IST-2000-25406

Contribution to IST-2000-7.1.2:

End-to-End Application Experiments

DAMIEN

Distributed Applications and Middleware for Industrial Use of European Networks

D9.5 / Pallas

2nd Dissemination Report

All Project Partners

Editor:
Ullrich Becker-Lemgau Pallas GmbH

Classification:

Status:	public
Version:	1.0
Planned for:	PM24
Actual:	PM25
Distribution:	Project Reviewers
	Project Partners

Table of Contents

1. Introduction

DAMIEN is one of the GRID projects funded by the European Commission. Due to the project structure interest can be generated about the DAMIEN GRID-technology in lots of different areas. High Performance Computing and Metacomputing states one of this areas which was mainly targeted during the reporting period of this Dissemination Report. Concepts, structure and especially the tools of the DAMIEN Metacomputing environment were disseminated through several talks, papers, and exhibitions.

Additionally DAMIEN supports the GRIDSTART project with publications and during dissemination actions.

This document will list the dissemination events in detail. Additionally dissemination material used at the exhibitions are attached.

2. Dissemination actions

2.1 Talks

09.01.2002-10.01.2002: ESTEC Workshop on Computational Integrated Design and Analysis for Space Vehicles (CINDA), Noordwjik, The Netherlands:

DAMIEN and the role of MpCCI for Metacomputing was presented by Karl Solchenbach at the CINDA Workshop held at ESA-ESTEC. Grid Computing and especially coupled applications will be a key technology for analysis of space vehicles. The talk raised great interest among the participants of this workshop. The coupling technology of MpCCI for distributed systems developed in DAMIEN is considered to be part of the CINDA initiative of ESA-ESTEC.

28.02.2002: University of Houston, USA:

Presentations by Peggy Lindner about the Configuration Manager.

28.02.2002: MpCCI User Forum 2002, Sankt Augustin, Germany:

Ullrich Becker-Lemgau held a presentation in the research session of the MpCCI User Forum with the title: “DAMIEN – Metacomputing at work”.

29.03.2002: University of Houston:

Presentation by Rainer Keller (HLRS) about ‘Communication in GRID Environment with PACX-MPI’.

14.05.2002-17.05.2002: 2nd Asian Pacific Gridforum Workshop (APGRID), Taipei, Taiwan:

Edgar Gabriel (HLRS) gave a one hour presentation about the DAMIEN project presenting the overall goals of the project, the tools developed in the frame of DAMIEN as well as the current status.

28.05.2002: 1st Public DAMIEN Workshop, Stuttgart, Germany:

Several talks about the project DAMIEN and the tools and technologies developed within DAMIEN were given during this event:
- Edgar Gabriel (HLRS): DAMIEN Overview
- Laurent Vervisch (CRIHAN): Technology and Status of QoS-Manager
- Karl Solchenbach (Pallas): MpCCI technology for the meta-computing environment
- Rosa M. Badia (Cepba): Performance prediction for grid-environments
- Hans-Christian Hoppe (Pallas): Performance analysis with Vampir for the GRID.
- Rainer Keller (HLRS): Communication Library PACX-MPI
- Peggy Lindner (HLRS): Configuration Manager for PACX-MPI applications.

24.06.2002-25.06.2002: TROPHY Project Meeting, Brussels, Belgium:

Ullrich Becker-Lemgau (Pallas) gave a presentation about MpCCI at a project meeting of TROPHY, also funded by the European Commission. During this presentation also the role of MpCCI in DAMIEN was described.

11.07.2002-12.07.2002: 2nd Grid Forum Korea Meeting, South Korea:

Edgar Gabriel (HLRS) was invited to give a tutorial about “Problem Solving Environments and Application Development in the Grid”, including a one hour DAMIEN project presentation. This consisted of presenting the overall goals of the project and the tools developed in the frame of DAMIEN.

13.07.2002: Seoul National University, Seoul, South Korea:

Presentation by Edgar Gabriel, The DAMIEN Software Development Environment

15.07.2002: Korean Institute of Science and Technology, Daejeon, South Korea:

Presentation by Edgar Gabriel, The DAMIEN Software Development Environment.

18.08.2002-23.08.2002: Workshop on Performance Analysis and Distributed Computing, Dagstuhl, Germany:

For this workshop Hans–Christian Hoppe was invited to present a paper on "Performance Analysis in GRID Environments", presenting the work on performance analysis and evaluation within the DAMIEN project. The workshop was organized by the Technical University of Munich and the University of Wisconsin at Madison in the US, and a total of 50 well–known researchers in the field had been invited.

30.10.2002: Inauguration of the GRID-center at Forschungszentrum Karlsruhe, Karlsruhe, Germany:

Michael M. Resch (HLRS) gave a presentation with the title 'Middleware: The Heart of GRID-Computing'.

11.12.2002: JAERI, Tokio, Japan:

Michael M. Resch gave a presentation about 'The State of GRID Computing in Europe and the US'.

13.12.2002: JAERI, Kyoto/Kansai, Japan:

Michael M. Resch gave a presentation about "The Bennefits and Effectiveness of GRID Computing".

09.12.2002-10.12.2002: Grid Benchmarking Workshop, Edinburgh, United Kingdom:

Jesus Labarta (CEPBA) gave a talk which was centered on performance prediction for grid environments. The presentation described the philosophy of Dimemas as a simulator to estimate the performance of parallel programs on hypothetical target machines. The presentation focused on how Dimemas functionalities are being extended as part of the project DAMIEN to model the characteristics of GRID systems.

16.12.2002-18.12.2002: Euroweb 2002, Oxford, United Kingdom:

EuroWeb 2002 is a major international forum at which research on GRIDs and Web Services is presented. Peggy Lindner (HLRS) gave a presentation about the Configuration Manager, which is one component of the DAMIEN tools.

2.2 Papers

- Ullrich Becker-Lemgau: *"DAMIEN – Metacomputing at work"*, Book of Abstracts, 3rd MpCCI User Forum, February 27th & 28th , 2002.
- Edgar Gabriel, Peggy Lindner, Rainer Keller, Matthias Müller, Michael M. Resch: "The DAMIEN software development environment", submitted to the 3rd International Conference on Grid Computing, Baltimore, USA, November 18th, 2002.
- Allan Snavely, Laura Carrington, Nicole Wolter, Jesus Labarta, Rosa M. Badia, Avi Purkayastha: "A Framework for Performance Modelling and Prediction", submitted to Supercomputing 2002 in Baltimore, USA, November 17-21, 2002.
- Peggy Lindner, Natalia Currle-Linde, Michael M. Resch and Edgar Gabriel, *"Distributed Application Management in Heterogeneous Grid"*, submitted and accepted for publication at the Euroweb 2002 conference, Oxford, UK, December 17th-18th, 2002.

- Rosa M. Badia, Francesc Escale, Edagr Gabriel, Judit Gimenez, Rainer Keller, Jesus Labarta and Matthias M. Müller *"Performance Prediction in a Grid environment"*, submitted to the 1st European Across Grids Conference, Santiago de Compostela, Spain, February 13th-14th, 2003

2.3 Exhibitions

27.02.2002-28.02.2002: 3rd MpCCI User Forum, Sankt Augustin, Germany

The MpCCI User Forum was organized by the Fraunhofer Institute SCAI. More than 70 people from industry and research interested in solutions for multidisciplinary problems were attending technical sessions and the poster exhibition. With a presentation in the research session and a poster in the poster exhibition The DAMIEN project and especially the metacomputing capabilities were presented to the audience. The abstract in the Book of Abstract will be distributed at least for one year and thereby create awareness about the DAMIEN project among industry and research institutions.

08.04.2002-10.04.2002: the 3rd Worldwide Aerospace Conference at Toulouse, France:

Pallas was present with a booth at the 3rd Worldwide Aerospace Conference and Technology Showcase. This event was held in Toulouse on April 8th to 10th and organized by MSC.Software. Besides several other topics the project DAMIEN and especially the role of MpCCI in it was presented by a slide show directly at the booth

28.05.2002: 1st Public DAMIEN Workshop, Stuttgart, Germany

Organized and hosted by HLRS on behalf of the DAMIEN project consortium the first public DAMIEN workshop was held in conjunction with the 5th HLRS Metacomputing Workshop in Stuttgart, Germany. The workshops attracted around 35 participants, the overall resonance to the DAMIEN workshop and the presentations were excellent. (see Deliverable D9.4)

19.06.2002-22.06.2002: ISC2002 Conference and Exhibition, Heidelberg, Germany:

The DAMIEN project was represented at this Exhibition by a booth together with the Eurogrid and the Datagrid project. DAMIEN was respresented with two posters, handouts and demonstrations. The exhibition attracted large attention to the audience.

