

04451 Abstracts Collection
Future Generation Grids
— Dagstuhl Seminar —

Michel Cosnard¹, Vladimir Getov², Domenico Laforenza³
and Alexander Reinefeld⁴

¹ INRIA, Sophia Antipolis Research Unit
06902, Sophia Antipolis, France
`Michel.Cosnard@inria.fr`

² University of Westminster, School of Computer Science
HA1 3TP, London-Harrow, Great Britain
`v.s.getov@wmin.ac.uk`

³ Consiglio Nazionale delle Ricerche, Istituto di Scienza e Tecnologia per
l'Informazione (ISTI)
56124, Pisa, Italy
`domenico.laforenza@cnuce.cnr.it`

⁴ Konrad-Zuse-Zentrum für Informationstechnik and Humboldt-Universität zu Berlin
14195, Berlin, Germany
`ar@zib.de`

The Dagstuhl Seminar 04451 “Future Generation Grid” was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl from 1st to 5th November 2004. The focus of the seminar was on open problems and future challenges in the design of next generation Grid systems. A total of 45 participants presented their current projects, research plans, and new ideas in the area of Grid technologies. Several evening sessions with vivid discussions on future trends complemented the talks. This report gives an overview of the background and the findings of the seminar.

Keywords. Grid Architectures, Autonomous Computing, Knowledge Grid, Peer-to-Peer, WSRF, OGSF, SOA

1 Background

The Internet and the Web have had a major impact on society. By allowing us to discover and access information on a global scale, they have created entirely new businesses and brought new meaning to the term “surf”. Yet, simply being able to offer and access information on the Web is ultimately unsatisfactory: We want processing and, increasingly, we want collaborative processing within distributed teams. This need has led to the creation of the Grid - an infrastructure that enables us to share capabilities, integrate services and resources within and across enterprises, and allows active collaborations across distributed, multi-organizational environments. Powered by on-demand access to computer resources, seamless access to data, and dynamic composition of

distributed services, the Grid promises to enable fundamentally new ways of interacting with our information technology infrastructure, of doing business, and practicing science. It represents perhaps the final step in the great disappearing act that will take computing out of our homes and machine rooms into the fabric of the society, where it will stand alongside telephone switches, power generators, and the other invisible technologies that drive the modern world. Future applications will not only use individual computer systems, but a large set of networked resources. This scenario of computational and data grids is attracting a lot of attention from application scientists as well as from computer scientists. In addition to the inherent complexity of current high-end systems, the sharing of resources and the transparency of the actual available resources introduce not only new research challenges but also a completely new vision and approaches to designing, building, and using future generation Grid systems.

2 Overview of the Seminar

The seminar brought together 45 scientists and researchers in the Grid area in an attempt to draw a clearer picture of future generation Grids and to identify the most challenging problems on the way to achieving the invisible information Grid ideas in our society. The participants came from France (12), Germany (10), Italy (8), Great Britain (5), The Netherlands (3), Belgium (1), Cyprus (1), Czech Republic (1), Poland (1), Spain (1), Switzerland (1), and the U.S.A. (1). This was the first seminar in a series of workshops planned by the just-started EU Network of Excellence project CoreGRID – the “European Research Network on Foundations, Software Infrastructures and Applications for large scale distributed, GRID and Peer-to-Peer Technologies”. The CoreGRID Network of Excellence aims at strengthening and advancing scientific and technological excellence in the area of Grid and Peer-to-Peer technologies. To achieve this objective, the Network brings together a critical mass of well-established researchers (119 permanent researchers and 165 PhD students) from forty-two institutions, which have constructed an ambitious joint programme of activities. Additional impetus for the organization of the FGG-Seminar came from another EU project, the “ERA Pilot on a Coordinated Europe-Wide Initiative in Grid Research” (GridCoord). It targets at strengthening Europe’s position on grid research and its exploitation by overcoming the fragmentation and dispersion across the EU research programs. The FGG seminar helped in getting an overview on the various Grid initiatives and projects and thereby provided a good basis for drafting a compendium on Grid research programmes in major European states. Background information for the seminar came from the just recently published findings of a EU expert group on Next Generation Grids (<http://www.cordis.lu/ist/grids/index.htm>).

3 Results

The talks of the seminar were grouped into the broad topics middleware, software toolkits, Grid and peer-to-peer system architecture, data and information management, resource management, and scheduling. In an attempt to provide an overview on the status of the various national Grid Initiatives – a topic deemed important especially for the GridCoord project – the following Grid Initiatives were presented:

- Grid.it (Italy)
- D-Grid (Germany)
- DAS-2 (The Netherlands)
- SGIGrid (Poland)
- UK e-Science (UK)
- ACI GRID’s Grid’5000 project (France)

While the general goal of establishing a national Grid for the benefit of science and research in the respective countries is similar, each of these initiatives puts an emphasis of slightly different aspects. Most apparent is perhaps the “virtual laboratories” approach in the Netherlands, the more experimental character of the French Grid 5000 project as part of the ACI GRID initiative and the strong trend towards the deployment of productive application scenarios in the UK e-Science initiative. However, it is difficult to summarize the subtle differences in the initiatives in this brief summary and therefore, a more detailed analysis must be left for the future. The discussion session on next generation Grid technologies focussed largely on the importance of making Grid systems “autonomic” in the sense that future Grid components shall be able to autonomously cope with failures without affecting the other “healthy” components. Even more emphasis was put on the discussion of the newly established Web Services Resources Framework (WSRF) versus the previous Open Grid Service Infrastructure (OGSI), Web Services, and Service Oriented Architectures (SOA) in general.

4 Program

To allow for spontaneous changes, the schedule was not completely fixed before the start of the seminar. Rather, the participants were asked to submit titles of their presentations. These talks have been put into a tentative order by the organizers just before the seminar. Sometimes, the agenda departed from the planned order to allow for spontaneous discussions on newly emerged topics of general interest. The following list gives an overview of the actual program.

Tuesday, 02.11.2004

9:00 Introduction Organizational Matters (Alexander Reinefeld)

9:30 Middleware (Alexander Reinefeld)

- *MPICH-V: a multi-protocols fault tolerant MPI*
Franck Cappello, INRIA Futurs - Orsay
- *A middleware toolbox for clusters and grids*
Thomas Herault, Université Paris SudGuillaume Huard, Laboratoire ID-IMAG
- *Improving Grid Middleware with Adaptive On-line Compression*
Emmanuel Jeannot, LORIA, Univ. Henri Poincaré - Nancy 1

14:00 Software Toolkits (Vladimir Getov)

- *Recent advances in DIET, a Toolbox for the development of Network Enabled Servers Environments*
Frédéric Desprez, ENS - Lyon
- *GAT - API for scientific Applications on the Grid*
André Merzky, Vrije Univ. Amsterdam
- *Ibis: a Java-based Grid Programming Environment*
Thilo Kielmann, Vrije Univ. Amsterdam

16:30 Software Toolkits contd. (Vladimir Getov)

- *GRIDLE: A Way to Compose Applications out of Software Components on the Grid*
Diego Puppini, CNR - Pisa
- *Defining, Implementing, Deploying and Using a Parallel Component Model*
Christian Perez, INRIA - Rennes
- *ParoC++: An Object-oriented environment for adaptive HPC programming on the Grid*
Pierre Kuonen, Univ. of Applied Sciences - Fribourg

