



10th European PVM/MPI Users' Group Meeting
Sep 29 – Oct 2, 2003
San Servolo Island, Venice, Italy

Conference Program

Platinum Sponsors



Gold Sponsor



Silver Sponsors



Monday, September 29th, 2003

14.00	14.30	Registration	
		Tutorial A (Part 1) (Room A)	Tutorial B (Part 1) (Room B)
14.30	16.30	High-Level Programming in MPI <i>William Gropp, Ewing Lusk</i>	Programming Environments for Grids and Distributed Computing Systems <i>Vaidy Sunderam</i>
16.30	16.45	Coffee break	
		Tutorial A (Part 2) (Room A)	Tutorial B (Part 2) (Room B)
16.45	18.45	High-Level Programming in MPI <i>William Gropp, Ewing Lusk</i>	Programming Environments for Grids and Distributed Computing Systems <i>Vaidy Sunderam</i>

Tuesday, September 30th, 2003

9.10	9.30	Registration	
9.30	9.50	Opening Session (Room A)	
9.50	10.40	Invited talk (Room A - Chair: Jack Dongarra) Messaging Systems: Parallel Computing the Internet and the Grid <i>Geoffrey Fox</i>	
10.40	11.00	Coffee break	
11.00	11.50	Invited talk (Room A - Chair: Jack Dongarra) Progress towards Petascale Virtual Machines <i>Al Geist</i>	
11.50	13.10	Session 1A <i>Extensions, Improvements and Implementations of PVM/MPI</i> (Room A - Chair: Vaidy Sunderam) Improving Generic Non-Contiguous File Access for MPI-IO <i>J. Worringer, J. Larson Träff, H. Ritzdorf</i> A Component Architecture for LAM/MPI <i>J. M. Squyres, A. Lumsdaine</i> Improving the Performance of Collective Operations in MPICH <i>R. Thakur, W. Gropp</i> MPI on BlueGene/L: Designing an Efficient General Purpose Messaging Solution for a Large Cellular System <i>G. Almási, C. Archer, J. G. Castaños, M. Gupta, X. Martorell, J. E. Moreira, W. Gropp, et al.</i>	Session 1B <i>Parallel algorithms using message passing</i> (Room B - Chair: Yves Robert) CGMgraph/CGMlib: Implementing and Testing CGM Graph Algorithms on PC Clusters <i>A. Chan, F. Dehne</i> An Online Parallel Algorithm for Remote Visualization of Isosurfaces <i>A. Clematis, D. D'Agostino, V. Gianuzzi</i> Non-strict Evaluation of the FFT Algorithm in Distributed Memory Systems <i>A. Cristóbal-Salas, A. Tchernykh, J.-L. Gaudiot</i> Object-Oriented NeuroSys: Parallel Programs for Simulating Large Networks of Biologically Accurate Neurons <i>P. Pacheco, P. Miller, J. Kim, T. Leese, Y. Zabayaka</i>
13.10	14.30	Lunch	

