to one possible drawback of such an approach, its limited support for adaptive computations, has also been proposed using "Zoltan" – a tool that provides support for adaptive, parallel scientific computations, and easy development for dynamic and adaptive simulations.

The rapid development of an emerging technology in "cluster and grid computing" suggests a need for dynamic distribution of work and data that can be adapted to the runtime behaviour of the algorithm. A solution to that has been proposed, and its design, implementation, and evaluation have been presented using "task pools". In this approach, tasks are dynamically distributed to different processors (within nodes of a SMP, or among nodes on clusters of SMPs), and each task specifies computations to be performed and provides the appropriate data.

Some interesting presentations focussed on improving performance of irregular parallel applications via addressing sources of load imbalance at all levels of irregular behaviour (related to problem, algorithm or systemic factors). A general purpose tool for dynamic loop scheduling to address the stochastic load variations from a range of sources has also been introduced.

A number of interesting discussions took place regarding recent advances in cluster and grid computing through a successful migration of parallel programs (via checkpointing

and fault tolerance). In the future, the migration of parallel programs will allow parallel applications to "surf" the grid and adapt dynamically to its changeable environment.

A few interesting contributions presented challenges in BSP algorithm design, programming and software engineering to address adaptivity in scientific computations. Moreover, there were a few novel ideas and original concepts introduced on language support for irregular problems and adaptivity. The audience was delighted to discuss during the talk, as well as during our evening pleasant moments of get-together, some of the possible breakthroughs that could evolve from these ideas.

In conclusion, this seminar presentations and discussions addressed many complex issues including application requirements for adaptivity in space and time, as well as requirements for improving the capacity to effectively use resources in heterogeneous environments. The seminar topics span and integrate the work of many research areas: from irregular scientific applications, to adaptive algorithms, programming models and tools, problem solving environments for cluster and grid computing, and others.

We believe that these contributions, in addition to talks and many interesting discussions, will inspire the participants to continue their research efforts towards an integrated view of adaptivity, allowing them in this way to make significant contributions to the advancement of science.

Participants

- Balasubramaniam, Mahadevan (Mississippi State University)
- Banicescu, Ioana (Mississippi State University)
- Beckmann, Olav (Imperial College London)
- Carino, Ricolindo (Mississippi State University)
- Chronopoulos, Anthony (Univ. of Texas at San Antonio)
- Danelutto, Marco (University of Pisa)
- Gutierrez, Eladio D. (University of Malaga)
- Hendrickson, Bruce (Sandia National Labs – Albuquerque)
- Hippold, Judith (TU Chemnitz)
- Kessler, Christoph W. (Linköping University)
- Korch, Matthias (Universität Bayreuth)
- Kühnemann, Matthias (TU Chemnitz)
- Lucas, Philipp (Universität des Saarlandes)
- O’Donnell, John (University of Glasgow)
- Oksa, Gabriel (Slovak Academy of Sciences – Bratislava)
- Petcu, Dana (West University of Timisoara)
- Rauber, Thomas (Universität Bayreuth)
- Rünger, Gudula (TU Chemnitz)
- Scholtes, Carsten (Universität Bayreuth)
- Sips, Henk (TU Delft)
- Tiskin, Alexander (University of Warwick)
- Torczon, Virginia (College of William and Mary)
- van Albada, Dick (University of Amsterdam)
- Weberpals, Helmut (TU Hamburg-Harburg)
- Zimmermann, Wolf (Martin-Luther-Universität Halle-Wittenberg)