20:00 National Grid Initiatives

- Grid.IT (Domenico Laforenza, Marco Danelutto)
- DGrid (Alexander Reinefeld)
- DAS-2 (Thilo Kielmann)
- SGIGrid (Marian Bubak)
- UK e-Science (Jon MacLaren)
- Grid 5000 (Franck Cappello)

Wednesday, 03.11.2004**9:00 Architecture, P2P and Autonomic (Domenico Laforenza)**

- *Distributed Components for the GRID*
Denis Caromel, INRIA - Sophia Antipolis
- *Automatic Reconfigurable Component Model For Grid Systems*
Jeyarajan Thiyaalingam, Imperial College - London
- *HOCs: Higher-Order Components for Grids*
Sergei Gorlatch, Universität Münster
- *Managing Grid under the Hood*
Marco Danelutto, Università di Pisa
- *Benchmarking Grid Resources*
George Tsouloupas, University of Cyprus

14:00 Afternoon Event**16:00 Trends in Grid Technology**

- *NGG++: Towards Next Generation Grids: Research needed and Open Issues*
Domenico Laforenza, CNR - Pisa
- *Autonomic Computing - The Prerequisite for Grid Computing*
Alexander Reinefeld, Zuse-Institut Berlin

19:30

- *FGG: Scalability, Suitability, Reconfigurability, Interoperability*
Vladimir Getov, University of Westminster
- discussion -

Thursday, 04.11.2004**9:00 Architecture, P2P and Autonomic contd. (Uwe Schwiegelshohn)**

- *How to make grid heterogeneity harmless: a proposal*
Marco Aldinucci, CNR - Pisa
- *Remote administration and fault tolerance in large distributed computer infrastructures*
Volker Lindenstruth, Universität Heidelberg

6 M. Cosnard, V. Getov, D. Laforenza, A. Reinefeld

- *Building a grid problem solving environment: the CrossGrid experience and beyond*

Marian Bubak - ICS/CYFRONET AGH Cracow

- *Event-driven workflows*

Craig Lee, Aerospace Corp. - El Segundo

Data and Information Management

- *GDS: Building a Grid Data Service*

Gabriel Antoniu, IRISA - Rennes

13:30 Data and Information Management contd. (Marco Danelutto)

- *The ZIB Distributed Data Management System*

Florian Schintke, Zuse-Institut Berlin

- *Designing Services for Grid-based Knowledge Discovery*

Domenico Talia, Università della Calabria

Data and Information Management

- *Processor and Data Co-allocation in Grids*

Dick H. J. Epema, Delft Univ. of Technology

- *The Virtual Resource Manager: An Architecture for SLA-aware Resource Management*

Hans-Ulrich Heiss, TU Berlin

- *Automatic Application Execution on Grids*

Nicola Tonellotto, CNR - Pisa

- *Coordinated Grid Scheduling*

Ramin Yahyapour & Uwe Schwiegelshohn, U. Dortmund

17:00 - 18:00 Discussion on Next Generation Grids

20:00 Discussion contd.

Friday, 05.11.2004

9:00 Resource Management and Scheduling contd. (Jon McLaren)

- *Bringing Knowledge to Middleware: Grid Scheduling Ontology*

Wolfgang Ziegler, Fraunhofer Institut - St. Augustin

- *A general scheme for resource matching in the context of global computing*
Pierre Manneback, Faculté Polytechnique de Mons
- *Non - Centralized Grid Infrastructure Monitoring*
Ludek Matyska, Masaryk University
- *Resource Demand Prediction for Utility Computing*
Artur Andrzejak, Zuse-Institut Berlin

5 Abstracts of the Talks

Rendering Grid heterogeneity harmless

Aldinucci, Marco (CNR - Pisa)

<http://www.di.unipi.it/~aldinuc>

Grid computing is supposed to enable the development of large scientific applications on an unprecedented scale by using the aggregate power of distributed resources, thus benefiting from a computing power that falls far beyond the current availability threshold in a single site. However, the number and quality of problems to be resolved in order to draw high-performance from Grid-aware applications is quite large. Because of this, a widely accepted programming model for the Grid is still lacking. Nevertheless, the Grid research community fully agree on some basic assumption about the Grid, as for example the very uneven nature of computing power of a Grid in terms of its homogeneity and availability over time. Paradoxically enough, none of the classical performance measures inherited from parallel computing world, as for example scalability, efficiency, isoefficiency, etc., seem suitable for describing applications quality on such a scenario. We therefore, can hardly argue about the quality of the implementation of them or even about the expected/ideal completion time bound of them.

We present a novel way to model the aggregate computing power of Grid execution environments which provides the application designer with a very intuitive and fairly precise measure of the performance quality of a given run w.r.t. its optimal scheduling/mapping on the same running environment. Also, the model enables the very fast forecast of future performance scenarios w.r.t. dynamically changing state of the execution environment. We eventually presents a first implementation of the ASSIST Application Manager able to keep an ASSIST application within certain performance contract constraints by using a configurable strategy based on the presented performance model.

Keywords: *Grid, Application Manager, performance model, storage component*
Joined work of: *Aldinucci, Marco; Campa, Sonia; Danelutto, Marco; Zoccolo, Corrado*

Modeling and Predicting Demand of Applications in Clusters

Andrzejak, Artur (ZIB)

<http://www.zib.de/andrzejak/>

We present approaches for characterizing and predicting demand behavior of applications, especially in clusters. The motivation of the demand modeling in context of self-management and resource sharing is discussed. After surveying approaches based on data mining classifiers and on the ARIMA/Kalman filter methods, we focus on method based on mining periodicities in historical data. The corresponding algorithms are change-adaptive in the sense that they automatically adjust to new regularities in demand patterns. They have been designed as tools for automating management of systems, and thus feature low running time and compact output. The approach is particularly useful for applications in scientific computing clusters, enterprise data centers and Grid and Utility environments, as most of them exhibit periodical behavior and offer high potential for automation. A case study using data from an enterprise data center evaluates the effectiveness of the technique.

Keywords: *prediction, modeling, pattern mining, resource management, self-management*

Building a grid data-sharing service or How to handle data consistency and fault-tolerance in a grid environment?

Antoniou, Gabriel (IRISA - Rennes)

<http://www.irisa.fr/paris/pages-perso/Gabriel-Antoniou>

We address the challenge of sharing large amounts of numerical data within computing grids consisting of clusters federation. We focus on the problem of handling the consistency of replicated data in an environment where the availability of storage resources dynamically changes. We propose a software architecture which decouples consistency management from fault-tolerance management. We illustrate this architecture with a case study showing how to design a consistency protocol using fault-tolerant building blocks. As a proof of concept, we describe a prototype implementation of this protocol within JuxMem, a software experimental platform for grid data sharing, and we report on a preliminary experimental evaluation.

Keywords: *grid, P2P, data-sharing, data consistency, fault-tolerance*

Building a Grid Problem Solving Environment: The CrossGrid Experience and Beyond

Bubak, Marian (AGH Univ. of Science & Technology - Krakow)

<http://www.icsr.agh.edu.pl>

This talk presents an approach to development of a problem solving environment for grid applications.

CrossGrid Project aims to develop new Grid services and tool environment for representative set of interactive compute- and data-intensive applications. The new tools and grid services are based on the Globus Toolkit and the EU DataGrid middleware. Important component of CrossGrid is an application monitoring infrastructure aimed at enabling performance measurements for the application developer. The application monitoring infrastructure is composed of a distributed monitoring system, the OCM-G, and a performance analysis tool called G-PM. The OCM-G is an on-line, grid-enabled, monitoring system, while G-PM is an advanced graphical tool which allows to evaluate and present the results of performance monitoring.