21.07.2002-24.07.2002: Global Grid Forum 5, Edinburgh, United Kingdom:

GGF is the major event in the area of Grid computing. The participating partners (CEPBA: Rosa M. Badia; HLRS : Edgar Gabriel, Matthias Müller) were involved in GGF research group meetings. Furthermore, they participated in the first plenary meeting of the GRIDSTART project representing DAMIEN.

23.09.2002-26.09.2002: iGRID 2002, Amsterdam, Netherlands

Rainer Keller gave demonstrations at this event using several applications from the area of high performance computing and demonstrating tools developed in the frame of the DAMIEN project.

04.11.2002-06.11.2002: IST2002 conference and exhibition, Copenhagen, Denmark:

At the IST2002 conference and exhibition DAMIEN was presented together with other European funded GRID projects at the GRIDSTART booth (booth number P1) to inform a broad audience about the grid-activities in Europe. Among a poster and several flyers a demonstration was given. Additionally DAMIEN was presented at the Pallas booth (Title: Grid Technology at Work, booth number P3) together with other GRID related projects.

16.11.2002-22.11.2002: Supercomputing 2002, Baltimore, USA:

DAMIEN was presented with its own booth at the Supercomputing 2002 exhibition, the most important event in the field of high performance computing (approx. 7000 attendees). The booth was integrated into the European Research village next to booths from HLRS, CEPBA and the EUROGRID project and operated by DAMIEN project staff from HLRS, CEPBA, and Pallas. Two posters were produced to highlight the project goals and tools. A project flyer describing the project, the tools, the application, and the partners was produced and handed out to interested attendees. Links to the DAMIEN project and DAMIEN booth were shown on the booths of the partners HLRS, CEPBA and Pallas. Additionally Rainer Keller (HLRS) was elected to attend the HPC Challenge at SC02. This event provided the opportunity to demonstrate the DAMIEN tools.

2.4 Miscellaneous

3. Dissemination material

3.1 DAMIEN Project Flyer

DAMIEN

The Project

Although high performance computing facilities in Europe are becoming more and more powerful it is still true that numerical simulations could be faster, more accurate in time and space if the scientific community would have access to an even bigger and faster supercomputer. For each level of speed reached by a new HPC system there is a new level of complexity reached in simulations.

A very promising way to satisfy this request is to combine already existing supercomputers via high speed networks. This will form the so called Compu-tational Grids (short: Grids). But designing applications for such Grids is different because most of the stan-dards and tools available at single systems are not valid or available for distributed computing systems. The central objective of DAMIEN is therefore to provide these well known and highly accepted tools to Compu-tational Grids, too.

The Tools

All tools developed during the DAMIEN project are forming together the DAMIEN Environment for Meta-computing:

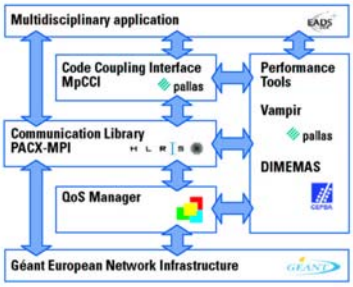

The **communication library PACX-MPI** is an MPI implementation that allows to couple distributed resources and still fully exploit their network capacity. The **Communication Manager** enables the user to configure and maintain the application on a metacom-puter.

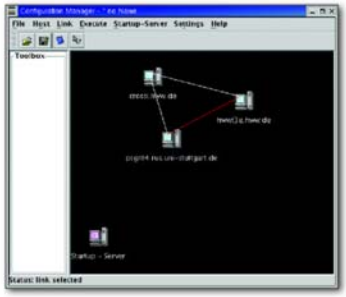

With the **code coupling interface MpCCI** two already existent codes can be combined to a coupled solver for multidisciplinary problems. The communication within MpCCI is based on PACX-MPI. Thereby the library is able to couple codes running on different systems.

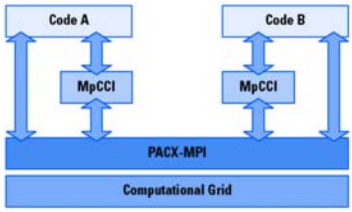

For message passing applications the **performance analysis tool Vampir** has been developed. It runs on all Unix-based workstations and analyses traces gener-ated by the **Vampirtrace** library. The library enables the detailed analysis of the performance of parallel applications and returns information like e.g. how much time is spent in communication and in computation.

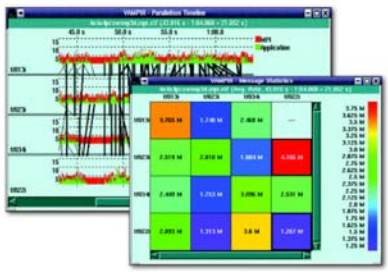

DAMIEN

The **performance prediction tool DIMEMAS** deals with the estimation of the performance that a given application will attain on a specific target machine. This information is necessary to predict how changes in the hardware configuration can affect the performance of the application in the distributed environment.

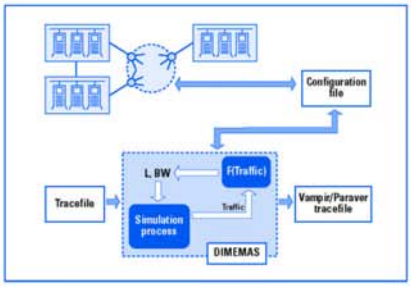

The main goal of the **Quality of Service Management (QoS)** is to provide characterisation and reservation of QoS over a network for an end user application. The QoS module will especially build on the IPv6 and IPv4 standards since these are the most important ones in wide area networks.

The Application

The **multi-physics coupling application** is provided by EADS CCR. Due to a high number of entities in the company, often remotely located, and the different domains they are involved, the whole industry group faces problems of interoperability between the teams, particularly problems of data exchange between applications.

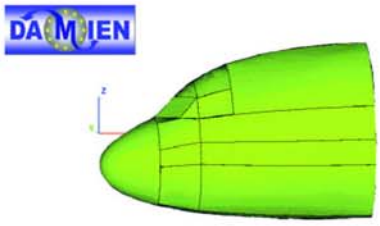

The Partners

The **German Federal High Performance Computing Center Stuttgart RUS/HLRS** is a technology and know-how provider with expertise in fields such as distributed simulation and software environments (PACX-MPI).

EADS CCR is the Corporate Research Center of AM Group. EADS acts as an end user in the project, which provide industrial applications and test cases to evaluate the industrial relevance of the project.

Centro Europeo de Paralelismo de Barcelona (CEPBA) is a research and service center within the Technical University of Catalonia (UPC). CEPBA provides the tool DIMEMAS and related technologies.

The **Centre de Ressources Informatiques de Haute-Normandie (CRIHAN)** is a regional computing resource center, known for its expertise in high-speed network technologies, Internetworking and parallel computing techniques.

Pallas GmbH is a software company specializing in the field of performance– and data–intensive computing. Pallas is supplier for high performance tools like Vampir and MpCCI and the UNICORE Grid system.

The DAMIEN project if funded by the **European Commission** under contract IST-2000-25406.

Contact

High performance Computing Center Stuttgart
Michael Resch
Email: **resch@hlrs.de**
URL: http://www.hlrs.de/organization/pds/projects/damien/

3.2 Poster for Supercomputing 2002

DAMIEN: The Project

DAMIEN

- Tools for Grid Computing
- Testbed based on High-Speed Networks
- Demonstration of Industrial Applications

Distributed Application

- Multidisciplinary Coupled Simulation
- Vibro-Acoustic Application

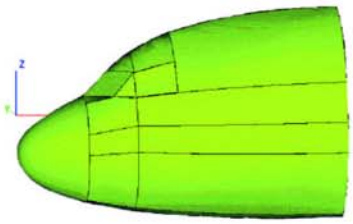

DAMIEN Project

- Funded by European Commission
- Partners: CEPBA-UPC, CRIHAN, EADS, HLRS, Pallas

DAMIEN: the Tools

Communication

- PACX-MPI: Communication Library
- Configuration Manager:
  Computational Grid Configurator
- MpCCI: Code Coupling Interface