14.30	16.10	<p>Session 2A <i>Evaluation and Performance Analysis</i> (Room A - Chair: Bernard Tourancheau)</p> <p>Evaluating the performance of MPI-2 dynamic communicators and one-sided communication <i>E. Gabriel, G. E. Fagg, J. Dongarra</i></p> <p>Performance Modeling and Evaluation of Java Message-Passing Primitives on a Cluster <i>G. L. Taboada, J. Touriño, R. Doallo</i></p> <p>Off-line Performance Prediction of Message-Passing Applications on Cluster Systems <i>E. Mancini, M. Rak, R. Torella, U. Villano</i></p> <p>Integrating New Capabilities into NetPIPE <i>D. Turner, A. Oline, X. Chen, T. Benjegerdes</i></p> <p>A Model for Performance Analysis of MPI applications on Terascale Systems <i>S. Chakravarthi, A. Skjellum, C.R. Krishna Kumar, B. Seshadri, Prahalad H. A.</i></p>	<p>Session 2B <i>Parallel algorithms using message passing</i> (Room B - Chair: Dieter Kranzlmüller)</p> <p>A Parallel Approach for the Solution of non Markovian Petri Nets <i>M. Scarpa, S. Distefano, A. Puliafita</i></p> <p>Parallel Algorithms for Computing the Smith Normal Form of Large Matrices <i>G. Jäger</i></p> <p>PageRank Computation using PC Cluster <i>A. Rungsawang, B. Manaskasemsak</i></p> <p>A Scalable Crystallographic FFT <i>J. Seguel, D. Burbano</i></p> <p>Efficient Parallel Implementation of Transitive Closure of Digraphs <i>C. E. R. Alves, E. N. Cáceres, A. A. Castro Jr, S. W. Song, and J. L. Szwarcfiter</i></p>
16.10	16.25	Coffee break	
16.25	17.45	<p>Session 3A <i>Extensions, Improvements and Implementations of PVM/MPI</i> (Room A - Chair: Rolf Hempel)</p> <p>Network Fault Tolerance in LA-MPI <i>R. T. Aulwes, D. J. Daniel, N. N. Desai, R. L. Graham, et al.</i></p> <p>PVMWebCluster: Integration of PVM Clusters using Web Services and CORBA <i>P. Czarnul</i></p> <p>Implementing Fast and Reusable Datatype Processing <i>R. Ross, N. Miller, W. Gropp</i></p>	<p>Session 3B <i>Monitoring and Programming Tools</i> (Room B - Chair: Peter Kacsuk)</p> <p>TOM - Efficient Monitoring Infrastructure for Multithreaded Programs <i>B. Balis, M. Bubak, W. Funika, R. Wismüller, G. Kaplita</i></p> <p>Distributed Configurable Application Monitoring on SMP Clusters <i>K. Furlinger, M. Gerndt</i></p> <p>Architecture of Monitoring System for Distributed Java Applications <i>M. Bubak, W. Funika, M. Smetek, Z. Kilianski, R. Wismüller</i></p>

		Fast and Scalable Barrier using RDMA and Multicast Mechanisms for InfiniBand-Based Clusters <i>S. P. Kini, J. Liu, J. Wu, P. Wyckoff, D. K. Panda</i>	A Communication API for Implementing Irregular Algorithms on SMP Clusters <i>J. Hippold, G. Runger</i>
17.45	19.15	Poster Session: Posters	Special Session: ParSim 03
19.15	21.15	Welcome Reception	

Wednesday, October 1st, 2003

9.20	10.10	<p>Invited talk (Room A - Chair: Salvatore Orlando)</p> <p>Future Developments in MPI <i>William Gropp</i></p>	
10.10	11.00	<p>Invited talk (Room A - Chair: Salvatore Orlando)</p> <p>ASSIST High-performance Programming Environment: Application Experiences and Grid Evolution <i>Marco Vanneschi</i></p>	
11.00	11.20	<p>Coffee break</p>	
11.20	13.00	<p>Session 4A <i>Grid and Heterogeneous Computing</i> (Room A - Chair: Craig Lee)</p> <p>Component-Based System for Grid Application Workflow Composition <i>M. Bubak, K. Górká, T. Gubala, M. Malawski, K. Zajac</i></p> <p>Mapping and load-balancing iterative computations on heterogeneous clusters <i>A. Legrand, H. Renard, Y. Robert, F. Vivien</i></p> <p>Executing and Monitoring PVM Programs in Computational Grids with Jini <i>G. Sipos, P. Kacsuk</i></p> <p>Monitoring Message Passing Applications in the Grid with GRM and R-GMA <i>N. Podhorszki, P. Kacsuk</i></p> <p>Resource Monitoring and Management in Metacomputing Environments <i>T. Wrzosek, D. Kurzyniec, D. Drzewiecki, V. Sunderam</i></p>	<p>Session 4B <i>Extensions, Improvements and Implementations of PVM/MPI</i> (Room B - Chair: Jesper Larson Träff)</p> <p>Using a self-connected Gigabit Ethernet adapter as a memcopy () low-overhead engine for MPI <i>G. Ciaccio</i></p> <p>The network agnostic MPI - Scali MPI Connect <i>L. P. Huse, O. W. Saastad</i></p> <p>Efficient Message-Passing within SMP Systems <i>X. Chen, D. Turner</i></p> <p>PC/MPI: Design and Implementation of a Portable MPI Checkpointer <i>S. Ahn, S. Han</i></p> <p>Evaluation of an Eager Protocol Optimization for MPI <i>R. Brightwell, K. Underwood</i></p>
13.00	14.30	<p>Lunch</p>	
14.30	15.30	<p>Vendor Session (Room A - Chair: Salvatore Orlando)</p>	
15.30	23.00	<p>Excursion & Social Dinner</p>	