A Grid HLA Management System supports execution of HLA based interactive applications in a Grid environment. This subsystem consists of following components: the HLA Speaking Service for multiple federates that interface the HLA application to the system, Monitoring Service integrated with the OCM-G monitoring system and HLA based Benchmark Services informing the Broker Service of what can be expected from application behavior.

We have also elaborated a system to support the user in composing an application workflow from existing Grid services. The flow composition system builds workflows on an abstract level with semantic and syntactic descriptions of services available on the Grid and it consists of the flow composer and the distributed Grid service registry.

As the next step, we have proposed a universal architecture for porting legacy code to web services environment. This solution provides support for process migration, checkpointing and transactional processing. Both concurrent and asynchronous method invocation patterns are supported. A framework was developed to facilitate usage of the proposed architecture.

Finally, we outline the directions of ongoing and future research aiming on building a grid problem solving environment.

MPICH-V: a multi-protocols fault tolerant MPI

Cappello, Franck (INRIA Futurs - Orsay)

<http://www.lri.fr/~fci/>

High performance computing platforms like Clusters, Grid and Desktop Grids are becoming larger and subject to more frequent failures. MPI is one of the most used message passing library in HPC applications. These two trends raise the need for fault tolerant MPI. The MPICH-V project focuses on designing, implementing and comparing several automatic fault tolerance protocols for MPI

applications. We present the generic framework for fault tolerant protocol comparison. We detail four fault tolerant protocols implemented in MPICH-V, covering a large spectrum of known approaches from coordinated checkpoint, to uncoordinated checkpoint associated with causal message logging. We measure the performance of these protocols on a micro-benchmark and compare them for the NAS benchmark, using an original fault tolerance test. Finally, we outline the lessons learned from this in depth fault tolerant protocol comparison for MPI applications.

Distributed Components for the GRID

Caromel, Denis (INRIA - Sophia Antipolis)

<http://www.inria.fr/oasis/denis>

A key idea is the capability to define a distributed component: a piece of Java program that can be mapped (deploy) onto several JVMs, running on several processors or machines.

We based the component model on a precise distributed and parallel object model, available within the Open Source ProActive library <http://ProActive.ObjectWeb.org>.

The ProActive library puts at work the reuse potential of object-oriented languages using techniques such as polymorphism and dynamic method resolution for distribution. Crucial design choices are to completely avoid syntax and language extensions, and to not change the Java VM, but rather to develop tools that are 100% pure Java.

ProActive features the following: o Asynchronous calls and messages o Migration, Active objects, Mobile Agents o Group Communications with dynamic group management o Transparent, dynamic code loading (up and down) o XML descriptors for deployment o Interface with rsh, ssh, lsf, pbs, Sun grid engine, gsissh, Globus, oar, prun.

The library features a small set of robust migration primitives that allows mobile computations to seamlessly move from one Java virtual machine to another, while still communicating. Together with the library, the graphical environment IC2D (Interactive Control and Debugging of Distribution) allows to monitor and steer distributed and metacomputing applications. IC2D features graphical visualisation and drag-and-drop migration of remote objects.

Based on this underlying model and infrastructure, we propose a parallel and distributed component framework for building Grid applications, adapted to the hierarchical, highly distributed, highly heterogeneous nature of Grids. We have extended ProActive by implementing a hierarchical and dynamic component model, named Fractal, so as to master the complexity of composition, deployment, re-usability, and efficiency of grid applications. This defines a concept of Grid components, that can be parallel, made of several activities, and distributed. These components communicate using typed one-to-one or collective invocations.

A major achievement is the capacity to handle hierarchical components, in order to master complexity, composability, and efficiency.

Managing grid under the hood

Danelutto, Marco (Università di Pisa)

<http://www.di.unipi.it/~marcod/>

A key question when considering grids as target architectures in the development of complex, possibly multidisciplinary applications, is who's responsible of dealing with all the grid related features and problems. That is, who's in charge of handling/managing dynamicity, heterogeneity, variable network delays, resource discovery and so on? Is the application programmer the right person to charge of all the duties related to the handling of these issues, or should we move the task to somebody else, possibly someone more expert of the grid mechanisms and middleware tools?

We claim that most of the specific tasks needed when dealing with grids can be taken away from application programmers, provided that a structured programming environment is available. The structured programming environment must provide all the tools needed to express the grid application at the highest abstraction level possible. Then, the compiler and the run time support tools must take care of solving all the grid related problems exploiting the qualitative knowledge provided by the programmer.

We discuss a set of experiences developed at our department that show how this approach can be effectively adopted. In particular, we show how dynamic grid behavior and heterogeneity can be taken into account and how the related problems can be automatically solved in the compiler/run time support of the structured programming environment. Also, we will discuss how the grid targeting activity can be suitably embedded within a component framework through the introduction of a hierarchical application manager concept.

The results achieved while performing these preliminary experiences show that the concept of "invisible grid", that is a grid that is used but not actually perceived by the programmers developing grid applications, is not just a myth.

Keywords: *structured programming models, skeletons, adaptivity, heterogeneity, grid architecture*

Joined work of: *Danelutto, Marco; Aldinucci, Marco*

Recent advances in DIET, a Toolbox for the development of Network Enabled Servers Environments

Desprez, Frédéric (ENS - Lyon)

<http://graal.ens-lyon.fr/~desprez/>

DIET (Distributed Interactive Engineering Toolbox) is a set of hierarchical components to design Network Enabled Server systems. These systems are built upon servers managed through distributed scheduling agents for a better scalability. Clients ask to these scheduling components to find servers available (using some performance metrics and information about the location of data already on the network). Our target architecture is the grid which is highly heterogeneous and dynamic. In this presentation, we will describe its architecture and

some recent developments around scheduling, deployment, and monitoring of the platform.

Processor and Data Co-allocation in Grids

Epema, Dick H.J. (Delft University of Technology)

<http://www.pds.ewi.tudelft.nl/~epema/>

Jobs submitted to a grid may require more resources than are available at any point in time in any single subsystem making up a grid. Therefore, grid schedulers may employ co-allocation, that is, the simultaneous allocation of possibly multiple resources in multiple subsystems to a single job. We will report on the design and implementation of a grid scheduler that incorporates co-allocation of both processors and data on our DAS multicluster and on more heterogeneous testbeds, and on extensive simulations of processor co-allocation in multicluster systems.

Keywords: *co-allocation, scheduling, design, implementation, simulation*

Towards Building a Generic Grid Services Platform: A Component-Oriented Approach

Getov, Vladimir (Univ. of Westminster - London)

Grid applications using modern Grid infrastructures benefit from a rich variety of features, because they are designed with built-in exhaustive set of functions. As a result, the notion of a lightweight platform has not been addressed properly yet, and current systems cannot be transplanted, adopted or adapted easily. With the promise of the Grid to be pervasive, it is time to re-think the design methodology for next generation Grid infrastructures. Instead of building the underlying platform with an exhaustive rich set of features, in this chapter, we describe an alternative strategy following a component-oriented approach. Having a lightweight reconfigurable and expandable core platform is the key to our design. We identify and describe the very minimal and essential features that a modern Grid system should always offer and then provide any other functions as pluggable components. These pluggable components can be brought on-line whenever necessary as demanded implicitly by the application. With the support of adaptiveness, we see our approach as a solution towards a flexible dynamically reconfigurable Grid platform.

