CITI – Centro de Informática e Tecnologia da Informação

Parallel and Distributed Processing Systems

External Review Report

January/2008

Parallel and Distributed Processing Systems

http://asc.di.fct.unl.pt/pdp

- **Mission**: to develop research in parallel and distributed computing with a focus on the following main research themes:

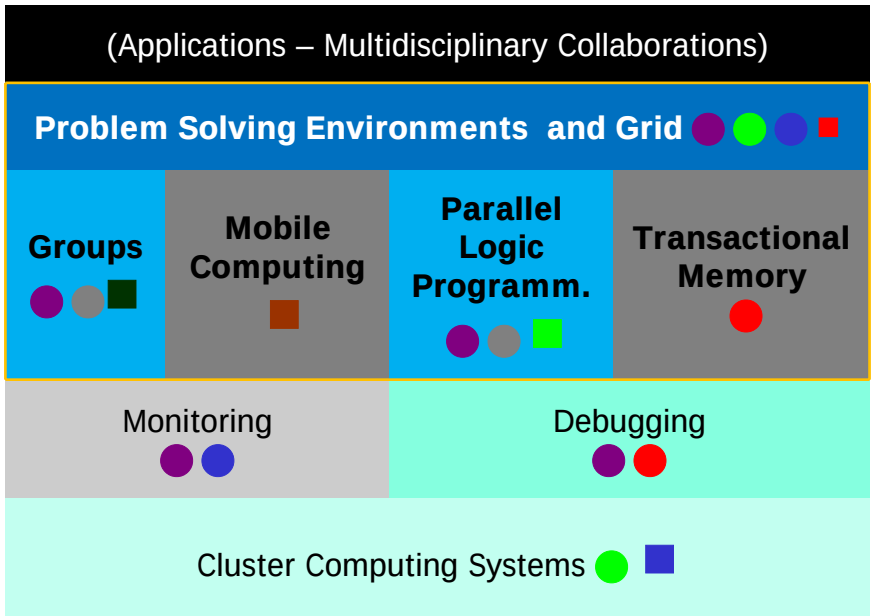

- **Key Researchers**
  - José Cardoso e Cunha, PhD
    - Team coordinator
  - Pedro D. Medeiros, PhD
  - João Lourenço, PhD Nov. 04
  - Fernanda Barbosa, PhD Dec. 04
  - Vitor Duarte, PhD May. 05
  - Hervé Paulino, PhD Dec. 06
  - Rui Marques, PhD Nov. 07
  - Cecília Gomes, PhD Nov. 07
  - Carmen Morgado, MSc, PhD in 08
  - Paulo Lopes MSc PhD in 08
  - **Average # PhD in 2003-2006: 3.5**
- **PhDs**
  - # PhD Completed in 2003-2006: 4
  - # PhD Completed in 2003-2007: 6
  - # PhD Submitted in 2007: 1
  - # PhD Ongoing 2 [+4 starting in 2008]
- **Other researchers**
  - # MSc completed: 1 (2003) + 3 (2007)
  - # MSc ongoing: 15

Challenges

Applications

Large volumes of data, requiring

- Efficient and intelligent management and search
- Parallel and distributed processing

Dynamic, distributed, and mobile application entities, requiring management of

- Structure, interaction, and coordination

Integration of distributed, heterogeneous components in dynamic and interactive environments

Organisation of small / medium / large scale collections of distributed, intelligent entities

Computing Environments

Increased levels of **concurrency, parallelism & distribution**

New forms of **dynamic behavior**

- increasing levels of interaction among system and application components
- more frequent changes in interaction and behavior
- more frequent changes in system and application configurations
- mobility

Increasing **scale** in terms of system and application components

Main Goal to address the challenges

- The uptake of **new generations of parallel and distributed computing applications** is being restricted by the availability of suitable abstractions, methodologies, and tools.
- Our goal is to identify:
  - **Abstractions, models and tools** that can be integrated in **parallel and distributed computing environments**,
  - along with problem-domain tools,
  - to enable the deployment of applications.
  - with **improved functionality and efficiency**

Parallel and Distributed Processing Systems

http://asc.di.fct.unl.pt/pdp

Problem-Solving Environments and Grid Computing

Abstractions & Models

Group-oriented Abstractions and Models

Service-oriented Mobile Computations

Parallel and Distributed Logic Programming

Software Transactional Memory

Tools for Parallel and Distributed Program Development

Cluster Computing Systems

Problem-Solving Environments and Grid Computing

Software Environments for Grid Computing

Collaboration with Omer Rana, Cardiff U., Distributed Collab. Group (2003-...)