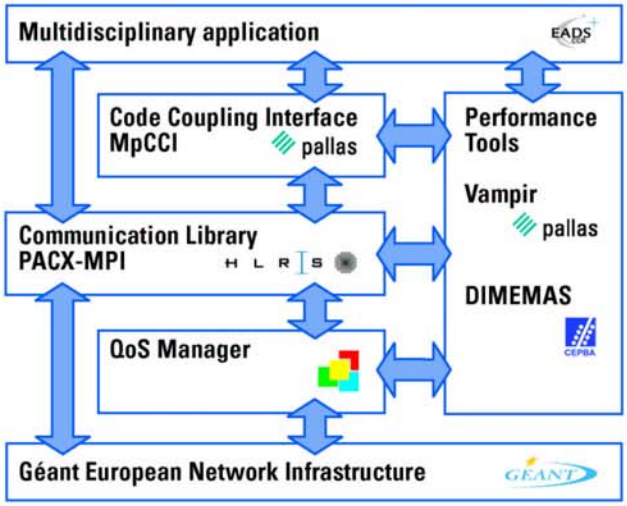

Performance

- Vampir: Performance Analysis
- Dimemas: Performance Prediction
- QoS Manager: Quality of Service for
  IPv4/IPv6

Performance Prediction: from Parallel Platforms to Grid environments

- Support for development, tuning and mapping
- Configurable architecture and network parameters
- Modeling of WANs
- Visualization with Paraver
- On your desktop workstation

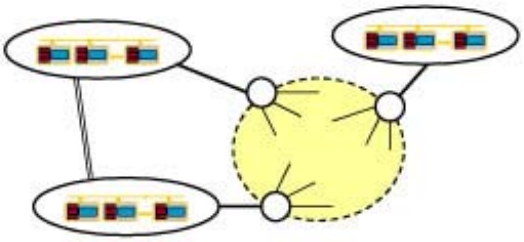

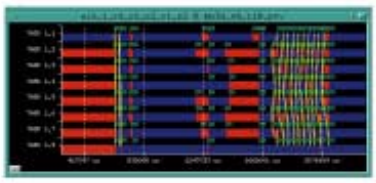

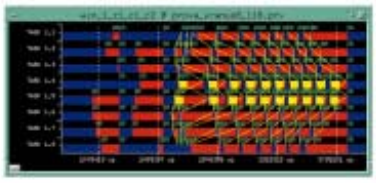

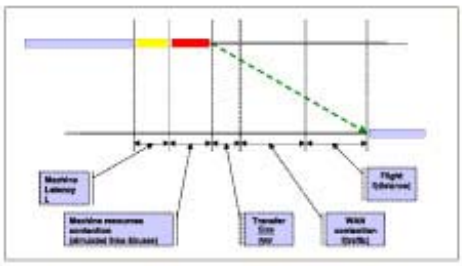

Dimemas is partially funded by DAMIEN project (IST-2000-25406)

Information Society Technologies

www.cepba.upc.es/dimemas

4. Papers

4.1 Abstract "DAMIEN – Metacomputing at work"

Published in the Book of Abstracts from the 3rd MpCCI User Forum 2002
(http://www.mpcci.org/download.htm)

pallas

DAMIEN

DAMIEN – Metacomputing at work

Dr. Ullrich Becker-Lemgau
Pallas GmbH

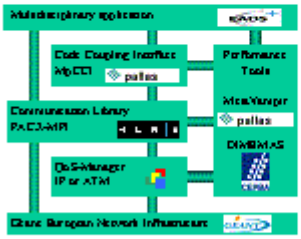

Figure 1: DAMIEN Environment for Metacomputing

Figure 2: MpCCI concept

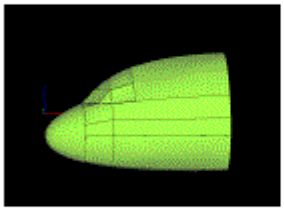

Figure 3: Airbus A320 cockpit

MCCI

The Project

Although high performance computing facilities in Europe are becoming more and more powerful it is still true that numerical simulations could be faster, more accurate in time and space if the scientific community would have access to an even bigger and faster supercomputer. For each level of speed reached by a new HPC system there is a new level of complexity reached in simulations. Multidisciplinary applications are a very good example for this: Because computers are faster and more powerful than ever we are able to look behind mono-disciplinary simulation. Our world is multi-disciplinary and this is the new target for the future. But also this will require much more powerful computers.

A very promising way to satisfy this request is to combine already existing supercomputers via high speed networks. This will form the so called Computational Grids (short: Grids). But designing applications for such Grids is different because most of the standards and tools available at single systems are not valid or available for distributed computing systems. The central objective of the project is therefore to provide these well known and highly accepted tools to Computational Grids, too.

The Tools

All tools developed during the DAMIEN project are forming together the DAMIEN Environment for Metacomputing (Fig. 1):

- Communication Library PACX-MPI
- Code Coupling Interface MpCCI
- Quality of Service Management QoS
- Performance Analysis Tool Vampir
- Performance Prediction Tool Dimemas

With the code coupling interface MpCCI two already existent codes can be combined to a coupled solver for multidisciplinary problems. MpCCI is a library which has to be linked with the simulation codes to couple (Fig. 2). It provides all the functionality required for coupled computations like communication between codes, neighbourhood search between different meshes, and interpolation of several types.

Within DAMIEN MpCCI will be ported to the PACX-MPI communication library. PACX-MPI is an MPI implementation that allows to couple distributed resources and still fully exploit their network capacity. Thereby the library will be able to couple codes running on distributed computer systems.

The Application

EADS CCR provides a multi-physics coupling application significant for their activities. The application handles the coupling of mechanic and acoustic software for simulation vibrations of a Airbus A320 cockpit (Fig. 3). Due to a high number of entities in the company, often remotely located, the whole industry group faces problems of interoperability between the teams, particularly problems of data exchange between applications.

The Partners

In DAMIEN five partners from France, Spain and Germany are working together: HLRS/Uni Stuttgart, EADS CCR, Cepba, Crihan, and Pallas.
The DAMIEN project is funded by the European Commission.

Contact:
Dr. Ullrich Becker-Lemgau
Pallas GmbH
Hermülheimer Straße 10
50321 Brühl, Germany
Tel: +49-2232-1896-17
ullrich.becker-lemgau@pallas.com

4.2 AcrossGrid Paper by Badia et al.

Performance Prediction in a Grid environment

Rosa M. Badia[2], Francesc Escalé[2], Edgar Gabriel[1,3], Judit Gimenez[2], Rainer Keller[1], Jesús Labarta[2], Matthias S. Müller[1]

[1] High Performance Computing Center Stuttgart, Germany
[2] European Center for Parallelism of Barcelona (CEPBA), UPC, Spain
[3] Innovative Computing Laboratories, Univ. of Tennessee, USA

{keller, gabriel, mueller}@hlrs.de {rosab, fescale, jesus, judit}@cepba.upc.es

Abstract. Knowing the performance of an application in a Grid environment is an important issue in application development and for scheduling decisions. In this paper we describe the analysis and optimisation of a computation- and communication-intensive application from the field of bioinformatics, which was demonstrated at the HPC-Challenge of Supercomputing 2002 at Baltimore. This application has been adapted to be run on an heterogeneous computational Grid by means of PACX-MPI. The analysis and optimisation is based on trace driven tools, mainly Dimemas and Vampir. All these methodologies and tools are being extended in the frame of the DAMIEN IST project.

1 Introduction

The efficient execution of a complex scientific application in a distributed, heterogeneous environment where resources are shared with others is a challenge for the developer and user. The DAMIEN [2] project (Distributed Applications and Middleware for Industrial Use of European Networks) aims to produce a tool-chain to support developers and scientific users of Grid-applications. To achieve this goal, new tools are being created and widely accepted tools extended.

For the HPC-Challenge at SC'02, we demonstrated a computation- and communication-intensive application of the area of bioinformatics, RNAfold, running on a computational Grid, successfully employing the DAMIEN tool-chain. This computational Grid consisted of 22 high-performance computers at the sites of our partners, constituting a distributed, heterogeneous metacomputer.

In this paper we focus on two of the DAMIEN-tools, namely PACX-MPI, which will be introduced in the next section, and Dimemas, described in section 3. Section 4 then describes experiments done with these two tools in conjunction with the application RNAfold.

2 PACX-MPI

The middleware PACX-MPI [3] is an optimized MPI implementation and enables MPI-conforming applications to be run on a heterogeneous computational Grid without requiring the programmer to change the source code.

The hosts making up the computational Grid are coupled through an interconnecting network, e. g. the internet. Communication between MPI-processes within the host is done with the optimized vendor MPI, while communication between MPI-processes on different hosts is carried out over the interconnecting network. For this communication to run efficiently, two MPI-processes on each host, the so called daemons, are employed. Any communication between two remote MPI-processes is done via these daemons. While the internal connections provide very high bandwidth in the range of several hundred MB per second with low latencies, the external connection between hosts is of much lower quality. The main interest within the DAMIEN project is therefore to speed up these external connections.