Joined work of: *Getov, Vladimir; Thiyagalingam, Jeyarajan; Isaiadis, Stavros*

HOCs: Higher-Order Components for Grids

Gorlatch, Sergei (Universität Münster)

<http://pvp.cs.tu-berlin.de/gorlatch/>

Our talk deals with the problem of application programming for future generation grids. We argue that grid programming is still too complex because of the big gap between the currently used grid middleware, e.g., the Globus Toolkit, and the application level. We propose to bridge this gap by using Higher-Order Components (HOCs) – recurring patterns of parallel behaviour that are provided to the user as program building blocks with pre-packaged implementation.

The presentation is illustrated with a case study.

Our experiments demonstrate that HOCs simplify grid application programming significantly, without serious performance loss.

RPC-V : Toward Fault Tolerant RPC for Grid with Volatile Nodes

Herault, Thomas (Université Paris Sud)

<http://www.lri.fr/~herault/>

RPC is one of the programming models envisioned for the Grid. In Internet connected Large Scale Grids such as Desktop Grids, nodes and networks failures are not rare events. This paper provides several contributions, examining the feasibility and limits of fault-tolerant RPC on these platforms.

First, we characterize these Grids from their fundamental features and demonstrate that their applications scope should be safely restricted to stateless services.

Second, we present a new fault-tolerant RPC protocol associating an original combination of three-tier architecture, passive replication and message logging. We describe RPC-V, an implementation of the proposed protocol within the XtremWeb Desktop Grid middleware.

Third, we evaluate the performance of RPC-V and the impact of faults on the execution time, using a real life application on a Desktop Grid testbed assembling nodes in France and USA. We demonstrate that RPC-V allows the applications to continue their execution while key system components fail.

Joined work of: *Djilali, Samir; Herault, Thomas; Lodygenski, Oleg; Morlier, Tanguy; Fedak, Gilles; Cappello, Franck*

A middleware toolbox for clusters and grids

Huard, Guillaume (Laboratoire ID-IMAG)

This talk presents some of the software and algorithms we develop at ID laboratory for the efficient management of large scale parallel architectures.

This includes : – Taktuk 2: a library that perform efficient and adaptative parallel distant execution on large scale platform.

Taktuk uses a combination of local parallelization and global distribution to deploy near-optimally a parallel distant execution.

Taktuk is able to use a static as well as a fully dynamic deployment scheme which makes it suited for both homogeneous and heterogeneous architecture. – Ka-tools 2: a set of tools for the management of distant environment (operating system + software suite) installation.

Using Ka-tools 2, one can make its own environment image, insert it to the system and reconfigure on-the-fly any set of nodes using the newly created image. In particular Ka-tools 2 can install any image on any node and handle console redirection as well as hard/soft reboot of nodes. – Oar and CiGri: two batch scheduler respectively designed for clusters and grids. Oar is a batch scheduler for cluster built upon high level tools such as MySQL and Perl. It emphasises simple and concise code rather than complex functionalities. CiGri is built upon Oar and is designed with Grids in mind.

Thus, it can handle large scale or parametrized jobs submissions and has built-in mechanisms for automatic error handling.

Joined work of: *Huard, Guillaume; Richard, Olivier*

Improving Grid Middleware with Adaptive On-line Compression

Jeannot, Emmanuel (LORIA, Univ. Henri Poincaré - Nancy 1)

<http://www.loria.fr/~ejeannot>

In this talk, we present the AdOC (Adaptive Online Compression) library. It is a user-level set of functions that enables data transmission with compression. The compression is performed dynamically during the transmission and the compression level is constantly adapted according to the environment. In order to ease the integration of AdOC into existing software the API is very close to the read and write UNIX system calls and respect their semantic. Moreover this library is thread-safe and is ported to many UNIX-like systems. We have tested AdOC under various conditions and with various data types. Results show that the library outperforms the POSIX read/write system calls on a broad range of network (up to 100 Mbit LAN), whereas on Gbit Ethernet, it provides similar performances. We have integrated this library into several grid Middleware, experiments on netSolve show an interesting increase of performance on various platforms.

Keywords: *lossless compression; communication library; adaptive service; grid middleware*

Ibis: a Java-based Grid Programming Environment*Kielmann, Thilo (Vrije Universiteit Amsterdam)*<http://www.cs.vu.nl/~kielmann/>

Ibis provides a Java-based programming environment for Grid applications.

Its runtime system deals with connectivity issues like firewalls, TCP performance and encryption. Ibis provides for APIs, namely RMI, RepMI (replicated objects), GMI (group method invocation), and the divide-and-conquer based Satin. Satin's load balancing algorithm, CRS, executes applications almost as efficient in Grid environments as on single clusters. For long-running codes, we have extended Satin with an efficient fault-tolerance mechanism.

ParoC++: An Object-oriented environment for adaptive HPC programming on the Grid*Kuonen, Pierre (Univ. of Applied Sciences - Fribourg)*<http://www.eif.ch/~kuonen/>

Adaptive utilization of resources in a highly heterogeneous computational environment such as the Grid is a difficult question. We address an object-oriented approach to the solution an original model of parallel objects. Each parallel object is a self-described, shareable and passive object that resides in a separate memory address space. The allocation of the parallel object is driven by the constraints on the resource on which the object will live. A new parallel programming paradigm is presented in the context of a new parallel object-oriented programming environment (called ParoC++) for high performance distributed computing. ParoC++ extends the C++ for supporting our new parallel objects model. It includes a runtime system that provides services to run ParoC++ programs in distributed heterogeneous environments.

Keywords: *Programming model, GRID-Computing, Object oriented model*

Joined work of: *Kuonen, Pierre; Tuan Anh Nguyen*

Towards Next Generation Grids: Research needed and Open Issues*Laforenza, Domenico (CNR - Pisa)*<http://www.miles.isti.cnr.it/~lafo/domenico/domenico.html>

This talk aims to present the outcome of a group of independent experts convened by the European Commission with the objective to identify potential European Research priorities for Next Generation Grid(s) in 2005 - 2010 [1]. The Next Generation Grid Properties ("The NGG Wish List") will be presented. The current Grid implementations do not individually possess all of the properties reported in the NGG document. However, future Grids not possessing them are unlikely to be of significant use and, therefore, inadequate from both research and commercial perspectives. In order to realise the NGG vision much research is needed.

During the last few years, several new terms such as Global Computing, Ubiquitous Computing, Utility Computing, Pervasive Computing, On-demand Computing, Autonomic Computing, Ambient Intelligence, etc., have been coined. In some cases, these terms describe very similar computational approaches. Consequently, some people are raising the following questions: Are these computational approaches facets of the same medal? The last part of this talk will explore the relationship of these approaches with Grid.

References: 1. *EU Expert Group Report, Next Generation Grid(s) 2005 - 2010, Brussels, June 2003, ftp://ftp.cordis.lu/pub/ist/docs/ngg-eg_final.pdf*

Keywords: *Next Generation Grids, New Research Issues, EU Strategies*

Event-driven workflows

Lee, Craig (Aerospace Corp. - El Segundo)

This talk examines the integration of content-based event notification systems with workflow management. This is motivated by the notion of Dynamic, Data-Driven Application Systems which endeavor to dynamically discover, ingest, and interact with sensors and actuators across the spectrum of physical systems. This requires workflows that can be dynamically reconfigured on-the-fly on the receipt of important events.

Such a capability also supports fault tolerance, i.e., reconfiguring workflows on the receipt of failure events. Key technologies for content-based event notification are examined, e.g., distributed hash tables, and composable name spaces. Centralized and decentralized workflow engines are examined along with relevant programming paradigms, e.g., active messages, agent coordination languages, etc.