Thursday, October 2nd, 2003

9.20	10.10	<p>Invited talk (Room A - Chair: Domenico Laforenza)</p> <p>Integrating Scalable Process Management into Component-Based Systems Software <i>Ewing Lusk</i></p>	
10.10	11.00	<p>Invited talk (Room A - Chair: Domenico Laforenza)</p> <p>Programming High Performance Applications using Components <i>Thierry Priol</i></p>	
11.00	11.20	<p>Coffee break</p>	
11.20	13.00	<p>Session 5A <i>Extensions, Improvements and Implementations of PVM/MPI</i> (Room A - Chair: Umberto Villano)</p> <p>ORNL-RSH package and Windows '03 PVM 3.4 <i>P. Pfeiffer, S. L. Scott, H. Shukla</i></p> <p>MPI for the Clint Gb/s interconnect <i>N. Fugier, M. Herbert, E. Lemoine, B. Tourancheau</i></p> <p>Remote exception handling for PVM processes <i>P. Kaczmarek, H. Krawczyk</i></p> <p>A Comparison of MPICH Allgather Algorithms on Switched Networks <i>G. D. Benson, C.-W. Chu, Q. Huang, S. G. Caglar</i></p> <p>Porting P4 to Digital Signal Processing Platforms <i>J. A. Rico, J. C. Díaz Martín, J. M. Rodríguez García, J. M. Álvarez Llorente, J. L. García Zapata</i></p>	<p>Session 5B <i>Applications in Science and Engineering</i> (Room B - Chair: Casiano Rodriguez-Leon)</p> <p>Hierarchical MPI+OpenMP implementation of parallel PIC applications on clusters of Symmetric MultiProcessors <i>S. Briguglio, B. Di Martino, G. Fogaccia, G. Vlad</i></p> <p>A Parallel Split Operator Method for the Time Dependent Schrödinger equation <i>Jan P. Hansen, T. Matthey, T. Sørsvik</i></p> <p>Parallel Pencil-Beam Redefinition Algorithm <i>P. Alderson, M. Wright, A. Jain, R. Boyd</i></p> <p>Flow Pattern and Heat Transfer Rate in Three-Dimensional Rayleigh-Benard Convection <i>T. Watanabe</i></p> <p>A Parallel Software for the Reconstruction of Dynamic MRI Sequences <i>G. Landi, E. Loli Piccolomini, F. Zama</i></p>