3 Performance prediction: Dimemas

Dimemas is a performance prediction simulator for message passing applications. Dimemas reconstructs the time behavior of a parallel application on a target architecture. The inputs to Dimemas are: a tracefile of an execution of the application on a source machine that captures the CPU bursts versus the communication pattern information; and a configuration file that contains a set of parameters that model the target architecture. Complemented with a visualization tool [5, 6], it allows the user to gain insight into the application behavior.

The initial Dimemas architecture model considered networks of Shared Memory Processors (SMP). In this architecture, each node is composed of a set of processors and local memory. The interconnection of these nodes is modeled with input/output links and a set of buses.

In the IST project DAMIEN [2], Dimemas is being extended to predict for distributed or Grid architectures. Therefore, the target architecture is now a set of machines, each of them can be a network of SMPs. Different machines are connected through an external Wide Area Network (WAN) or through dedicated connections. Dedicated connections are considered to have a certain QoS, with guaranteed latency and bandwidth, while the WAN latencies and bandwidth can be affected by the traffic present on the network. Both situations may be modeled with the extended Dimemas.

Dimemas has different models for point to point communications and collective operations, with user-configurable parameters to fit different communication library implementations and networks characteristics.

4 Experiments

Every organism's code of live, the genes, are stored in deoxyribonucleic acid (DNA), which is available in every cell of the organism. Still the ribonucleic acid (RNA) plays an even more important role in gene expression and protein synthesis.

The application RNAfold calculates the secondary structure of RNA. It was developed as part of the ViennaRNA package at the University of Vienna [4]

and MPI-parallelized. For iGrid2002 and the HPC-Challenge at SC2002, this application was enhanced by Sandia National Labs and HLRS, namely the energy parameters were updated, the application was integrated into the framework of a virtual reality environment based on Covise [1] and the communication pattern was improved in several consecutive steps to make it more efficient. Still it's computationally expensive in the order of $O(n^3)$, and also communication intensive: $O(n^2)$, n being the number of bases in the sequence.

4.1 Experiments within CEPBA-UPC and between both sites

A first set of experiments was performed at the CEPBA on a heterogeneous metacomputer composed of two machines: an SGI Origin O2000 server with 64 processors (Karnak) and an IBM RS-6000 SP with 8 by 16 Nighthawk Power3 processors (Kadesh). The network measurement between these two machines were more accurate, since the conditions of the network were more predictable.

Dimemas tracefiles of local executions on the SP3 were extracted with mpidtrace, the Dimemas tracefile generator. These tracefiles were then used to feed Dimemas simulator with different input configurations. To estimate the characteristics of the connection between the two machines the ping program was used.

Afterwords, measurements of distributed executions between these two machines were performed to validate Dimemas predictions. The first part of the table 1 shows some of the results obtained in this environment. The prediction for these cases was obtained by setting the network bandwidth to 0.5 MB/s and the flight time (network latency) to 4.9 ms.

For the prediction of Dimemas on links with high-latency, tests were performed on a distributed heterogeneous metacomputer consisting of a 512 processor Cray T3e/900 (Hwwt3e) installed at HLRS and the IBM-SP3. The second part of table summarizes the results of the tests performed. For this case, some of the measurements were done more than one time (columns Measure 1 and Measure 2) showing how influent the network state can be on the achieved performance.

The different columns for prediction are obtained by setting the network bandwidth and flight time parameters to different values in order to compare runtime with changed network situation: for column *Predict 1*, the bandwidth was set to 70KB/s and the flight time to 10ms; for column *Predict 2*, to 100KB/s and 10ms and for *Predict 3* to 200KB/s and 1ms respectively.

5 Conclusion

The use of performance analysis and prediction tools for parallel codes is a must on computational Grids due to the complexity of the environment and the codes itself. The DAMIEN project aims at the development of tools for message passing applications on heterogeneous computational Grids. This abstract presents some of the work performed in this project.

Benchmark	*# processors*		*Time*				
	Kadesh	Karnak	Measurement		Prediction		
Test_5000.seq	4	4	245s		289s		
Test_5000.seq	14	14	107s		116s		
Yeast DNA seq	4	4	338s		409s		
	Hwwt3e	Karnak	Measurement		Predict 1	Predict 2	Predict 3
Test_1000.seq	4	4	79s		58s	42s	20s
Test_2000.seq	4	4	168s		226s	160s	80s
Test_3000.seq	4	4	377s		503s	356s	180s
Test_4000.seq	4	4	1058s		890s	629s	320s
Test_5000.seq	4	4	1409s		1388s	981s	502s
	Hwwt3e	Kadesh	Measure 1	Measure 2	Predict 1	Predict 2	Predict 3
Test_1000.seq	6	6	29s	49s	57s	41s	20s
Test_2000.seq	6	6	85s	204s	223s	158s	78s
Test_3000.seq	6	6	185s	364s	499s	351s	175s
Test_4000.seq	6	6	353s	–	883s	621s	310s
Test_5000.seq	6	6	910s	–	1377s	968s	484s
Test_6000.seq	6	6	1597s	–	1980s	1391s	697s
Test_1000.seq	14	14	30s	47s	55s	39s	19s
Test_2000.seq	14	14	125s	–	220s	154s	78s
Test_3000.seq	14	14	240s	–	493s	346s	173s
Test_4000.seq	14	14	424s	–	876s	614s	307s
Test_5000.seq	14	14	705s	1067s	1377s	968s	480s
Test_6000.seq	14	14	651s	–	1968s	1380s	691s

Table 1. Results obtained with a distributed metacomputer

Acknowledgments

This work was supported by the European Commission under contract number DAMIEN IST-2000-25406.

References

1. Covise. Internet, October 2002. http://www.hlrs.de/organization/vis/covise.
2. DAMIEN. Internet, October 2002. http://www.hlrs.de/organization/pds/projects/damien.
3. E. Gabriel, M. Resch, T. Beisel, and R. Keller. Distributed Computing in a Heterogeneous Computing Environment. In *PVM/MPI*, pages 180–187, 1998.
4. I. L. Hofacker, W. Fontana, L. S. Bonhoeffer, M. Tacker, and P. Schuster. Vienna RNA Package. Internet, October 2002. http://www.tbi.univie.ac.at/~ivo/RNA.
5. Paraver. Internet, October 2002. http://www.http://www.cepba.upc.es/paraver.
6. Vampir. Internet, October 2002. http://www.pallas.com/e/products/vampir/index.htm.

5. Distributed Application Management in Heterogeneous Grids

Peggy Lindner, Natalia Currle-Linde, Michael M. Resch
High Performance Computing Center Stuttgart (HLRS),
Stuttgart, Germany

Edgar Gabriel
High Performance Computing Center Stuttgart (HLRS)
Stuttgart, Germany
and
Innovative Computing Laboratory, Computer Science Department
University of Tennessee, Knoxville, TN, USA

Abstract

Distributing an application on several machines is one of the key aspects of Grid-computing. In the last few years several groups have developed solutions for the occurring communication problems. However, users are still left on their own when it comes to the handling of a Grid-computer, as soon as they are facing a mix of several Grid software environments on target machines. This paper presents a tool, which eases the handling of a Grid-computer, by hiding some of the complexity of heterogeneous Grid-environments from the user.