The key observation is that systems truly providing autonomic, reconfigurable workflows should not have to rely on any form of priori knowledge, i.e., static information that is “compiled-in” to an application. The increased reliance on only a posteriori information will require semantic analysis and planning to reach abstract goal states – classic topics in artificial intelligence. These observations indicate that FGs could be pushed into more declarative programming methods and the use of artificial intelligence. For practical grid systems, these are risk areas that must be approached carefully.

Remote administration and fault tolerance in large distributed computer infrastructures

Lindenstruth, Volker (Universität Heidelberg)

<http://www.ti.uni-hd.de>

Independent on the level of built in resilience large distributed computer infrastructures will become unreliable if scaled to an appropriate size. Experiments, currently prepared in the physics field, foresee for instance PC clusters

of sizes reaching 10k nodes. A remote sensing device is outlined and a fault tolerance and monitoring infrastructure is presented, which was developed, having named applications in mind, in order to enable self healing autonomous compute clusters.

Keywords: *autonomous computing, clusters, remote administration*

A general scheme for resource matching in the context of global computing

Manneback, Pierre (Faculté Polytechnique de Mons)

The resource matching problem in global computing environment involves assigning volatile resources to tasks in order to satisfy as best as possible task requirements. The volatility and heterogeneity of resources have to be taken into account to ensure a good match. We propose a simple scheme which partitions dynamically the resources and the tasks into hierarchical classes, depending on statistical analysis.

Resources request corresponding tasks if available. A downgrading or upgrading of tasks is also proposed when some classes are empty. This scheme is compared to FCFS approach with the SimGrid simulator. Results show that the number of completed tasks in a given duration is much higher with the proposed method.

Keywords: *Resource matching, global computing, grid computing*

Joined work of: *Beauwens, Jean-François; Manneback, Pierre*

Non - Centralized Grid Infrastructure Monitoring

Matyska, Ludek (Masaryk University)

Monitoring of Grid infrastructure is an indispensable part of the Grid management. However, scalability of centralized approaches is rather limited, both in terms of central site(s) overload and slow reaction time. We present a new approach based on a distributed monitoring infrastructure using application-like agents (“worms”) managed by shepherds capable of autonomous behavior. It is combined with a middleware instrumentation, that allows to collect information related specifically to actual Grid use patterns. To cover various security and reliability requirements, we extend the basic Grid Monitoring Architecture (GMA) with Quality of Service like features. The combination is flexible and scalable environment for grid infrastructure monitoring.

Keywords: *Grid monitoring architecture, Grid monitoring scalability, Distributed Grid monitoring*

GAT - API for scientific Applications on the Grid

Merzky, André (Vrije Universiteit Amsterdam)

<http://www.zib.de/merzky>

Core grid technologies are rapidly maturing, but there remains a shortage of real grid applications. One important reason is the lack of a simple and high-level application programming toolkit, bridging the gap between existing grid middleware and application-level needs. The Grid Application Toolkit (GAT), as currently developed by the EC-funded project GridLab, provides this missing functionality. As seen from the application, the GAT provides a unified simple programming interface to the grid infrastructure, tailored to the needs of grid application programmers and users. A uniform programming interface will be needed for application developers to create a new generation of ‘grid-aware’ applications. The GAT implementation handles both the complexity and the variety of existing grid middleware services via so-called adaptors.

We present the GridLab software architecture, consisting of the GAT, and environment-specific adaptors. We elaborate the concepts underlying the GAT and outline the corresponding API, and demonstrate how a dynamic grid application can easily benefit from the GAT. All GridLab software is open source and can be downloaded from the project web site.

Defining, Implementing, Deploying and Using a Parallel Component Model

Perez, Christian (INRIA Rennes)

<http://www.irisa.fr/paris/pages-perso/Christian-Perez/>

The potential computational power of grids is very attractive for code coupling applications. However, the grid complexity should be hidden as much as possible by a programming model while still providing efficient execution. Component models promise to handle such a complexity but some issues are not yet tackled by standard component models. First, a parallel component model is needed in order to embed a parallel code (for example a MPI code) into a component. A key issue is to efficiently handle data redistribution that may occur in a parallel to parallel component communication. For example, such communications should be scalable and should not saturate the network.

Second, while some component models, such as CCM, provide a deployment model, they do not deal with the actual selection of resources where to install components. A grid-aware deployment model is needed to find and select resources and to compute a mapping of components onto them according to some criteria (time, cost, ...). Last, communications need to be abstracted from the actual network technology to remove a burden from the user (no requirement to compile the codes for every technology) and to offer more flexibility to scheduler. Hence, network technology constraints should be replaced by communication property constraints.

These different issues are being investigated in three prototypes.

GridCCM is a parallel extension of CCM, Adage is a grid-aware component deployment tool and PadicoTM is an open integration framework for communication middleware and runtimes.

Keywords: *component model, middleware, runtime, component deployment, communication*

GRIDLE: A Way to Compose Applications out of Software Components on the Grid

Puppin, Diego (CNR - Pisa)

<http://hpc.isti.cnr.it/~diego>

The Grid and its related technologies enable large-scale sharing of resource of various type. We envision that in the near future applications will be completely built in a bottom-up fashion using software components residing on various locations and interconnected to form a workflow graph. Following this direction, in this paper we make some proposals on the design of a component search service, enabling users to locate the components they need to deploy an application.

Joined work of: *Puppin, Diego; Silvestri, Fabrizio; Orlando, Salvatore; Laforrenza, Domenico*

Autonomic Computing - A Prerequisite for Grid Computing

Reinefeld, Alexander (ZIB and Humboldt-Universität zu Berlin)

<http://www.zib.de/reinefeld/>

“A distributed system is one in which the failure of a computer you didn’t even know existed can render your own computer unusable”. Leslie Lamport’s famous statement is especially true for Grid systems. The need to integrate many independent and heterogeneous subsystems into a well-organized Grid introduces new levels of complexity. IBM tries to tackle the problem with the often-quoted five “selves”: self configuration, self-healing, self-optimization, self-protection, and, as a combination of all, self-management. Early prototype implementations exist, but they are mostly confined to simple domains with one-dimensional optimization regimes.

Control loops, which are known in optimization theory since (at least) the mid of the last century, can be used to apply the five selves. They split the optimization process into (1) a monitoring phase, where the symptoms are collected, (2) an analysis phase, where the current status is checked against the goal status, (3) a planning phase, where a plan is created to enact the desired alterations according to some policies, and (4) the execution phase, which provides the mechanisms to schedule and perform the necessary changes to the system.

This seems to be a reasonable approach that may be applied in a straightforward manner. But is it enough for truly autonomic Grids? To our opinion,

autonomic computing also means to be able to cope with unknown errors – at least to a certain extent. What we mean with this can best be illustrated by a simple example: When an error occurs in the execution of a computer program the program crashes. It is the users task to decide what to do next. But there exist also simple cases where the system can decides by itself. In a divide-by-zero exemption, for example, the runtime system issues a trap and a pre-determined code sequence is executed. The user is freed from the need to intervene. Similarly, in Grid environments many error cases can be handled autonomously by selecting and enforcing pre-determined action schemes. The main difference to existing solutions lies in the fact that policies in the Grid are made external and they are executed by independent mechanisms or agents. So, if a program or a Grid node crashes, let it just crash. Let somebody else do the cleanup.