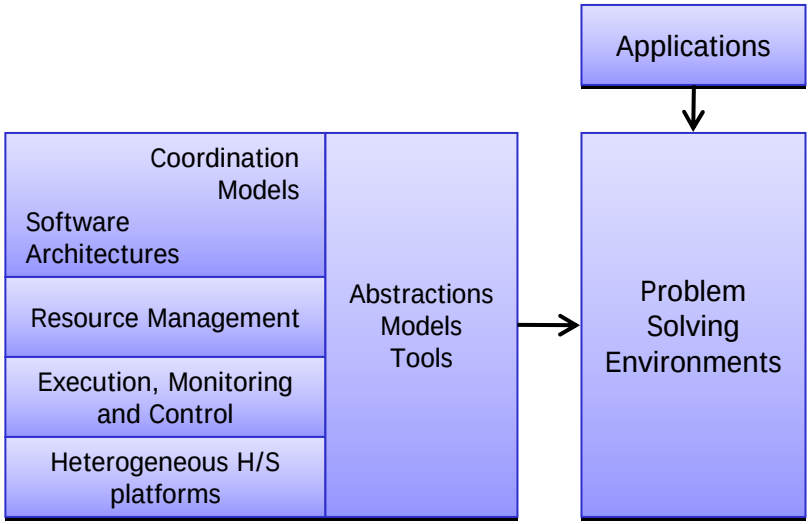

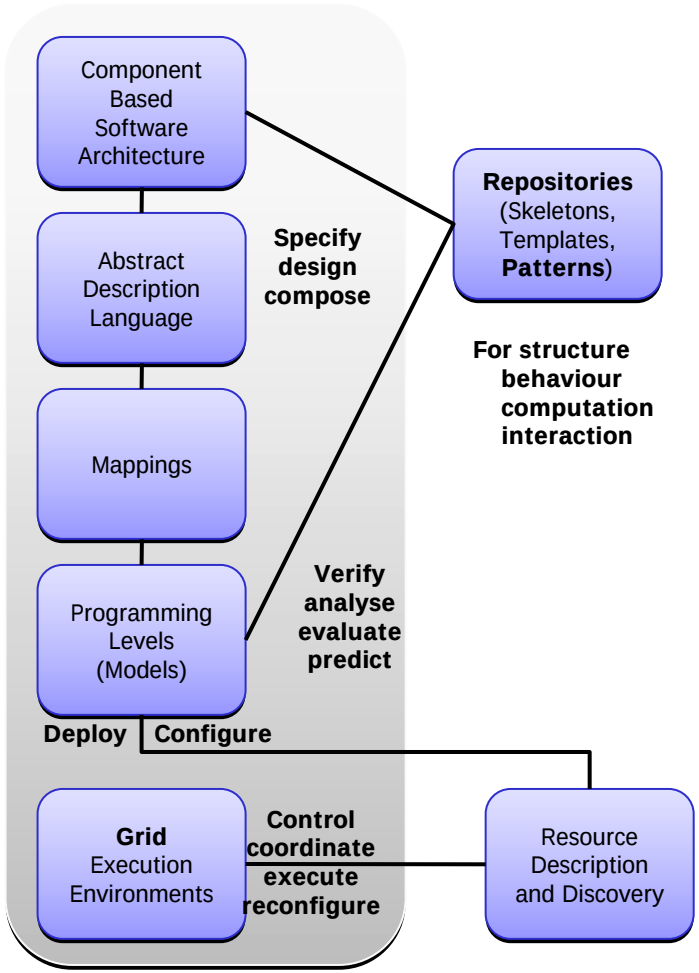

Results and Directions:

Book: Grid Computing: Software Environments and Tools, Editors: Jose Cunha and Omer Rana, Springer (2006)

Journal paper: Future Trends in Distributed Applications and Problem-Solving Environments, Jose Cunha, Omer Rana, Pedro Medeiros, Future Generation Computer Systems, Elsevier (2005)

One completed PhD thesis: Cecilia Gomes, 2007

Projects:

- EU AsiaLink Curriculum Development for HPC and Grid Computing (2005-08)
- Participation CoreGRID as Associate Member WP3-Programming, (mid-2007-)

Design Patterns and Operators for Distributed Grid Environments

M. Cecília Gomes, PhD Thesis, November 2007, FCT/UNL

Goal:

To abstract common structures and behaviours

To integrate them in software environments

Approach:

Model for reuse and flexible composition:

Structural Patterns: *Pipeline, Star, Ring, Proxy, Facade, Adapter, ...*

Behavioural Patterns: *Master/Slave, Producer/Consumer, Peer-to-peer, ...*

Patterns as *first-class entities* in application life-cycle:
- development
- execution
- reconfiguration

Model integrated in a prototype extending Triana workflow system, in collaboration with Cardiff Un.

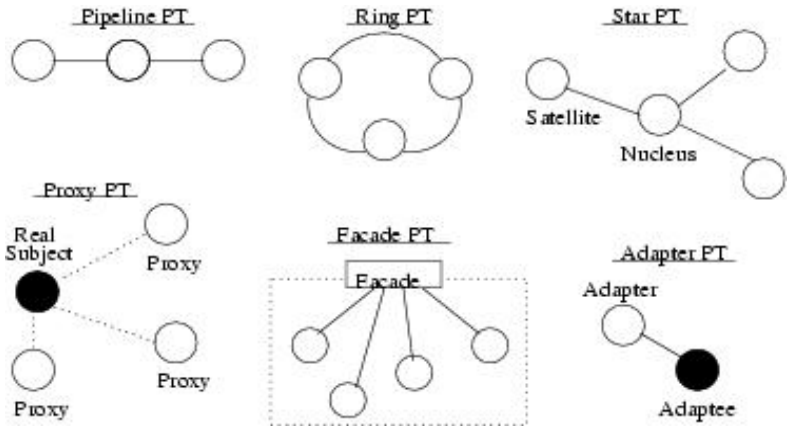

Structural Operators: *Increase, Decrease, Extend, Reduce, Embed, Extract, Group, Rename/Reshape*, ...

Behavioural Operator: *Start, Terminate, Log, Stop, Resume, Restart, Limit, Repeat,* ...