5.1 Introduction

The Grid is generally seen as a concept for 'coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organization' [5] . The original idea came from scientists, who were mainly interested in the scientific solution of their problem. Their jobs can be executed on any machine, respectively on any set of machines, to which they have access. Like in Metacomputing, which was a popular concept in the Mid 90s [4] , the distribution of jobs onto several machines is an important part of the Grid concept. The need for doing so mainly comes from applications, which have a high demand on computational resources [10] [2] or for increasing throughput on the machines and reducing turnaround times.
Distributing an MPI job onto several machines imposes however many problems to the user and the application. Most of the problems stem from the fact, that the Grid is heterogeneous from several aspects. This starts with different data representations on different machines, which requires data conversion either in the MPI-library or in the application. Various processor speeds and differences in the available memory require the application to have a smart initial load distribution and/or dynamic load balancing. The differences on the communication characteristics for communication between processes located on potentially different sites require distinct programming techniques for hiding the wide area latency and dealing with the low bandwidth. The latter aspects have been handled in the last couple of years by several projects, which focused on the development of MPI libraries for Grid environments, like MPICH-G2 [12] , STAMPI[11] , MPI_Connect [6] and PACX-MPI [9] .
Another aspect of heterogeneity in the Grid has been up to now not really handled. The user has to deal also with several mechanisms of how to access a machine and how to start a job on it. While some computing centers run ssh [17] , rsh or similar protocols, Grid software environments like Globus [5] , Legion [16] or UNICORE [1] are getting more common. There is a high probability that the user has to deal with a mix of these access mechanisms when distributing an application onto several hosts. In the following this is the aspect of heterogeneity which we mean, if we refer to the Grid as being heterogeneous.
Beside troubling the user, this aspect of the Grid affects also the startup procedure of MPI libraries. Library developers tend to integrate their communication library and the startup procedure into a specific Grid software environment. For example, MPICH-G2 is part of the the Globus software toolkit, and uses the Globus mechanisms for authentication, authorization and communication. However, this implies, that it is hard to execute an MPICH-G2 job on a machine, where Globus is not installed.
The default startup procedure of PACX-MPI requires just a login on each machine. In addition, further methods for starting PACX-MPI jobs are provided, which interface directly with several Grid computing environments. Together with the Globus group, an interface has been defined and implemented, which allows a PACX-MPI job to use all Globus features mentioned in the previous section [22] . Furthermore, a plug-in for the UNICORE Grid-computing environment for PACX-MPI is currently under development.
However, none of these solutions support the user, if he has to deal with a heterogeneous Grid-computing

environment. If Globus is installed on the first machine, UNICORE on the second one, and ssh is the only way to access the third machine, the user currently cannot use a unique startup procedure with any of the libraries mentioned above. He might run a PACX-MPI job on all machines, but in this case, he is not taking advantage of any of the features of the Grid software environments.
This situation is not satisfactory. Our experience in the last years doing distributed simulations all over the world (see e.g. [13]) showed furthermore, that starting with a certain number of machines, it is not easy to handle a terminal for each machine. Especially, PACX-MPI requires some configuration files to be available on all hosts. When modifying the configuration of the metacomputer, the user has to update the configuration files of PACX-MPI on all hosts. While this is not a problem for two or three machines, it is an error-prone procedure for more.
To solve the problems described, a tool is currently under development in the frame of the DAMIEN project [19] sponsored by the European Commission called the 'Configuration Manager', and it is designed to handle distributed MPI-applications in heterogeneous Grid-environments. The primary goal of the Configuration Manager is to ease the handling of resources and applications for PACX-MPI jobs.
The structure of this paper is as follows: in the next section we describe PACX-MPI and the required configuration files for the library. Then, we present the goals, concept and the implementation of the Configuration Manager. The following section gives some usage examples of the Configuration Manager. Finally, we summarize the paper and give a brief overview about the ongoing work in this area.

5.2 PACX-MPI

This section describes the communication library PACX-MPI [9] . The focus is on the startup procedure and the configuration files of PACX-MPI, since these aspects are necessary for understanding the concept of the Configuration Managers.
PACX-MPI is the implementation of the message passing standard MPI that is optimized for Meta- and Grid-computing environments. Key aspect of these environments is the difference in the characteristics of the communication layers regarding latency and bandwidth. If communication occurs between two processes within the same machine, the latency is usually in the range of microseconds, and the bandwidth in the range of tens to hundreds of MBytes/s. However, if the communication partners are distributed on different machines, latencies tend to be in the range of tens of milliseconds, and bandwidth drops to some MBits/s. PACX-MPI takes these characteristics into account by using two independent layers for handling communication between processes on the same machine *(internal communication)* and for handling communication between processes on different machines *(external communication)*. Internal communications are handled by using the vendor MPI-library on each machine, since this a highly efficient and portable communication method. For external communications two communication daemons have been introduced, which translate the MPI messages into TCP/IP messages or vice versa. Furthermore, PACX-MPI has optimized the high level functions of the MPI-standard, e.g. the collective operations [7] or the derived datatypes [8] , to these environments. For more details, see e.g. [9] .
When starting an MPI-application on several hosts simultaneously, PACX-MPI reads two configuration files, which answer the following questions:

- Which machines are participating in the current run ?
- Which is the rank of this machine in the given configuration ?
- Which network protocol is a machine using for communicating to each other host ?

The first question is important for determining which machines have to be contacted at all. The second question is dealing with the problem of how to assign a global, unique rank to each process in the global configuration. This problem is solved in PACX-MPI by using for each participating machine a unique host-rank and the local MPI-rank of each process. Finally, the third question clarifies the communication protocol between each pair of hosts and the attributes of the protocol. This information consists of a list of protocols supported by PACX-MPI (currently TCP/IP, a native ATM protocol and a secure protocol based on SSL), and some protocol dependent attributes, e.g. the port number under which a machine can be reached for the TCP/IP protocol.
This information is split into two configuration files: the hostfile , which contains the information regarding the participating machines, and the netfile , which deals with the network parameters of the current run. While the hostfile is mandatory for each run, the second configuration file is optional. In case the library does not find this file, it assumes, that all machines are using the default networking protocol (TCP/IP) with default port-numbers.
In the following, we would like to describe the configuration files using a simple example of three machines, called host1, host2 and host3, with 16 processors on the first two machines, and 24 processors on the last one.

While the first two machines are within the same institution and protected by a firewall, the third machine is located at a different site. Therefore, one can use standard TCP/IP between host1 and host2, while a secure protocol is required to the third machine. This configuration is also shown in figure[1].
In the most simple case, the according hostfile looks like the following:

```
host1 16
host2 16
host3 24
```

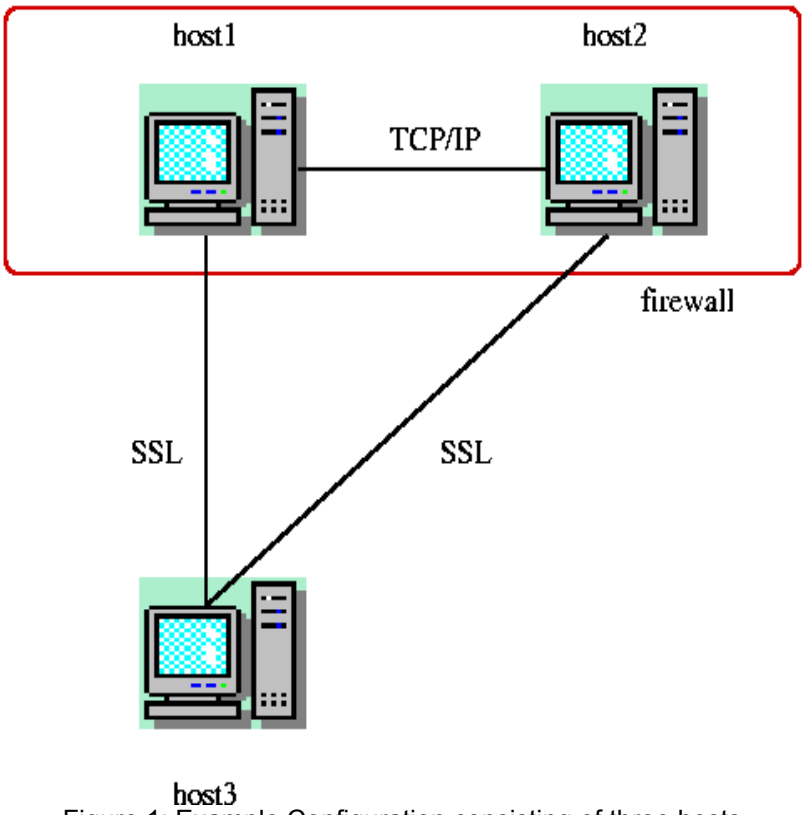

Figure 1: Example Configuration consisting of three hosts

The hostfile consists of a list of the machines and the number of processes used for the application on each machine. The format is straight forward and based on the well-known machinefile of MPICH [15] . This information however is not enough for cluster-like machines (e.g. IBM RS6000SP, PC-Clusters), where each processor has its own IP-address, and therefore the machine can not be identified by a single name or IP-address. In this case, the format of the hostfile and the whole startup procedure is different: it contains the keyword Server, the name or the IP-address of the Startup-Server and the rank of the machine in the global configuration.

```
Server <Startup-Server > <host-rank>
```

The Startup-Server is an independent process. Its purpose is to collect relevant startup-information, e.g. the IP-addresses of all PACX-MPI daemons, and redistribute this in a second step to all hosts.
For the example described above the second configuration file, the netfile looks like the following:

```
BEGIN 1;
```

```
HOST=3, PROTOCOL=SSL;
END;
BEGIN host2;
HOST=3, PROTOCOL=SSL;
END;
BEGIN 3;
HOST=1, PROTOCOL=SSL;
HOST=2, PROTOCOL=SSL;
END;
```

The file contains one section for each participating machine. A section starts with a BEGIN statement and either the rank of the machine in the global configuration or the name of the host and ends with an END statement. Inside each section, the network connections of this host to all other machines are described. To simplify the usage of this file, just those connections between hosts have to be described, which are not using default parameters. Therefore, the connection between host1 and host2 need not be described.