There is another property that makes it difficult to design and deploy autonomic computing systems: Existing computer systems are designed and implemented through a well-defined software development process that tries to plan the necessary actions on each possible error situation in advance. In other words, classical system architecture is a design-time artifact. Autonomic behavior, in contrast, is a run-time property. This exactly makes it difficult to develop autonomic Grid systems.

Keywords: *autonomic systems, grid computing*

The ZIB Distributed Data Management System

Schintke, Florian (ZIB)

<http://www.zib.de/schintke/>

ZIBDMS is a system for distributed data management which is focused on the efficient management of high volume replicated scientific data sets in the Grid. It uses the peer-to-peer (P2P) paradigm in the way that single components of ZIBDMS are entire P2P systems by themselves. To support flexible schemas of metadata ZIBDMS does not use a static database schema to store metadata, but allows to have arbitrary key-value pairs as metadata. Included in the system are aspects of autonomic computing, that for example, can ensure a certain file availability by actively creating additional replicas, if this becomes necessary. An analytical model is used inside peers to calculate the number of replicas that are needed to reach the user specified file availability.

Keywords: *distributed data management, peer-to-peer computing, autonomous computing, availability*

Designing Services for Grid-based Knowledge Discovery

Talia, Domenico (Università della Calabria)

<http://si.deis.unical.it/~talia>

The use of computers is changing our way to make discoveries and is improving both speed and quality of the discovery processes. In this scenario the Grid can provide an effective computational support for distributed knowledge discovery from large data sets. To this purpose we designed a system called Knowledge Grid. This talk describes how to design and implement distributed knowledge discovery services, according to the OGSA model, by using the Knowledge Grid environment starting from searching Grid resources, composing software and data elements, and executing the resulting application on a Grid.

Keywords: *data mining, distributed knowledge discovery, knowledge grid*

Automatic Reconfigurable Component Composition Model for Grid Systems

Thiyagalingam, Jeyarajan (Imperial College - London)

<http://www.doc.ic.ac.uk/~jeyan/>

Contemporary Grid systems offer services through assemblies of software components/services operating on resources, rarely tracking and adapting to changes in the system. Adaptive software components and services are necessary to adapt accordingly to the changes in a Grid system. Future generation Grid systems should exhibit the property of “autonomic adaptation to changes”. Adaptive software components and services can aid in achieving this property. However, automatically tracking and adapting to changes in the presence of large number of adaptive software components, requires a clever software component composition mechanism. The research has demonstrated that many applications can benefit from the grid infrastructure, including collaborative engineering, data exploration, high throughput computing, and of course distributed super-computing.

In this paper, we present a model for component composition to enable runtime composition of adaptive components and services, given that the composition is governed by a set of constraints. We then discuss the possible software technologies, which permit us to implement autonomic adaptivity in future generation Grid systems.

Joined work of: *Thiyagalingam, Jeyarajan; Getov, Vladimir; Kelly, Paul*

Automatic Application Execution on Grids

Tonello, Nicola (CNR - Pisa)

<http://www.isti.cnr.it/php-pers/iselpers.php?Tonello+Nicola>

The research has demonstrated that many applications can benefit from the grid infrastructure, including collaborative engineering, data exploration, high throughput computing, and of course distributed supercomputing.

This benefit is somewhat weakened by the fact that writing grid applications as well as porting existing ones to the grid is a difficult and often tedious and error-prone task. lifetime.

Our approach intends to automatize the common tasks needed to start grid-enabled and grid-aware applications, in order to allow the largest user community to gain the full benefit from the grid, and at the same time to give the maximum generality and applicability.

This approach, combined with the adoption of high-level programming tools, can greatly simplify the task of writing grid-aware applications.

The ASSIST programming environment can help to write applications capable of adapting to dynamic resource availability, because it a reconfigurable runtime support that can interact with the proposed resource discovery and selection automatic mechanisms.

The Grid Abstract Manager presented here can seamlessly run complex grid applications, taking care of all the details needed to interface the middleware: resource discovery and selection, data and executable files transfers, interaction with local GRAMs, execution monitoring and result collection.

Joined work of: *Tonello, Nicola; Fagni, Tiziano; Paccosi, Alessandro; Zoccolo, Corrado*

Distributed Query Processing over Dynamically Changing Resources

Watson, Paul (University of Newcastle)

<http://www.cs.ncl.ac.uk/people/home.php?name=paul.watson>

Distributed Query Processing supports the integration and processing of distributed information through the use of declarative query languages. In this talk, we will show how this can hide from the client concerns such as: programming, location, communication, acquisition and exploitation of dynamic resources, coordination and parallelisation. Some recent work on fault-tolerance will also be described.

Keywords: *Databases, Distributed Query Processing, Fault-Tolerance, Parallelism, Data*

Grid Resource Management for Future Generation Grids

Yahyapour, Ramin (Universität Dortmund)

<http://www-ds.e-technik.uni-dortmund.de/>

Current research in the area of Grid resource management produces a rich set of architectural approaches and corresponding algorithms. Most of this research is based on the background of scientific computing and often finds its actual implementations in dedicated and sometimes specialized Grids. It has to be further discussed what the actual requirements will be for Future Generation Grids. Clearly, quite different application scenarios come from other areas as from commercial environments with business services or as from ubiquitous computing. While the scenarios seem to significantly deviate from scientific computing, a lot of similarities can be seen. The convergence of these approaches could be a key factor for the success of future generation grids. It is acknowledged that the different application domains will have special requirements that will yield to different implementations. However, the continuous re-invention of several different architectures is highly inefficient but may also be necessary. The question is raised whether the proliferation of Grid technologies can be increased if common interfaces for core functionalities are identified. That is, different architectures might be implemented but the inter-operability between these systems is maintained. Approaches as in SOA, OGSA, CoG can be seen as examples in that direction. As a starting point for discussion, several core functionalities for Grid resource management have been identified in this talk.

Keywords: *Grids; Scheduling; Resource Management*

Joined work of: *Yahyapour, Ramin; Schwiigelshohn; Uwe*

Bringing Knowledge to Middelware : Grid Scheduling Ontology

Ziegler, Wolfgang (Fraunhofer Institut - St. Augustin)

<http://www.scai.fraunhofer.de/ziegler.0.html>

Hiding complexity of Next Generation Grids from the end-user is one major task to be solved to make the use of Grids a real option. Relieving the system administrator of the burden to manually tweak his Grid system to integrate and deal with external, unknown resources is the other side same medal making Grids a technology for a broader community than the usual experts. As the Grids will more and more stretch across heterogeneous hardware and operating systems located in different administrative domains scheduling of a single resource is already becoming a challenging task. Automatic scheduling of several resources with different properties to execute a complex workflow with temporal dependencies is not possible if the resources and their capabilities and constraints are not known in advance. This knowledge has several facets: the site-dependent way how resources are described, how local scheduling systems describe and manage resources and how a super-scheduler tries to map the single user request to the individual requests addressed to these local scheduling

systems. At this point resource ontologies come into play as potentially useful instruments providing a sophisticated approach to categorizing and drawing relationships between the various ways resources and services are described today, finally leading to a shared and common understanding of resources and services. Here a local scheduler or super-scheduler is just another resource (or service) with dedicated attributes and behaviour. A Scheduling Domain Ontology seems to be a promising way to introduce knowledge exploitable by machines into the scheduling process in Grids. Given such an ontology shared among the resource brokers, resource management systems and schedulers finding resources or services and mapping user requests to resources in the Grid may become less arbitrarily. At the same time negotiation between sites with different resources available under different scheduling policies will become more distinct and may be carried out by automatically by a super-scheduler.