Results :

PhD thesis:	**1** (Nov´07)
Journal articles:	**1 [+1]**
Conf. papers:	**4** (1 as a book chap)
Prototypes:	**1**

Pattern-based Methodology for Application Development

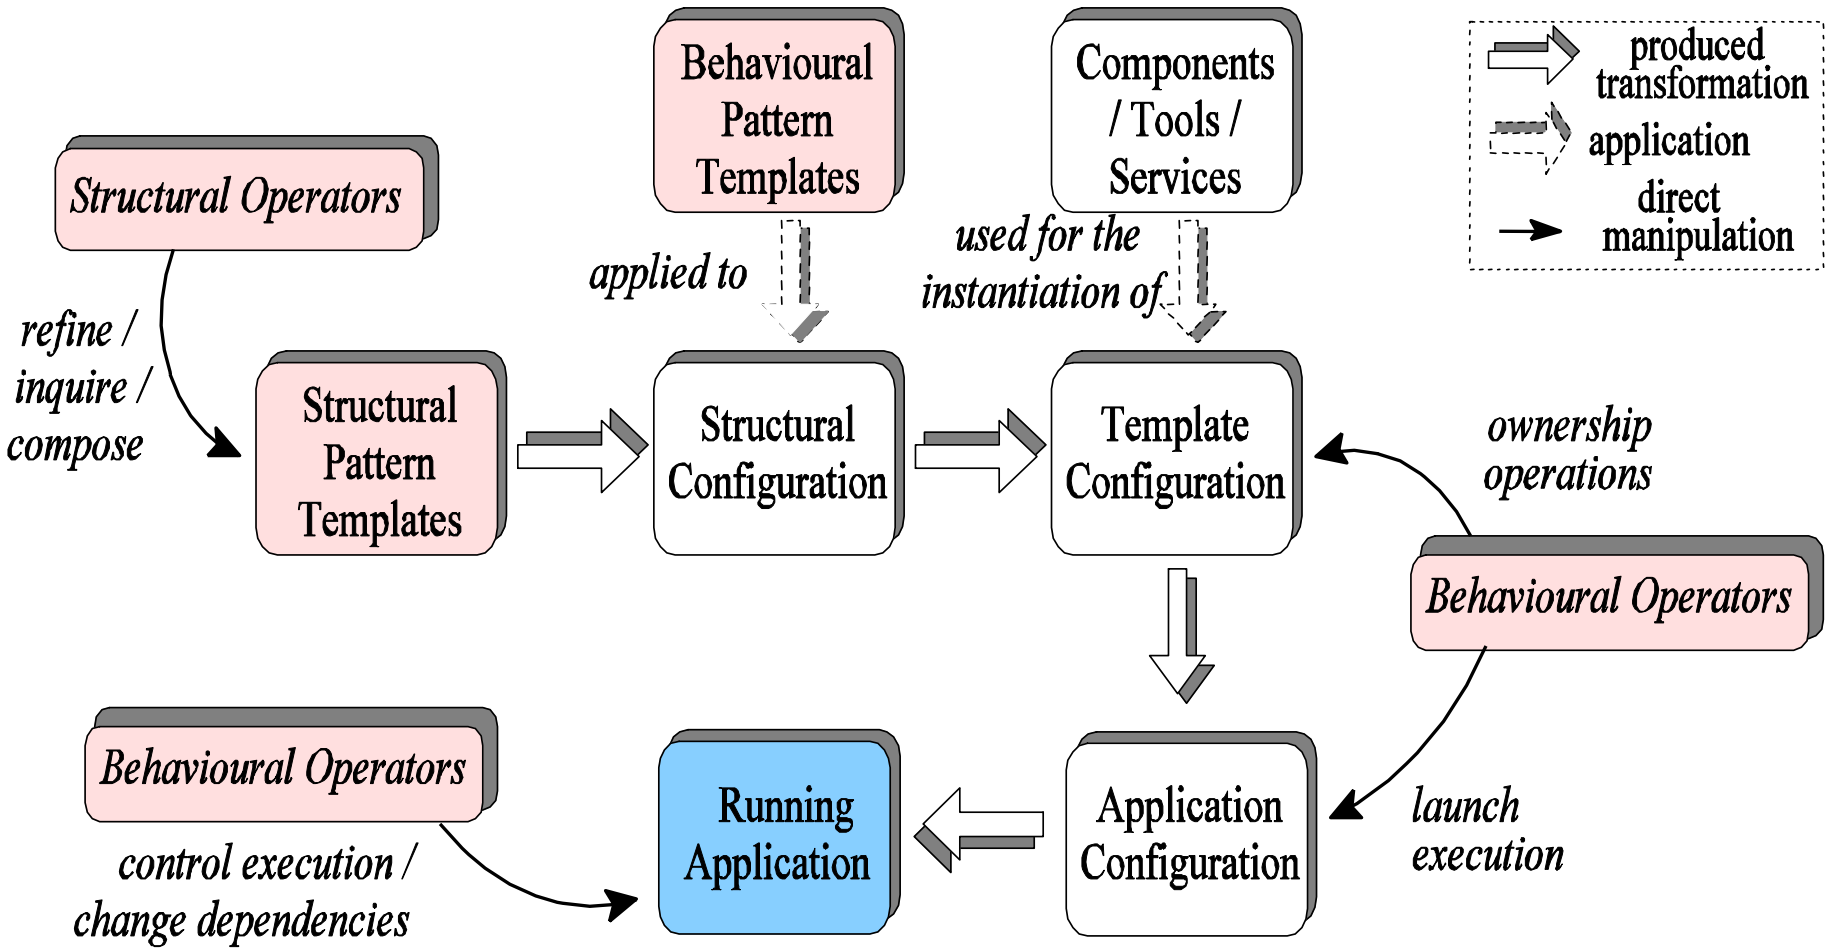

A GeoScience View: Application of computer algorithms for reservoir and aquifer modelling