5.3 Goals, Concept and Implementation of the Configuration Manager

To automatize the handling of resources for a PACX-MPI job, the Configuration Manager has to provide the following basic functionality:

- Create the configuration files for PACX-MPI: Based on a GUI, the user describes the required resources for the current run.
- Distribute the configuration files onto all machines: Since the configuration files of PACX-MPI have to be available on all machines, the Configuration Manager has to transfer the most recent version of these files to all machines.
- Start the application on all hosts: Among the information, which the user has to provide when specifying a machine in the Configuration Manager, is a command which will be used to start the application on the given host.

As described in the introduction, the heterogeneity of the Grid will not only consist of different types of machines with different data representations, but also the method how the user can access the machines. Grid-software like Globus, Legion and UNICORE are meanwhile well established. On the other hand, some machines will probably continue to be accessible just by ssh, rsh or similar methods. Therefore, another goal of the Configuration Manager is to hide the complexity of how to access each individual machine for the every-day use.
To hide the complexity of heterogeneous Grids from the user, the Configuration Manager requires the following functionality for each access-mechanism:

- Authentication and Authorization: This can be either done by a combination of username and password for ssh, or a certificate for Globus and UNICORE.
- File-transfer: This is required for sending the current versions of the PACX-MPI hostfile and netfile to all hosts. Furthermore, file transfer is necessary for executable staging or transferring input files to the according machines respectively moving output files back onto local machines.
- Job submission: This is used to start each part of the distributed job on all machines.

The job submission again can be detailed to the following items:

- Name and path of the application: This has to be handled separately for each machine, since the requirement of having the same name and path for the application on all systems can often not be fulfilled (e.g. the name is <myappl>.<architecture> .
- Number of processors: Again this is a machine and run-specific information and has to be stored for each host separately.

- Generic MPI command: Specifying the command how to start an MPI application on a certain machine. In fact, the user often has to face the problem, that several MPI libraries are installed on a machine with different MPI start commands (e.g. mpirun, mpiexec, poe).

In the concept of the Configuration Manager foreseen is the possibility to store the settings for a certain machine. Typically, an end-user will have a limited number of machines which he uses for his distributed runs. This feature allows him also to store the same machine using different MPI implementation (e.g. <machinename>.<vendormpi> and <machinename>.<mpich1.2.3>).

5.3.1 Implementation Issues

The Configuration Manager is implemented in Java, following Model-View-Controller (MVC) principles. This concept separates input, processing and output in graphically oriented software. Therefore, the programmer using the MVC principles has to deal with three different components:

- The model, which represents the application object
- The view, which realizes the representation of the model on the screen and
- The controller, which defines the reaction on input of the user-interfaces.

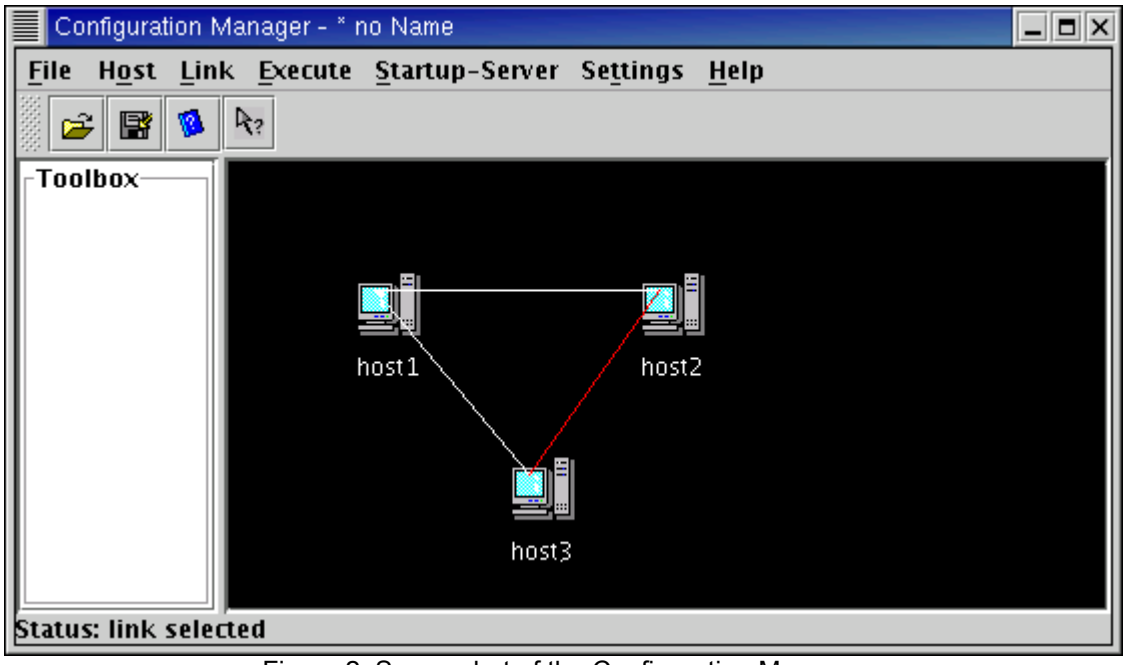

Figure 2: Screenshot of the Configuration Manager

The current implementation of the Configuration Manager supports two different methods to save and open a configuration. In case the same user wants to reuse the configuration, it is convenient for him not to enter the passwords or to specify the certificates for each machine every time. Therefore, the Configuration Manager encrypts all passwords and certificates when saving the configuration. Thus, the user has to enter a password when opening a configuration. The encryption functionality is handled using the Java Cryptography Extensions (JCE), using an implementation of Bouncy Castle [21] .
However, it is sometimes desirable to pass a configuration to another person. In this case passwords and certificates should not be included in the configuration. The Configuration Manager supports this scenario by providing export/import functions.
For the integration of the startup-server mechanism of PACX-MPI for handling cluster-like machines, an implementation of the startup-server in Java has been developed . When using this startup method of PACX-MPI, the GUI automatically changes the content of the hostfile . The user can then launch the startup-server, which will be executed in a separate thread. Thus, the handling of different formats of the PACX-MPI configuration files is nearly seamless for the end-user, improving the usability of PACX-MPI.

5.3.2 An ssh-based Grid-terminal

The first access-mechanism implemented in the Configuration Manager was ssh. Due to its current wide

usage, ssh is still considered to be the most widespread method to access a high-performance computer. The implementation in the Configuration Manager is based on the Java ssh-implementation of Mindbright [18] . Based on this library, a *grid-terminal* (gridterm) has been developed, which allows the user to execute a command simultaneously on one, a subset or all hosts in the current configuration (see fig.[3]). Stdout and stderr are merged for all hosts into a single window. This helps to avoid the situation mentioned in the introduction, where starting with a certain number of machines, the user has problems to manage regular terminals on all hosts. It is therefore a major improvement of a regular terminal for Grid-environments. At the current stage, the Configuration Manager supports password based ssh-access to the different machines.