Keywords: *Grid Scheduling Ressource Ontology Meta-Scheduling*

6 Participants

Marco Aldinucci
Consiglio Nazionale delle Ricerche
Istituto di Scienza e Tecnologie per l'Informazione (ISTI)
Via Moruzzi 1
I-56124 Pisa, (I)
Tel: +39-050-2212-728
Fax: +39-050-2212-726
E-Mail: aldinuc@di.unipi.it
URL: <http://www.di.unipi.it/~aldinuc>

Artur Andrzejak Konrad-Zuse-Zentrum für Informationstechnik
Computer Science Research
Takustr. 7
D-14195 Berlin, (D)
Fax: +49-30-8418-5311
E-Mail: andrzejak@zib.de
URL: <http://www.zib.de/andrzejak/>

Gabriel Antoniu
IRISA
Campus de Beaulieu
F-35042 Rennes, (F)
Fax: +33-2-9984-7244
E-Mail: gabriel.antoniu@irisa.fr
URL: <http://www.irisa.fr/paris/pages-perso/Gabriel-Antoni>

Olivier Aumage

Université Bordeaux I
LaBRI
351 Cours de la Libération
F-33405 Talence, (F)
Fax: +33-5400-066-69
E-Mail: aumage@labri.fr
URL: <http://dept-info.labri.fr/~aumage/>

Luc Bougé
IRISA - Ecole Normale Supérieure de Cachan
Campus Ker Lann
Avenue Robert Schumann
F-35170 Bruz, (F)
Tel: +33-2-9905-5283
Fax: +33-2-9984-7171
E-Mail: Luc.Bouge@bretagne.ens-cachan.fr
URL: <http://www.bretagne.ens-cachan.fr/DIT/People/Luc.Bouge/>

Marian Bubak AGH - University of Science & Technology
Department of Computer Integrated Manufacturing
al. Mickiewicza 30
PL-30-059 Krakow, (PL)
Tel: +48-12-617-3964
Fax: +48-12-633-9406
E-Mail: bubak@uci.agh.edu.pl
URL: <http://www.icsr.agh.edu.pl>

Sonia Campa
Università degli Studi di Pisa
Dipartimento di Informatica
Via Buonarroti 2
I-56127 Pisa, (I)
Tel: +39-050-221-2797
Fax: +39-050-221-2726
E-Mail: campa@di.unipi.it
URL: <http://www.di.unipi.it/~campa/>

Franck Cappello
INRIA Futurs
4 rue Jean Monod
ZAC des Vignes
F-91893 Orsay, (F)
Fax: +33-169-156-586
E-Mail: fci@lri.fr
URL: <http://www.lri.fr/~fci/>

Denis Caromel
INRIA
Sophia Antipolis
I3S - CNRS URA 1376 - Projet OASIS
2004 Route des Lucioles
B.P. 93
F-06902 Sophia Antipolis, (F)
Tel: +33-4-9238-7631
Fax: +33-4-9238-7971
E-Mail: caromel@unice.fr
URL: <http://www.inria.fr/oasis/denis>

Marco Danelutto
Università degli Studi di Pisa
Dipartimento di Informatica
Via Buonarroti 2
I-56127 Pisa, (I)
Tel: +39-050-22-127-42
Fax: +39-050-22-127-26
E-Mail: marcod@di.unipi.it
URL: <http://www.di.unipi.it/~marcod/>

Frédéric Desprez
Ecole Normale Supérieure de Lyon
Laboratoire LIP
46 allée d'Italie
F-69364 Lyon Cedex 07, (F)
Fax: +33-4-72 72 80 80
E-Mail: Frederic.Desprez@ens-lyon.fr
URL: <http://graal.ens-lyon.fr/~desprez/>

Dick H.J. Epema
Delft University of Technology
Dept. of Electrical Engineering
Postbus 3051
Mekelweg 4
P.O. Box 5031
NL-2600 GA Delft, (NL)
Tel: +31-15-278-3853
Fax: +31-15-278-7141
E-Mail: epema@ewi.tudelft.nl
URL: <http://www.pds.ewi.tudelft.nl/~epema/>

Vladimir Getov

University of Westminster
School of Computer Science
Northwick Park
Watford Road
HA1 3TP London-Harrow, (GB)
Tel: +44-171-911-5917
Fax: +44-171-911-5906
E-Mail: v.s.getov@wmin.ac.uk
URL: <http://perun.hscs.wmin.ac.uk/~vsg/>

Sergei Gorlatch
Westfälische Wilhelms-Universität Münster
FB 10 Mathematik und Informatik
Einsteinstr. 62
D-48149 Münster, (D)
Tel: +49-251-833-2741
Fax: +49-251-833-2742
E-Mail: gorlatch@math.uni-muenster.de
URL: <http://pvp.cs.tu-berlin.de/gorlatch/>

Armin Größlinger
Universität Passau
FB Informatik
LST Prof. Ch. Lengauer
D-94030 Passau, (D)
Tel: +49-851-509-3073
Fax: +49-851-509-3092
E-Mail: groessli@fmi.uni-passau.de
URL: <http://www.infosun.fmi.uni-passau.de/cl/staff/groesslinger/>

Hans-Ulrich Heiß
TU Berlin
Fakultät IV - ETTS / FOKUS
Franklinstr. 28-29
D-10623 Berlin, (D)
Tel: +49-30-314-73160
Fax: +49-30-314-25156
E-Mail: heiss@cs.tu-berlin.de
URL: <http://kbs.cs.tu-berlin.de/staff/heiss/heiss.htm>

Thomas Herault
Université Paris Sud
Laboratoire de Recherche en Informatique
Bat. 490
F-91405 Orsay, (F)

Tel: +33-1-6915-4222
Fax: +33-1-6915-6586
E-Mail: herault@lri.fr
URL: <http://www.lri.fr/~herault/>

Guillaume Huard
Laboratoire ID-IMAG
Information & Distribution 51
Avenue J. Kuntzmann
F-38330 Montbonnot St. Martin, (F)
Tel: +33-476-612-053
Fax: +33-476-612-099
E-Mail: Guillaume.Huard@imag.fr
URL: http://www-id.imag.fr/Laboratoire/Membres/Huard_Guillaume

Stavros Isaiadis
University of Westminster
School of Computer Science
Watford Road
HA1 3TP London-Harrow, (GB)
Tel: +44-77-6619-7351
Fax: +30-231-0888-646
E-Mail: s.isaiadis@westminster.ac.uk

Emmanuel Jeannot
INRIA - LORIA
Université Henri Poincaré
BP 239
F-54506 Vandoeuvre-lès-Nancy, (F)
Fax: +33-3-8327-8319
E-Mail: ajeannot@loria.fr
URL: <http://www.loria.fr/~ajeannot>

Thilo Kielmann
Vrije Universiteit Amsterdam
Dept. of Computer Science
De Boelelaan 1081a
NL-1081 HV Amsterdam, (NL)
Tel: +31-20-4447-789
Fax: +31-20-4447-653
E-Mail: thilo.kielmann@acm.org
URL: <http://www.cs.vu.nl/~kielmann/>

Pierre Kuonen
University of Applied Sciences of Fribourg
Bd. de Pérolles 80
1705 Fribourg, (CH)
Tel: +41-26-429-6565
Fax: +41-26-429-6600
E-Mail: pierre.kuonen@eif.ch
URL: <http://www.eif.ch/~kuonen/>