Collaboration with Center for Geological Science FCT/UNL

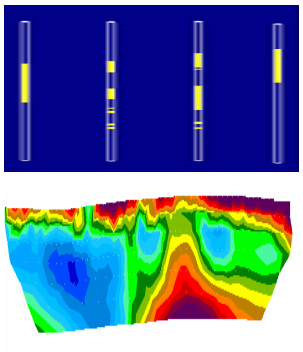

Data

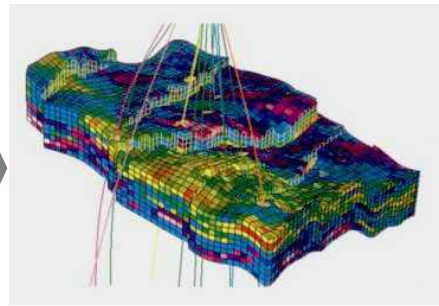

Static model

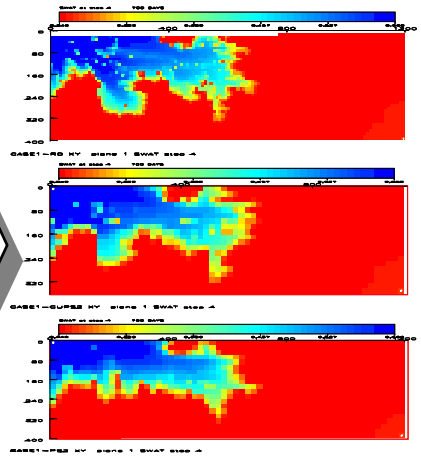

Dynamic model

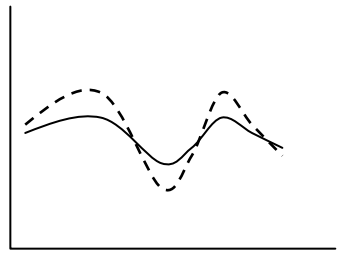

Outputs

$output = f(i_1, i_2, i_3 \ldots)$

- Geological units
- Permeability
- Porosity
- Fluid saturation
- Exploitation and recovery

Ongoing and future work:

Two new PhD theses start Jan´08 at CITI - Scientific Workflows for Parallel and Distributed Computing

Post-doctorate on HPC and Grid Computing - GeoInfo Project at CITI starting March´´08, 5-years contract, Portuguese FCT,

New Project proposal planned with Cardiff Un.

Scientific Workflows Abstractions and Models for Parallel and Distributed Computing

Center for Geological Sciences FCT/UNL:

Numerical simulation of properties with sequential geostatistical algorithms (static model).

Simulation of reservoirs with integration of seismic information (static model).

Upscaling of permeability (interface static-dynamic model) and simulation of fractures.

Stochastic inversion problems (conditioning of static models to dynamic data

CITI / Parallel and Distributed Processing Stream:

Workflow models are increasingly used to specify scientific applications;

Preliminary experiments with *Triana*, *Kepler* and *Taverna* have shown that many issues are still open;

Expressiveness of Workflow Models: Large-scale parallel and distributed / Event-driven and asynchronous notifications / Sharing state / Data access and transfer / Adaptive and dynamic

Flexible mappings to the execution support environments: Configuration/Deployment / New requirements for workflow engines

Two PhD Theses (Planned to start Jan. 2008)

Luís Assunção / Carlos Gonçalves

Supervisor: José C. Cunha

Results in more detail

PhD thesis completed: 1

Maria Cecília Gomes, *Pattern Operators for Grid Environments*, Nov´07, sup: José Cunha. **[PSE and Grid]**

PhD thesis planned: 2

Luís Assunção (M.Sc) starting as PhD student Jan 2008, sup: José Cunha.

Carlos Gonçalves (M.Sc.) starting as PhD student Jan 2008, sup: José Cunha.

MSc thesis ongoing: 2

Paulo Carmo, *Design of a Grid-based Automated Processing System*, sup: José. Cunha.

Tânia Sapage, *Methodologies, Environments and Tools for Grid Application Development*, sup: José. Cunha.

Publications: 7 [+1]

Book

J. Cunha and O. Rana, Grid Computing: Software Environments and Tools, Springer-Verlag. 2006.

Articles

M. C. Gomes, O. Rana, and J. C. Cunha. Extending grid-based workflow tools with patterns/operators. IJHPCA - The International Journal of High Performance Computing Applications, 22(2). 2008.

J. C. Cunha, O. Rana, P. Medeiros. Future Trends in Distributed Applications and Problem-Solving Environments, FGCS - Future Generations of Computer Systems, Special Issue on "Complex Problem-Solving Environments for Grid Computing", Vol 21 No. 6. Elsevier. 2005

M. C. Gomes, O. Rana, and J. C. Cunha. Pattern operators for grid environments. Scientific Programming, 11(3):237-261, 2003.

Book Chapter

M. C. Gomes, J. C. Cunha, and O. Rana. A Pattern-based Software Engineering Tool for Grid Environments, pages 213-222. NATO Science Series III: Computer and Systems Sciences. IOS PRESS, 2005.

International Conferences and Workshops (with peer review)

M. C. Gomes, O. Rana, and J. C. Cunha. Pattern/operator based problem solving environments. In M. Danelutto, D. Laforenza, and M. Vanneschi, editors, Euro-Par 2004 Parallel Processing, 10th International Euro-Par Conference, number 3149 in Lecture Notes in Computer Science, pages 960-971. Springer-Verlag, 2004.

O. Rana, M. C. Gomes, and J.C. Cunha. Patterns and operators for grid software development. In Proceedings of the IADIS International Conference WWW/Internet 2003. IADIS Press, 11 2003.