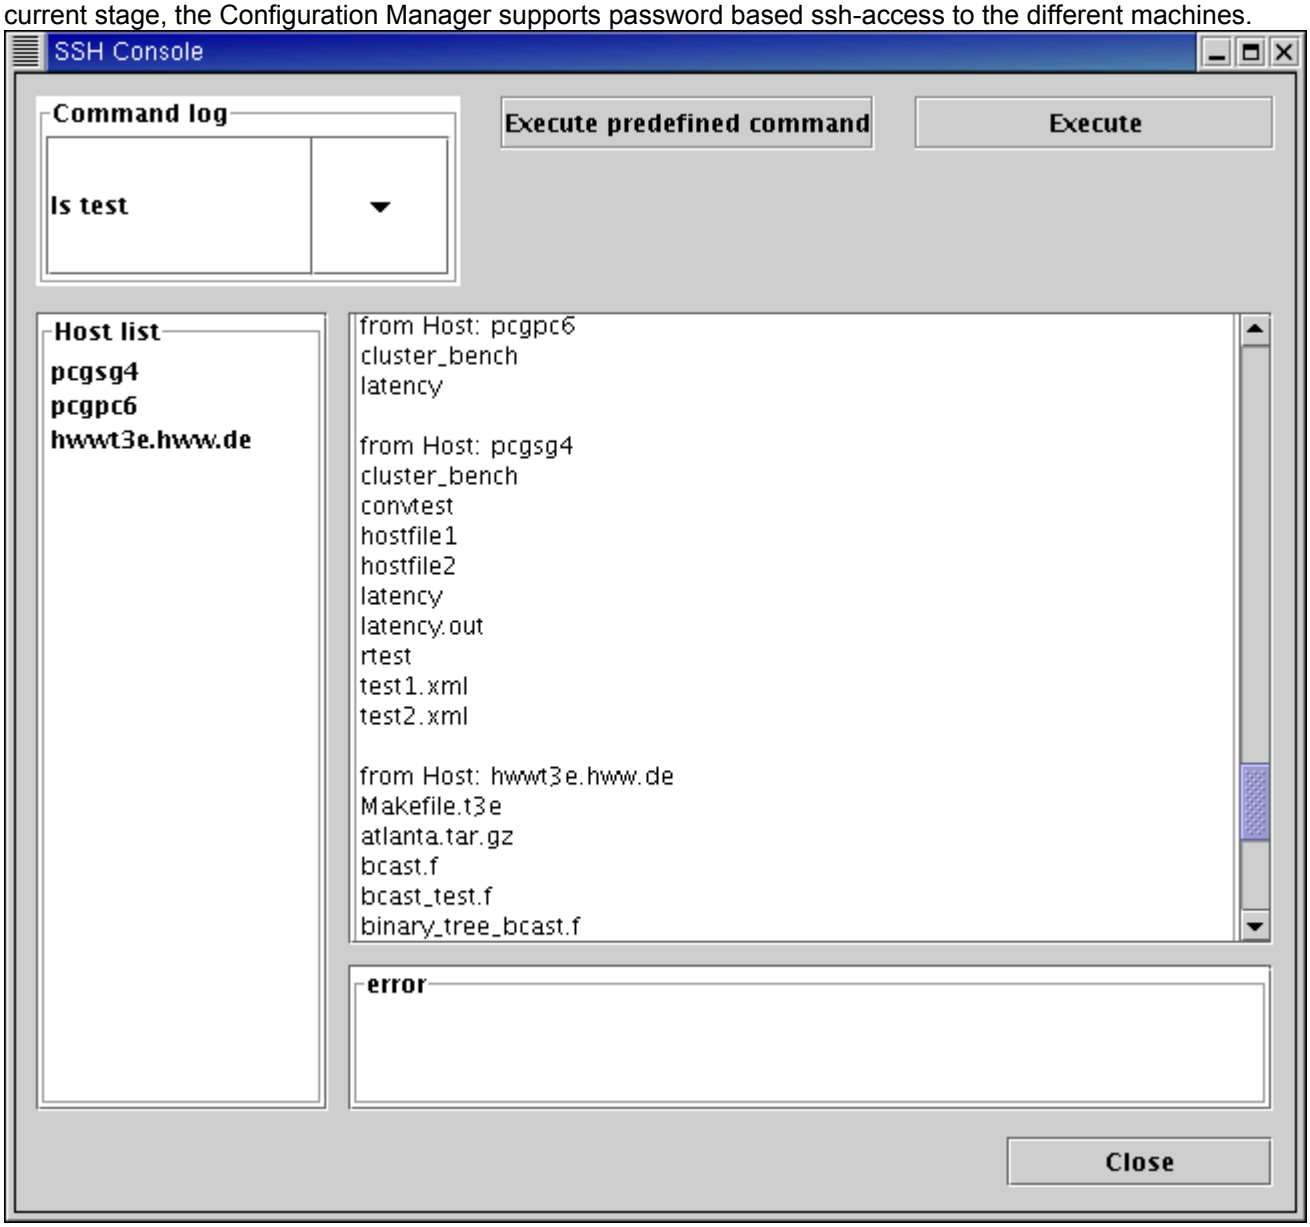

Figure 3: Screenshot of the grid-terminal

5.3.3 Integration with UNICORE

Currently ongoing work is the implementation of UNICORE as a second access mechanism. UNICORE (Uniform Interface to Computer Resources) is a software infrastructure having the goal to provide the user an easy way to submit jobs or chains of jobs to supercomputers. Furthermore, UNICORE supports the idea of specifying a multi-site job, where the first part of an application is executed on the first machine, and the second part is executed on the second machine. In the frame of the UNICORE Plus project, the handling of Metacomputing jobs with PACX-MPI is already supported. To realize this, the plug-in mechanism of UNICORE is used for creating the PACX-MPI configuration files and for launching the jobs. However, as mentioned previously in the introduction, this mechanism requires UNICORE to be installed on all machines,

which are part of the Metacomputer. Therefore, another solution for handling mixed environments is currently being developed for the Configuration Manager.
Like ssh, UNICORE supports file-transfer and the starting of a job on a machine. Two major differences to ssh are however important when integrating UNICORE into the Configuration Manager:

- UNICORE does currently not support interactive access to computing resources. While this is ongoing work in the frame of the Eurogrid project [14] , this does mean, that all jobs (even a simple ls-command) will be executed as batch jobs. Since co-allocation of resources across several sites is still an unsolved problem, this does also impose currently, that it is up to the user to make sure, that all resources are available at the same time.

- The access to UNICORE resources is based on X509 certificates, and not on passwords. Since this is also part of the ssh2 protocol, the Configuration Manager is currently being extended to support the management of certificates as-well.

5.4 Example of Usage scenarios

The Configuration Manager has been used in several demonstrations and scenarios. It is currently tested under industrial conditions by the European Aeronautic Defense and Space Company (EADS) in the frame of the DAMIEN project. In this section, we will give brief descriptions of another scenario presented in the frame of the DAMIEN project, where computational resources in Stuttgart/Germany, Rouen/France and Barcelona/Spain were connected for a distributed simulation between these sites. The machines used were an SGI Origin 2000 in both Rouen and Stuttgart and an IBM RS6000SP in Barcelona.
During the preparation of the demonstration, the network between the sites was continuously modified in order to improve the connection of each site to the new European high-speed backbone network Geant. This time, the Configuration Manager turned out to be extremely useful, since the launching of the application could be done within seconds. Therefore, an immediate response to the networking experts about the quality of the network and the impact of their modifications could be provided.
The application used for the simulation was a flow-solver called URANUS [3] . This code simulates the flow around a space vehicle during the reentry phase. In the final run, a simulation of the MIRKA-vehicle has been successfully demonstrated for a medium size problem.
Of importance for the Configuration Manager is one technical issues of URANUS, its directory structure. During compile-time, a directory will be created for each architecture. Thus, the executables are not located in the same directory for the IBM RS6000SP and for the SGI Origin2000. While the whole run could also be done without the Configuration Manager, its usage eased the management of the application and provided an easy to use tool for Grid- and Metacomputing.

5.5 Conclusions and future plans

In this paper we presented a tool which helps to ease the usage of distributed resources in Grid-environments. The Configuration Manager creates the configuration files for PACX-MPI, transfers them to the participating machines and starts the application on all hosts simultaneously. Furthermore, it hides parts of the complexity of the Grid from the user by supporting several mechanisms of how to reach different machines. Currently implemented are ssh and (partially) UNICORE. This tools has proved its usability in several demonstrations and projects, and is currently being tested in an industrial environment.
A further aspect of the Configuration Manager has proved its usefulness beyond the submission of PACX-MPI jobs. The grid-terminal (gridterm) is an extension of a regular terminal, which enables simultaneous execution of commands and applications on several hosts in a Grid environment.
Further plans for the tool involve the support for other Grid software environments such as Globus or Legion. Together with the already implemented mechanisms, we believe we will then cover the most relevant software environments for Grid-computing.
In the frame of the DAMIEN project, PACX-MPI is extended to handle Quality-of-Service parameters. The Configuration Manager will support the specification of these parameters as-well.
An omnipresent question currently is the convergence of all Grid-tools toward the Open Grid Software Architecture (OGSA) [23] , since most of the Grid software environments, like e.g. Globus or UNICORE, have committed to convert toward this new architecture. In this concept Grid-computing is constructed out of several Grid Services with defined interfaces and semantics. These services are accessible by several protocols e.g. the Simple Object Access Protocol (SOAP) [24] . The development and implementation of a tool like the Configuration Manager could potentially be eased by a common architecture, since all operations should be formulated as requests to an abstract system, independent of the implementation underneath.