13.00	14.30	Lunch	
		<p>Session 6A <i>Grid and Heterogeneous Computing</i> (Room A - Chair: Yannis Cotronis)</p> <p>Evaluating and Enhancing the Use of the GridFTP Protocol for Efficient Data Transfer on the Grid <i>M. Cannataro, C. Mastroianni, D. Talia, P. Trunfio</i></p> <p>Multiprogramming Level of PVM jobs in a non-dedicated Linux NOW <i>F. Giné, F. Solsona, P. Hernández, M. Hanzich, E. Luque</i></p> <p>Generating an Efficient Dynamics Multicast Tree under Grid Environment <i>T. Vorakosit, P. Uthayopas</i></p> <p>Dynamic Topology Selection for High Performance MPI in the Grid environments <i>K.-L. Park, H.-J. Lee, K.-W. Koh, O.-Y. Kwon, S.-Y. Park, H.-W. Park, S.-D. Kim</i></p> <p>Topology-Aware Communication in Wide-Area Message-Passing <i>C. A. Lee</i></p>	<p>Session 6B <i>Evaluation, Improvements, Performance Analysis</i> (Room B - Chair: Joao Gabriel Silva)</p> <p>Ring algorithms on heterogeneous clusters with PVM: performance analysis and modeling <i>A. Corana</i></p> <p>Complexity Driven Performance Analysis <i>García L., González J.A., González J.C., León C., Rodríguez C., Rodríguez G.</i></p> <p>Usefulness and Usage of SKaMPI-Bench <i>W. Augustin, T. Worsch</i></p> <p>The Performance of Parallel Disk Write Methods for Linux Multiprocessor Nodes <i>G. D. Benson, K. Long, P. S. Pacheco</i></p> <p>Lock-free collective operations <i>A. Supalov</i></p>
14.30	16.10		
16.10	16.25	Coffee Break	
		<p>Session 7A <i>Programming tools</i> (Room A - Chair: Giuseppe Ciaccio)</p> <p>MPI Farm programs on non-dedicated clusters <i>N. Fonseca, J. G. Silva</i></p> <p>An MPI Implementation Supported by Process Migration, Group Process Creation and Load Balancing <i>A. Maloney, A. M. Goscinski, M. Hobbs</i></p> <p>Integrating Multiple Implementations and Structure Exploitation in the Component-based Design of Parallel ODE Solvers <i>J. M. Mantas Ruiz, J. Ortega Lopera, J. Antonio Carrillo</i></p>	<p>Session 7B <i>Algorithms and Applications in Science and Engineering</i> (Room B - Chair: Paolo Palmerini)</p> <p>Advanced Hybrid MPI/OpenMP Parallelization Paradigms for Nested Loop Algorithms onto Clusters of SMPs <i>N. Drosinos, N. Koziris</i></p> <p>The AGEB Algorithm for Solving the Heat Equation in Two Space Dimensions and Its parallelization on A Distributed Parallel Machine <i>N. Alias, S. S. Mohamed, A. R. Abdullah</i></p> <p>Improving Wildland Fire Prediction on MPI Clusters <i>B. Abdalhaq, G. Bianchini, A. Cortés, T. Margalef, E. Luque</i></p>
16.25	17.45		

		Application Composition in Ensemble using Intercommunicators and Process Topologies <i>Y. Cotronis</i>	Building 3D State Spaces of Virtual Environments with a TDS-based Algorithm <i>A. Krenek, I. Peterlík, L. Matyska</i>
17.45	18.00	Closing Session (Room A)	

Tutorials

Tutorial (Sep 29th, 2003 - Room A - h. 14:30-18:45)

High-Level Programming in MPI

William Gropp, Ewing Lusk

*Mathematics and Computer Science Division
Argonne National Laboratory, USA*

MPI is often thought of as a low-level approach, even as a sort of "assembly language", for parallel programming. This is both true and false. While MPI is designed to afford the programmer the ability to control the flow of data at a detailed level for maximum performance, MPI also provides highly expressive operations that support high-level programming. MPI's design has also encouraged the design of parallel libraries, which can provide performance while shielding the user from the details of MPI programming. This tutorial will begin with the basics of MPI, so that attendees need not be familiar with MPI beforehand. We will focus, however, on using MPI as a high-level language by exploiting data type libraries, collective operations, and MPI-I/O. We will cover both MPI-1 and MPI-2 topics in this area, and will introduce some tools for analyzing performance of MPI programs. We will then familiarize attendees with libraries that have been built on MPI, such as ScaLAPACK, PETSc, and parallel NetCDF. We will conclude with a hands-on, "class participation" project. Using whatever laptops attendees can bring to the tutorial together with a wireless network we will supply, we will set up a parallel computer and supply attendees with copies of MPICH-2, our all-new implementation of MPI-1 and MPI-2. If all goes well, we should thus be able to have a quite capable parallel computer available for some interesting demonstrations and experiments. Both Windows and Linux laptops will be accommodated.