O. Rana, M. C. Gomes, and J.C. Cunha. A pattern based software engineering tool for grid environments. In D. Grigoras, A. Nicolau, and F.L.Tiplea, editors, Proceedings of NATO Advanced Research Workshop on Concurrent Information Processing and Computing, CIPC2003, pages 107-116. Al.I.Cuza University Press, 07 2003.

Problem-Solving Environments and Grid Computing

Abstractions & Models

Group-oriented Abstractions and Models

Service-oriented Mobile Computations

Parallel and Distributed Logic Programming

Software Transactional Memory

Tools for Parallel and Distributed Program Development

Cluster Computing Systems

Group-oriented Abstractions and Models

Groups at distinct abstraction levels:

- As Application Units
- As Programming Units
- As System Units

GroupLog, an abstract model: agents, groups, forms of interaction

Distinct instances, at distinct abstraction levels:

Logic-based instance : GroupLog

Java-based instance: JGroupSpace with cluster implementation

MAGO: distributed interactive applications

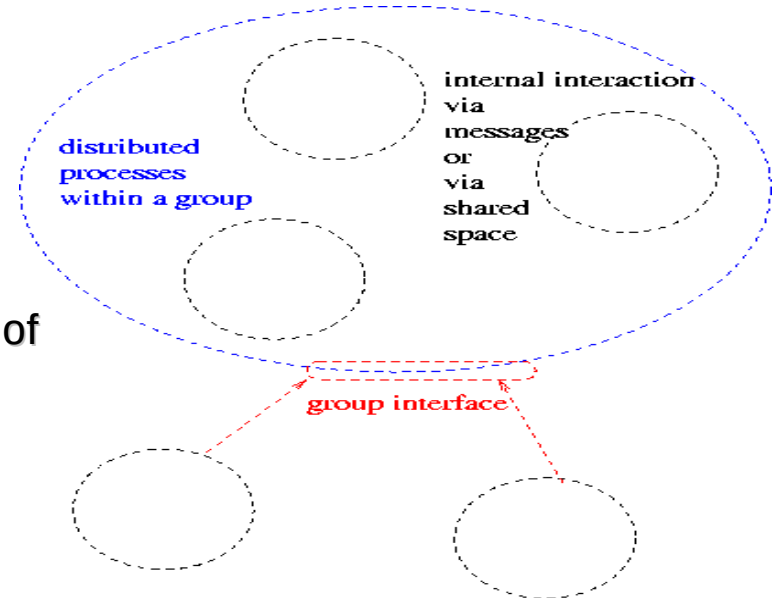

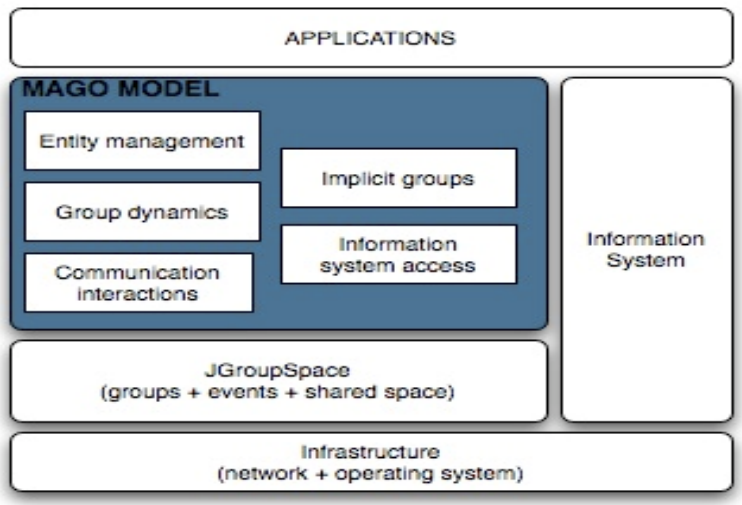

Results :

- PhD theses: **1** (on the GroupLog Model in 2004)
  [+1] subm. 2007 awaiting defense]
- Ongoing M.Sc thesis: **1**
- Conf. papers: **3** (**1** as book chapter) [**+1**]
- Prototypes: **3**

Project: **1** (InStory with IMG CITI Stream)