Domenico Laforenza
Consiglio Nazionale delle Ricerche
Istituto di Scienza e Tecnologie per l'Informazione (ISTI)
Via Moruzzi 1
I-56124 Pisa, (I)
Tel: +39-050-315-2992
Fax: +39-050-313-8091
E-Mail: domenico.laforenza@cnuce.cnr.it
URL: <http://www.miles.isti.cnr.it/~lafo/domenico/domenico.html>

Craig Lee
Aerospace Corporation
2350 East El Segundo Bl.
CA 90245 El Segundo, (USA)
Fax: +1-310-336-4402
E-Mail: lee@aero.org

Volker Lindenstruth
Universität Heidelberg
Kirchhoff-Institut für Physik
Im Neuenheimer Feld 227
D-69120 Heidelberg, (D)
Tel: +49-6221-549-801
Fax: +49-6221-549-809
E-Mail: ti@kip.uni-heidelberg.de
URL: <http://www.ti.uni-hd.de>

Jon MacLaren
The University of Manchester
Manchester Computing
Kilburn Building
Oxford Road
M13 9PL Manchester, (GB)
Tel: +44-161-275-6821
E-Mail: Jon.Maclaren@manchester.ac.uk

Pierre Manneback
CETIC
Rue Clément Ader, 8
6041 Charleroi, (B)
Tel: 065-374-050
Fax: 065-374-500
E-Mail: pierre.manneback@fpms.ac.be
URL: <http://www.cetic.bewww.fpms.ac.be>

Ludek Matyska
Masaryk University
Faculty of Informatics
Botanicka 68a
CZ-60200 Brno, (CZ)
Tel: +420-5-4949-2105
Fax: +420-5-4122-747
E-Mail: ludek@ics.muni.cz
URL: <http://www.data.muni.cz/people/matyska>

André Merzky
Vrije Universiteit Amsterdam
Dept. of Mathematics & Computer Science
De Boelelaan 1081 a
NL-1081 HV Amsterdam, (NL)
Tel: +49-30-841-85-339
Fax: +49-30-841-85 107
E-Mail: andre@merzky.net
URL: <http://www.zib.de/merzky>

Raymond Namyst
Université Bordeaux I
LaBRI
351 Cours de la Libération
F-33405 Talence, (F)
Fax: +33-54-000-6669
E-Mail: raymond.namyst@labri.fr
URL: <http://dept-info.labri.u-bordeaux.fr/~namyst/#contact>

Christian Perez
INRIA Rennes
IRISA
Campus de Beaulieu
F-35042 Rennes, (F)
Tel: +33-2-9984-7205
Fax: +33-2-9984-2528

E-Mail: Christian.Perez@irisa.fr
URL: <http://www.irisa.fr/paris/pages-perso/Christian-Perez/>

Thierry Priol
Université de Rennes
Campus de Beaulieu
F-35042 Rennes, (F)
Tel: +33-2-9984-7210
Fax: +33-2-9984-2528
E-Mail: thierry.priol@irisa.fr
URL: <http://www.irisa.fr/paris/pages-perso/Thierry.Priol/anglais/welcome.html>

Diego Puppini
Consiglio Nazionale delle Ricerche
Istituto di Scienza e Tecnologie per l'Informazione (ISTI)
Via Moruzzi 1
I-56124 Pisa, (I)
Tel: +39-050-315-3003
Fax: +39-050-313-8091
E-Mail: diego.puppini@isti.cnr.it
URL: <http://hpc.isti.cnr.it/~diego>

Alexander Reinefeld
Konrad-Zuse-Zentrum für Informationstechnik
Computer Science
Takustr. 7
D-14195 Berlin, (D)
Tel: +49 30 84185-130
Fax: +49 30 84185-311
E-Mail: reinefeld@zib.de
URL: <http://www.zib.de/reinefeld/>

Florian Schintke
Konrad-Zuse-Zentrum für Informationstechnik
Computer Science Research
Room 3155
Takustr. 7
D-14195 Berlin, (D)
Tel: +49-30-841-85-306
Fax: +49-30-841-85-311
E-Mail: schintke@zib.de
URL: <http://www.zib.de/schintke/>

Uwe Schwiegelshohn
Universität Dortmund

FB Elektro- und Informationstechnik
Otto-Hahn-Str. 4
D-44221 Dortmund, (D)
Tel: +49-231-755-26 34
Fax: +49-231-755-32 51
E-Mail: uwe.schwiegelshohn@udo.edu
URL: <http://www-ds.e-technik.uni-dortmund.de/new/CEI/de/staff/schwiegelshohn.shtml>

Raül Sirvent
Universitat Politècnica de Catalunya
Dept. d'Arquitectura de Computadors
Jordi Girona 1-3
Modul C6
E-08034 Barcelona, (E)
Tel: +34-9340-11650
Fax: +34-9340-17055
E-Mail: rsirvent@ac.upc.es

Domenico Talia
Università della Calabria
(DEIS)
Via Pietro Bucci, 41C-42C
I-87036 Rende, (I)
Tel: +39-0984-49 47 26
Fax: +39-0984-49 4713
E-Mail: talia@deis.unical.it
URL: <http://si.deis.unical.it/~talia>

Jeyarajan Thiyagalingam
Imperial College
London
E-Mail: jeyan@doc.ic.ac.uk

Nicola Tonellotto
Consiglio Nazionale delle Ricerche
Istituto di Scienza e Tecnologie per l'Informazione (ISTI)
Via Moruzzi 1
I-56124 Pisa, (I)
Tel: +39-050-315-3003
Fax: +39-050-313-8091
E-Mail: Nicola.Tonellotto@isti.cnr.it
URL: <http://www.isti.cnr.it/php-pers/iselpers.php?Tonellotto+Nicola>

George Tsouloupas

University of Cyprus
Dept. of Computer Science
75 Kallipoleos Street
P.O. Box 20537
CY-1678 Nicosia, (CY)
Fax: +357-2289-2701
E-Mail: georget@ucy.ac.cy

Paul Watson
University of Newcastle upon Tyne
School of Computing Science
Claremont Tower
NE1 7RU Newcastle-upon-Tyne, (GB)
Tel: +44-191-222-7653
Fax: +44-191-222- 8232
E-Mail: paul.watson@ncl.ac.uk
URL: <http://www.cs.ncl.ac.uk/people/home.php?name=paul.watson>

Ramin Yahyapour
Universität Dortmund
FB Elektro- und Informationstechnik
LST Datenverarbeitungssysteme
Otto-Hahn-Str. 4
D-44221 Dortmund, (D)
Tel: +49-231-755-27 35
Fax: +49-231-755-32 51
E-Mail: yahya@ds.e-technik.uni-dortmund.de
URL: <http://www-ds.e-technik.uni-dortmund.de/>

Wolfgang Ziegler
Fraunhofer Institut
SCAI - Inst. for Algorithms & Scientific
Computing Schloss Birlinghoven
D-53754 St. Augustin, (D)
Tel: +49-2241-14-2258
Fax: +49-2241-14-42258
E-Mail: wolfgang.ziegler@scai.fraunhofer.de
URL: <http://www.scai.fraunhofer.de/ziegler.0.html>

Corrado Zoccolo
Università degli Studi di Pisa
Dipartimento di Informatica
Via Buonarroti 2
I-56127 Pisa, (I)

E-Mail: zoccolo@di.unipi.it

URL: <http://www.di.unipi.it>

7 Further Information

The list of participants with abstracts of their talks and additional material (e.g., participants group picture) can be found on the seminar homepage <http://www.dagstuhl.de/04451/>.