Tutorial (Sep 29th, 2003 - Room B - h. 14:30-18:45)

Programming Environments for Grids and Distributed Computing Systems

Vaidy Sunderam

*Department of Mathematics and Computer Science
Emory University, USA*

Platforms for high-performance concurrent computing range from tightly-coupled multiprocessors to clusters, networks of workstations, and large-scale geographically distributed metacomputing systems. Effective application programming for each type of environment can be a challenge, given the wide variability and heterogeneity of the computational resources, communication interconnects, and level of dedicated access. In this tutorial, we will review and compare programming paradigms and associated runtime systems for high-performance concurrent computing in multiple types of distributed system environments. After a brief discussion of language-based approaches, we will analyze the underlying concurrent programming models embodied in MPI and PVM, with a particular focus on their suitability in loosely coupled distributed systems and grids. The emerging paradigm shift in computational grids to a standards-based service-oriented model will then be explored, and methodologies for developing and programming grid services will be discussed. Examples of writing and deploying grid services as well as clients will be included. The tutorial will conclude by summarizing the essential features and characteristics of programming different distributed computing environments and discussing approaches to application development using a combination of programming paradigms.

Invited Talks

Invited Talk (Sep 30th, 2003 - Room A - h. 9:50-10:40)

Messaging Systems: Parallel Computing the Internet and the Grid

Geoffrey Fox

*Indiana University
Community Grids Computing Laboratory
Computer Science, Informatics and Physics
USA*

We contrast the requirements and performance of messaging systems in parallel and distributed systems emphasizing the importance of the five orders of magnitude difference in network hardware latencies in the two cases. We note the importance of messaging in Grid and Web service applications in building the integrated system and the architectural advantages of a message based compared to a connection based approach. We illustrate these points using the NaradaBrokering system and its application to Audio-Video conferencing.

Invited Talk (Sep 30th, 2003 - Room A - h. 11:00-11:50)

Petascale Virtual Machines

AI Geist

*Oak Ridge National Laboratory
USA*

PVM continues to be a popular software package both for creating personal grids and for building adaptable, fault tolerant applications. We will illustrate this by describing a computational biology environment built on top of PVM that is used by researchers around the world. We will then describe or recent progress in building an even more adaptable distributed virtual machine package called Harness. The Harness project includes research on a scalable, self-adapting core called H2O, and research on fault tolerant MPI. The H2O core can be configured to support distributed web computing like SETI@home, parallel virtual machines like PVM, and OGSA compliant grid environments. The Harness software is being designed to scale to petascale virtual machines. We will describe work at Oak Ridge National Lab on developing algorithms for such petascale virtual machines and the development of a simulator to test these algorithms on simulated 100,000 processor systems.

Invited Talk (Oct 1st, 2003 - Room A - h. 9:20-10:10)

Future Developments in MPI

William Gropp

*Mathematics and Computer Science Division
Argonne National Laboratory\
USA*

The Message Passing Interface (MPI) has been very successful at providing a programming model for computers from small PC clusters through the world's fastest computers. MPI has succeeded because it

successfully addresses many of the requirements of an effective parallel programming model, including portability, performance, modularity, and completeness. But much remains to be done with MPI, both in terms of the performance of MPI and in the supporting the use of MPI in applications. This talk will look at three areas: programming models, implementations, and scalability.

The MPI programming model is often described as a supporting "only" basic message passing (point-to-point and collective) and (in MPI-2) simple one-sided communication. Such a description ignores the support in MPI for the creation of effective libraries built using MPI routines. This support has encouraged the development of powerful libraries that, working with MPI, provide a powerful high-level programming environment. This will be illustrated with two examples drawn from computational simulation.