MAGO - Modeling Applications with a Group-based Approach

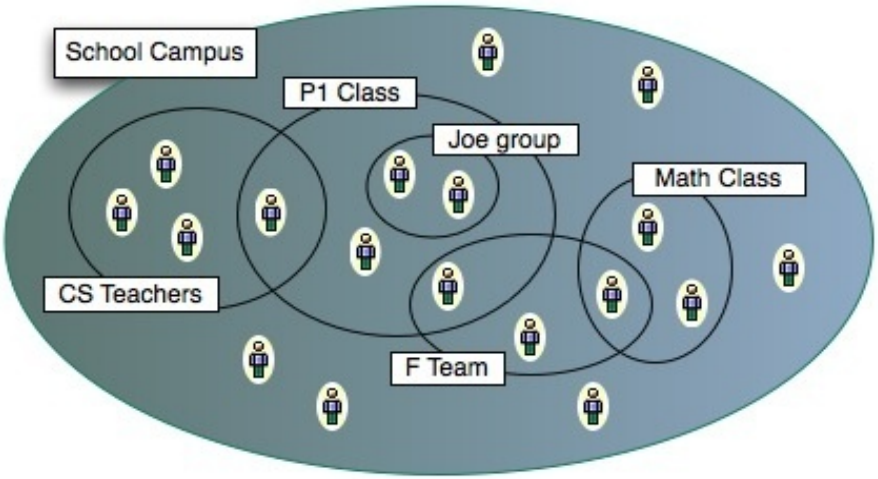

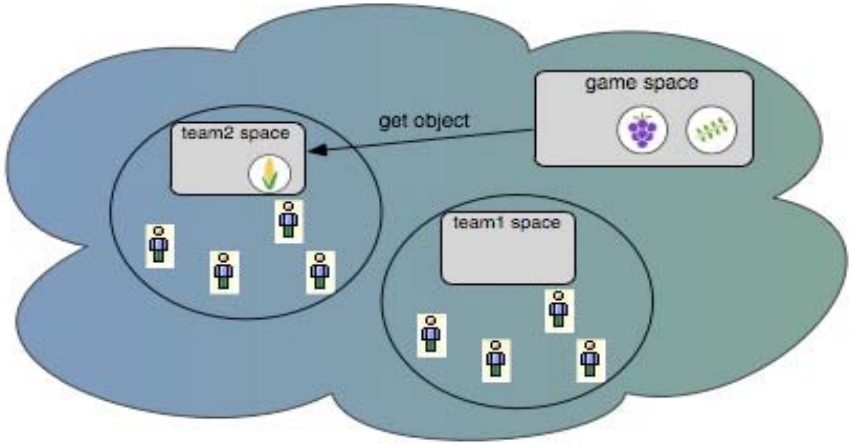

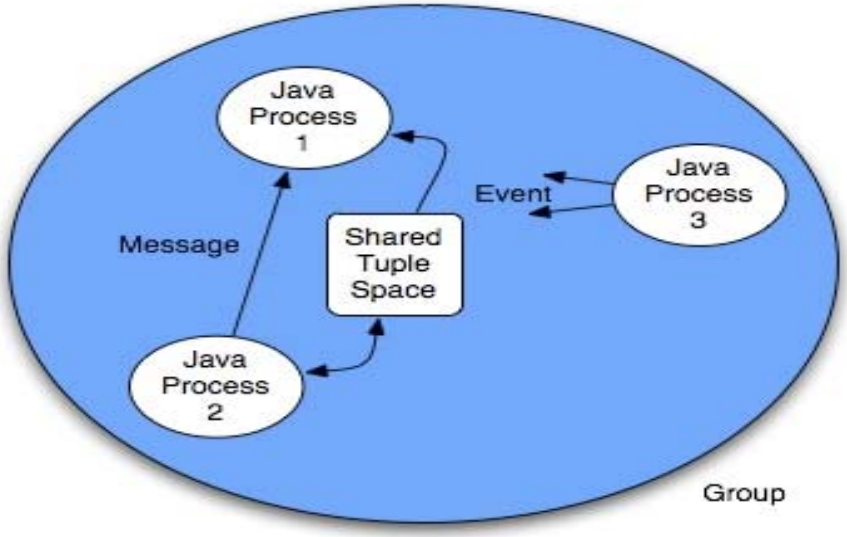

Results in more detail

PhD thesis completed: 1

Fernanda Barbosa, *Group-Based Distributed Programming Abstractions: the GroupLog Model*, Dec´04, sup: José Cunha.

PhD thesis submitted in 2007, awaiting defense : 1

Carmen Morgado, *A Group-based Model for Distributed Interactive Applications*, Thesis submitted Dec´07, sups: José Cunha / Nuno Correia.

MSc thesis ongoing : 1

Jorge Custódio, *JGroupSpace: an API to Support Group-oriented Programming*, subm. 2008, sup: J. Cunha.

Publications : 3 [+1]

Book Chapters

F. Barbosa, J. C. Cunha, O. Rana, and Stephen Lynden. Coordination in Utility Managed Multi-Agent Group, volume 15 of Advances in Computation: Theory and Practice, pages 209-219. Nova Science, l. t. yang and y. pan edition, 2004.

F. Barbosa, J. C. Cunha, O. Rana, and Stephen Lynden. Coordination in Utility Managed Multi-Agent Groups, pages 995-1000. Swets & Zeitlinger B.V., 2003.

International Conferences (with peer review)

C. Morgado, N. Correia, and J. C. Cunha. A group-based approach for modeling interactive mobile applications. International ACM Conference on Supporting Group Work - Group07, 11 2007.

C. Morgado and L. Soares. Mig21 api: multimedia interactive groups api. In ITiCSE '05: Proceedings of the 10th annual SIGCSE conference on Innovation and technology in computer science education, pages 399-399. ACM Press, 2005.

Problem-Solving Environments and Grid Computing

Abstractions & Models

Group-oriented Abstractions and Models

Service-oriented Mobile Computations

Parallel and Distributed Logic Programming

Software Transactional Memory

Tools for Parallel and Distributed Program Development

Cluster Computing Systems

Service-Oriented Computing and Mobility

Seamless combination of service and software mobility paradigms.

The **Mob** programming language:

- High level model for programming distributed applications
- Agents can provide/require services
- Strong mobility of agents on top of Java Virtual Machine
- Bindings for agents based on services provided by agents
- **Compiled to the DiTyCO process calculus**
- **Run-time system built on top of the one for DiTyCO**.

Protocols for reliable communication in the presence of mobility.