However, such a common architecture would not really affect the Configuration Manager for two reasons: first, none of the systems are currently exceeding the status of pure demonstrations, and second, the argument, that some systems will remain just accessible through protocols like ssh, remains still valid.

5.6 Acknowledgments

This work was supported by the European Commission in the DAMIEN project under contract number IST-2000-25406.

5.7 References

[1] Jim Almond, Dave Snelling, 'UNICORE: Secure and Uniform Access to Distributed Resources via the World Wide Web, A White Paper, October 1998.

[2] Gabrielle Allen, Thomas Dramlitsch, Ian Foster, Nicholas T. Karonis, Matei Ripeanu, Edward Seidel, Brian Toonen, 'Supporting Efficient Execution in Heterogeneous Distributed Computing Environments with Cactus and Globus', Supercomputing 2001, Denver 2001

[3] Thomas B. Boenisch, Roland Ruehle, ' Efficient Flow Simulation with Structured Multiblock Meshes on Current Supercomputers', ERCOFTAC Bulletin No. 50: Parallel Computing in CFD, 2001.

[4] Charly Catlett and Larry Smarr, 'Metacomputing', Communications of the ACM, 35(6):44-52, 1992.

[5] Ian Foster, Carl Kesselmann, Steven Tuecke, 'Anatomy of the Grid: Enabling Scalable Virtual Organizations', International Journal of Supercomputing Applications, 15(3), 2001.

[6]Graham E. Fagg, Kevin. S. London, and Jack J. Dongarra, 'MPI_Connect Managing Heterogeneous {MPI} Applications Interoperation and Process Control', In V.~Alexandrov and J.~Dongarra, editors, Recent advances in Parallel Virtual Machine and Message Passing Interface, volume 1497 of Lecture Notes in Computer Science, pages 93--96. Springer, 1998, 5th European PVM/MPI Users' Group Meeting.

[7] Edgar Gabriel, Michael M. Resch, and Roland Ruehle, 'Implementing MPI with Optimized Algorithms for Metacomputing', In Y.~S.~D. Anthony~Skjellum, Purushotham V.~Bangalore, editor, Proceedings of the Third MPI Developer's and User's Conference, 1999.

[8] Edgar Gabriel, Michael M. Resch, and Roland Ruehle, 'Implementing and Benchmarking Derived Datatypes in Metacomputing',In B. Hertzberger, A. Hoekstra, R. Williams (Eds).High Performance Computing and Networking, pp. 493-502, Springer, 2001.

[9] Edgar Gabriel, Michael M. Resch, Thomas Beisel, Rainer Keller, 'Distributed Computing in a Heterogeneous Computing Environment', Vassil Alexandrov, Jack Dongarra (Eds.) Recent Advances in Parallel Virtual Machine and Message Passing Interface, pp. 180-188, Springer, 1998.

[10] Edgar Gabriel, Marc Lange and Roland Ruehle, 'Direct Numerical Simulation of Turbulent Reactive Flows in a Metacomputing Environment', Proceedings of the 2001 ICPP Workshops, pp. 237-244, 2001.

[11] Toshiyuki Imamura, Yuichi Tsujita, Hiroshi Koide, and Hiroshi Takemiya, 'An Architecture of Stampi: MPI Library on a Cluster of Parallel Computers', J.~Dongarra, P.~Kacsuk, and N.~Podhorszki, editors, Recent Advances in Parallel Virtual Machine and Message Passing Interface, number 1908 in Springer Lecture Notes

in Computer Science, pages 200--207, Sept. 2000. 7th European PVM/MPI Users' Group Meeting.

[12] Nicholas T. Karonis and Brian Toonen, MPICH-G2 web page http://www.globus.org/mpi

[13] Steven M. Pickles, John M. Brooke, Fumie C. Costen, Edgar Gabriel, Matthias Mueller, Michael M. Resch, and S. M. Ord, 'Metacomputing Across Intercontinental Networks', Future Generation Computer Systems (17), pp. 911-918, Elsevier Science.

[14]EUROGRID web page: http://www.eurogrid.org

[15] William Gropp and Ewing Lusk, 'User's Guide for mpich, a Portable Implementation of MPI'.

[16] Andrew S. Grimshaw, Wm. A. Wulf, and the Legion team, 'The Legion vision of a worldwide virtual computer', Communications of the ACM, 40(1):39-45, 1997.

[17] OpenSSH web page http://www.openssh.org

[18] Mindterm Secure Shell web page: http://www.mindbright.se

[19] Matthias Mueller, Edgar Gabriel and Michael Resch, 'A Software Development Environment for Grid-Computing', Concurrency and Computation - Practice and Experience, special issue on Grid computing environments, to be published.

[20] DAMIEN - Distributed Application and Middleware for Industrial Use of European Networks. Web page to be found at http://www.hlrs.de/organization/pds/projects/damien

[21] Bouncy Castle Crypto package. Web page to be found at http://www.bouncycastle.org/documentation

[22] B. Toonen, D. Ashton, E. Lusk, I. Foster, W. Gropp, E. Gabriel, R. Butler, N. Karonis, 'Interfacing Parallel Jobs to Process Managers', Proc. of the Tenth IEEE International Symposium on High Performance Distributed Computing, pp. 431-432., August 7-9, 2001, San Francisco, USA.

[23] Ian Foster, Carls Kesselman, Jeffrey M. Nick, Steven Tuecke, 'The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration', June 22, 2002. Document available at http://www.globus.org/research/papers.html

[24] SOAP (http://www.w3.org/2000/xp/Group/)

Copyright Notices

6. Photos

6.1 IST2002, Copenhagen, Denmark


GRID Technology at Work
The technology presented at IST2002 was partly funded by the IST projects
EUROGRID (IST-1999-20247)
EUROGRID
INTONE (IST-1999-20252)
intone
DAMIEN (IST-2000-25406)
DAMIEN


GridLab
DataTAG

6.2 Supercomputing 2002, Baltimore, USA

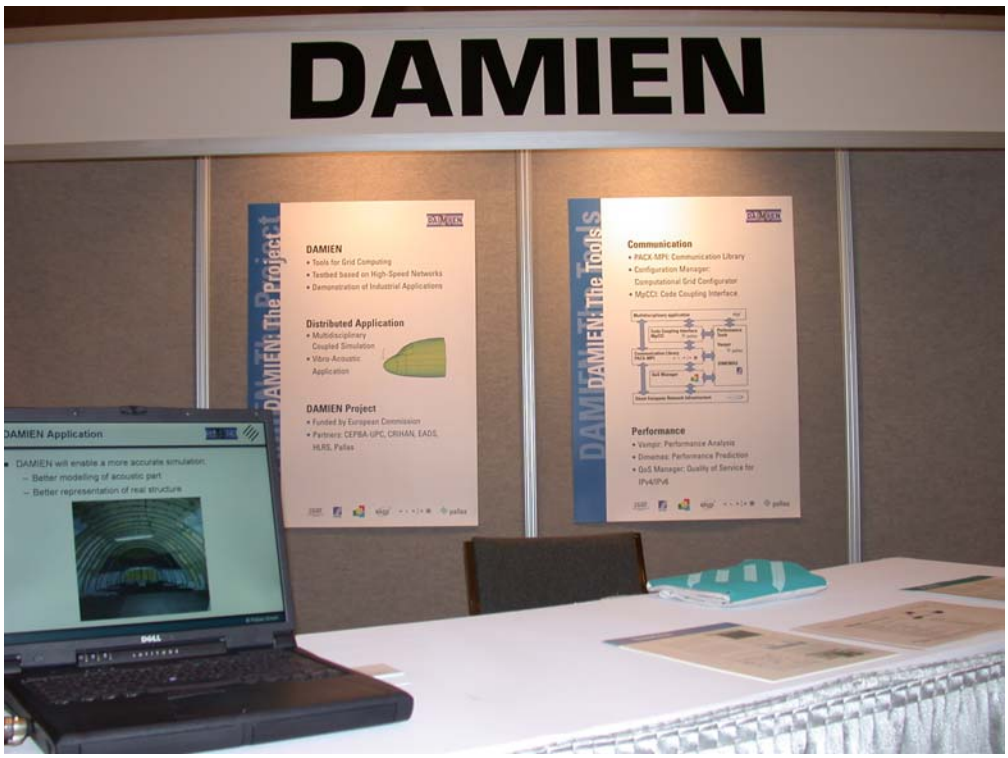


DAMIEN
vampir 3.0