MPI was designed to allow implementations to fully exploit the available hardware. It provides many features that support high performance, including a relaxed memory consistency model. While many MPI implementations take advantage of some of these opportunities, much remains to be done. This talk will describe some of the opportunities for improving the performance of MPI implementations, with particular emphasis on the relaxed memory model and both MPI's one-sided and parallel I/O operations.

Scalability is another goal of the MPI design and many applications have demonstrated scalability to thousands of processors. In the near future, computers with more than 64,000 processors will be built. Barriers to scalability in the definition and the implementation of MPI will be discussed, along with possible future directions for MPI developments. By avoiding a few very low usage routines and with the proper implementation, MPI should scale effectively to the next generation of massively parallel computers.

Invited Talk (Oct 1st, 2003 - Room A - h. 10:10-11:00)

ASSIST High-performance Programming Environment: Application Experiences and Grid Evolution

Marco Vanneschi

*Dipartimento di Informatica
Università di Pisa
Italy*

ASSIST (A Software development System based upon Integrated Skeleton Technology) is a new programming environment oriented to the development of parallel and distributed high-performance applications according to a unified approach. The main goals are: high-level programmability and software productivity for complex multidisciplinary applications, including data-intensive and interactive software; performance portability across different platforms, from homogenous parallel machines and possibly heterogeneous clusters to large-scale enabling platforms and computational Grids; effective reuse of parallel and distributed software; efficient evolution of applications through versions scalable according to the underlying technologies.

The programming model of ASSIST has been defined according to two main issues: i) evolution of structured parallel programming, starting from pros/cons of our past experience with the skeletons model, ii) joining the structured programming approach with the objects/components technology. This talk is organized in two parts. In the first part we discuss the utilization of ASSIST according to a set of application and benchmarking experiences performed during the last year. Such experiences have been done with the ASSIST 1.0 version of compiler and run-time system for Linux clusters; however, the run-time system is based upon ACE portable libraries, thus it is able to support ASSIST programs for heterogeneous clusters too. We show the expressive power and performance measures, and possible weaknesses to be overcome in the next versions of ASSIST.

In the second part we discuss the evolution of ASSIST for large-scale platforms and Grids. The development of Grid applications requires, in general, capabilities and properties beyond those needed in both sequential programming and in "classical" parallel/distributed programming, as it requires the management of computations and environments that are typically dynamic and heterogeneous in their nature and that include resource hierarchies with different features (e.g. memory and network). These issues are investigated in the Italian National Programme in Grid Computing (Grid.it project). Referring to such context, we discuss here some ideas and preliminary results.

Invited Talk (Oct 2nd, 2003 - Room A - h. 9:20-10:10)

Integrating Scalable Process Management into Component-Based Systems Software

Ewing Lusk

*Mathematics and Computer Science Division
Argonne National Laboratory\
USA*

The Scalable Systems Software Project is exploring the design of a systems software architecture based on separate, replaceable components interacting through publicly defined interfaces. This talk will describe how a scalable process manager has provided the implementation of the process management component of that design. We describe a general, implementation-independent definition of process management and how a scalable process management system was used to provide its implementation.

Invited Talk (Oct 2nd, 2003 - Room A - h. 10:10-11:00)

Programming High Performance Applications using Components

Thierry Priol

*IRISA/INRIA
Rennes Cedex
France*

Computational Grids promise to be the next generation of high-performance computing resources. However, programming such computing infrastructures will be extremely challenging. Current programming practices tend to be based on existing and well understood models such as message-passing and SPMD (single program multiple data). On-going works based on Web services (OGSA) aims at programming Grids by specifying the coordination through the expression of interaction and dependencies between already deployed web services. This talk will focus on another alternative that aims at programming Grids with software components. With such an approach, it will be possible to install components on remote resources and to express interaction between components. This presentation will describe work in progress to develop a component-based software infrastructure, called Padico, for computational Grids based on the CORBA Component Model (CCM) from the Object Management Group (OMG). The objective of Padico is to offer a component model targeting multi-physics simulations or any applications that require the coupling of several codes (simulation or visualisation) within a high-performance environment. Two issues will be addressed during the talk: encapsulation of codes into components and runtime support for efficient communication between components within a Grid. This talk will look at the GridCCM component model, an extension to the CCM model to address the encapsulation of SPMD parallel codes into a component and a communication framework, called PadicoTM, able to efficiently support various communication runtime and middleware in a heterogeneous networking environment.