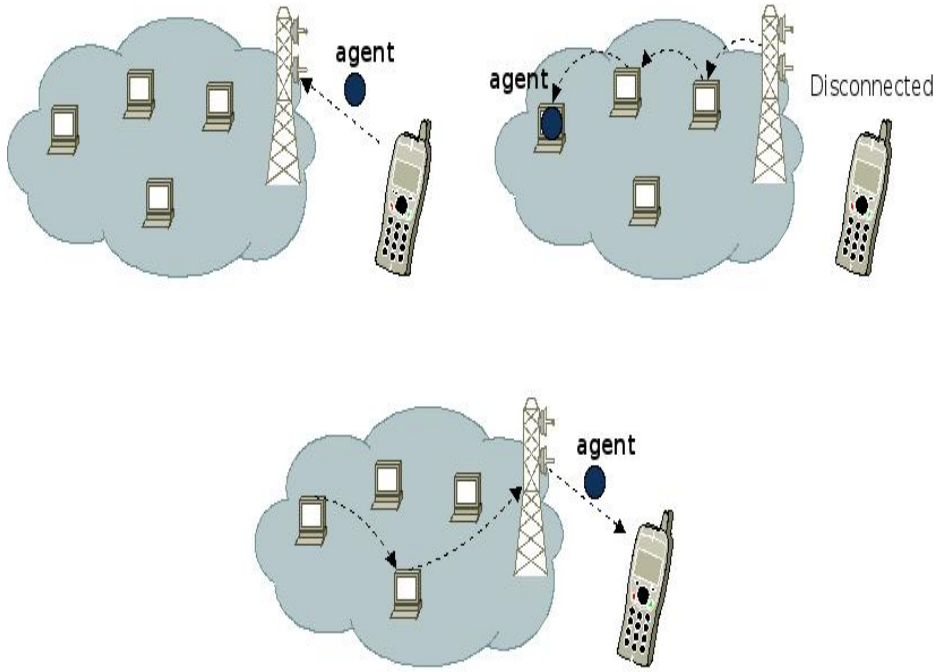

Results:

PhD thesis:	**1** (Hervé Paulino, Dec´06)
Ongoing MSc Thesis:	**1**
Journal articles:	**1 [+1]**
Conference papers:	**5 [+1]**
Workshop papers:	**2 [+1]**
Prototypes:	**3**
Projects:	**1** (MIMO, with U.Porto)

Ongoing directions :

To feature the model in a Java framework

Service-centric mobility: extension to express mobility in terms of services rather than hosts.

Service-oriented middleware for wireless sensor networks in the Callas national project with U. of Porto and Lisboa

Results in more detail

PhD thesis completed: 1

Hervé Paulino, *An Infrastructure for Mobile Service-Oriented Computing Encoded on a Process Calculus*, Dec.´06, sups: L.Lopes /F.Silva, U. Porto.

MSc thesis ongoing: 1

Carlos Tavares, *Towards Full Separation of Concerns in Software Mobility*, sup: H. Paulino.

Publications: 8 [+3]

Articles

H. Paulino and L. Lopes. A programming language and a run-time system for service-oriented computing with mobile agents. Software: Practice and Experience, 2008.

Á. Figueira, H. Paulino, L. Lopes, and F. Silva. Distributed typed concurrent objects: a programming language for distributed computations with mobile resources. Journal of Universal Computer Science, 8(9):745-760, 08 2003.

International Conferences and Workshops (with peer review)

H. Paulino. Reliable communication in the presence of agent Mobility. In 12th IEEE Symposium on Computers and Communications (ISCC 2007), pages 989-994. IEEE Computer Society, 07 2007.

H. Paulino. Mobile service development and deployment with remotely launched service-oriented mobile agents. In Proceedings of the Workshop on Mobile Services-oriented Architectures and Ontologies (MoSO 2007). IEEE Computer Society, 05 2007.

H. Paulino and L. Lopes. A mobile agent service-oriented scripting language encoded on a process calculus. In David Lightfoot and Clemens Szyperski, editors, 7th Joint Modular Languages Conference, JMLC 2006, number 4228 in Lecture Notes in Computer Science, pages 383-402. Springer-Verlag, 09 2006.

H. Paulino and L. Lopes. A service-oriented language for programming mobile agents. In Peter Stone and Gerhard Weiss, editors, Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems, pages 1294-1296. ACM Press, 05 2006.

L. Bettini, R. Nicola, D. Falassi, M. Lacoste, L. Lopes, L. Oliveira, H. Paulino, and V. T. Vasconcelos. A software framework for rapid prototyping run-time systems for mobile calculi. In Paola Priami, Corrado; Quaglia, editor, Global Computing: IST/FET International Workshop, GC 2004, Rovereto, Italy, March 9-12, 2004, Revised Selected Papers, number 3267 in Lecture Notes in Computer Science, pages 179-207. Springer-Verlag, 08 2005.

H. Paulino, L. Lopes, and Fernando Silva. Mob: a scripting language for programming web agents. In Sergio Luján-Mora Pedro J. Clemente, Miguel A. Pérez and Hans Reiser, editors, 13th Workshop for Phd Students in Object Oriented Programming: Summary and Accepted Papers, pages 21-29. Friedrich-Alexander-University, 10 2003.

Álvaro Figueira, H. Paulino, L. Lopes, and Fernando Silva. Distributed typed concurrent objects: a programming language for distributed computations with mobile resources. In Proceedings of the 7th Brazilian Symposium on Programming Languages (SBLP 2003). Universidade Federal de Minas Gerais, 05 2003.

H. Paulino, P. Marques, L. Lopes, V. Vasconcelos, and F. Silva. A Multi-threaded Asynchronous Language. In Victor E. Malyshkin, editor, Parallel Computing Technologies, 7th International Conference, PaCT 2003, volume 2763 of Lecture Notes in Computer Science, pages 316-323. Springer-Verlag, 2003.