EuroPVM/MPI'03: Poster Session

(Sep 30th, 2003 - h. 17:45-19:15)

Poster Session

(Chair: Raffaele Perego)

1. **An MPI Tool to Measure Application Sensitivity to Variation in Communication Parameters**
E. A. León, A. B. Maccabe, R. Brightwell
2. **Measuring MPI Latency Variance**
R. Riesen, R. Brightwell, A. B. Maccabe
3. **A Parallel Scheme for Solving a Tridiagonal Matrix with Pre-propagation**
A. Wakatani
4. **Competitive semantic tree theorem prover with resolutions**
C. K. Kim, M. Newborn
5. **Explicit Group Iterative Solver On A Message Passing Environment**
N. Hj. Mohd Ali, R. Abdullah, K. J. Lee
6. **Applying Load Balancing in Data Parallel Applications Using DASUD**
A. Cortés, M. Planas, A. Ripoll, M. A. Senar, E. Luque
7. **Performance Analysis of Approximate String Searching Implementations for Heterogeneous Computing Platform**
P. D. Michailidis, Ko. G. Margaritis
8. **Dynamic Load Balancing for the Parallel Simulation of Cavitating Flows**
F. Wrona, P. A. Adamidis, U. Iben, R. Rabenseifner, C.-D. Munz
9. **Message Passing Fluids: molecules as processes in parallel computational fluids**
G. Argentini
10. **Parallel Implementation of Interval Analysis for Equations Solving**
Y. Papegay, D. Daney, J.-P. Merlet
11. **Using MPI from C#**
H. Shukla, J. Zollweg - Cornell Theory Center - Cornell University (not included in the proceedings)
12. **Eigenanalysis of Finite Element 3D Flow Models by Parallel Jacobi-Davidson**
L. Bergamaschi, A. Martinez, G. Pini, F. Sartoretto
13. **Design and Implementation of Dynamic Process Management for Grid-enabled MPICH**
S. Kim, N. Woo, Heon, Y. Yeom, T. Park, H. Park
14. **Scheduling tasks sharing files on heterogeneous clusters**
A. Giersch, Y. Robert, F. Vivien
15. **PVM Over the CLAN Network**
R. Sohan
16. **Improving properties of a parallel program in ParJava Environment**
V. Ivannikov, S. Gaissaryan, A. Avetisyan, V. Padaryan

Special Session: ParSim 03

Current Trends in Numerical Simulation for Parallel Engineering Environments

(Sep 30th, 2003 - h. 17:45-19:15)

ParSim 03

(Room B - Chair: Carsten Trinitis)

Efficient and Easy Parallel Implementation of Large Numerical Simulations

R. Revire, F. Zara, T. Gautier

Toward a Scalable Algorithm for Distributed Computing of Air-Quality Problems

M. Garbey, R. Keller, M. Resch

A piloting SIMulator for maritime and fluvial NAVigation: SimNav

M. Vayssade, A. Pourplanche

Methods and Experiences for Parallelizing flood Models

L. Hluchy, V.D. Tran, D. Froehlich and W. Castaings

***padfem2* - An Efficient, Comfortable Framework or Massively Parallel FEM-Applications**

S. Blazy, O. Kao, O. Marquardt

AUTOBENCH/AUTO-OPT: Towards an Integrated Construction Environment for Virtual Prototyping in the Automotive Industry

A. Kuhlmann, C.-A. Thole, U. Trottenberg