H. Paulino, L. Lopes, and Fernando Silva. Mob: A scripting language for mobile agents based on a process calculus. In J.M. Cueva Lovelle, B.M. González Rodríguez, L. Joyanes Aguilar, J.E. Labra Gayo, and P. del Puerto, editors, Web Engineering - International Conference, ICWE 2003, number 2272 in Lecture Notes in Computer Science, pages 40-43. Springer-Verlag, 07 2003.

Problem-Solving Environments and Grid Computing

Abstractions & Models

Group-oriented Abstractions and Models

Service-oriented Mobile Computations

Parallel and Distributed Logic Programming

Software Transactional Memory

Tools for Parallel and Distributed Program Development

Cluster Computing Systems

Multi-threaded Tabling in XSB Prolog

Tabled Logic Programming

- Guarantees termination for datalog logic programs
- Reduces the temporal complexity for many logic programs

XSB – a Tabled Prolog System

xsb.sourceforge.net – developed in cooperation:

- SUNY at **S**tony **B**rook - USA
- Universidade Nova de Lisboa - Portugal
- Katholieke Universiteit Leuven - Belgium
- Uppsala Universitet – Sweden

Integrates the latest research results in tabling in a Prolog system for real world applications

Multi-Threaded Tabling extensions

•Explicit parallelism allowing several queries to be executed concurrently in XSB

•Programming model integrates thread-private and thread-shared tables

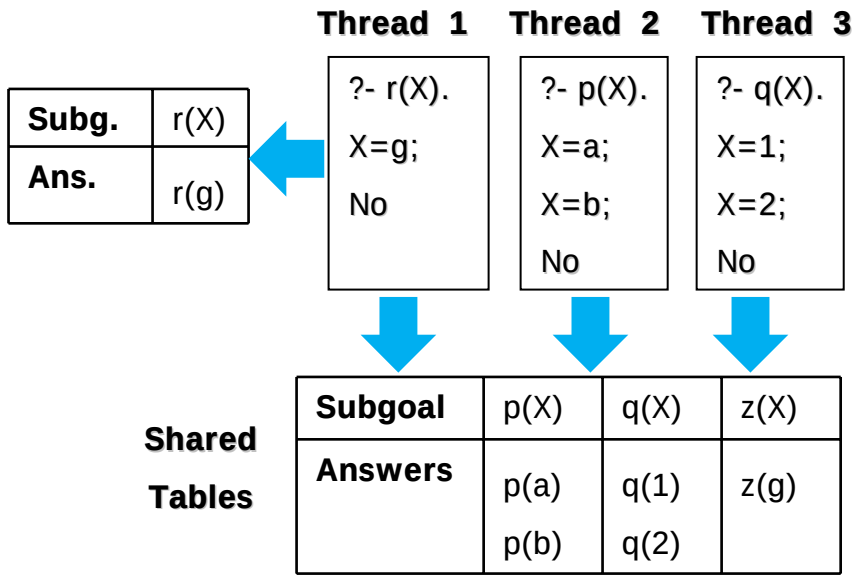

Results:

- **One PhD thesis**: Rui Marques, Concurrent Tabling: Algorithms and Implementation, November 2007 (supervisors: José C. Cunha and Terrance Swift, SUNY at Stony Brook)
- **Multi-threaded XSB is publicly available and is being used by the community**

Multi-Threaded Tabling in XSB

Strong aspects of the work

- Provided an useful extension to an already useful and intensively used system
- **XSB had more than 500 downloads per month in 2007**
- Demands continued maintenance and user support

Performance Results on a 2x2 Core

Average speedups on an 2x2 core machine
for some programs with no data dependencies

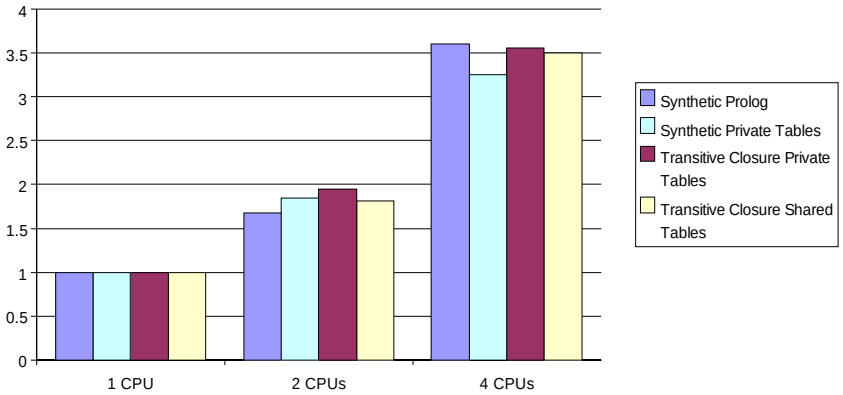

Open Research Directions

- Integrate more Tabling features with the multi-threading extensions – XSB keeps integrating new tabling features
- Extend XSB with distributed programming features (begin with MPI-based extensions)
- Explore Table Parallelism – a new form of implicit logic parallelism, other than the traditional and/or parallelism

Commercial Applications by XSB inc. (www.xsb.com)

- Weapon System Impact Tool, Warwick - Government Business Intelligence, Deductive Spreadsheet, etc
- Rely on CDF (Ontology Classifier program that uses Tabling in XSB) – starting to be Multi-threaded for Parallelism and XSB-Java Interface (already multi-threaded)

Research Applications by the user community

- Flora (Logic Object-oriented Transactional Language) – to be Multi-threaded for Distributed Agent Environments
- XMC (model checker) – to be Multi-threaded for Parallelism

An ultimate goal: Distributed Grid Agents in Logic Frameworks: this is still a vision...

To develop distributed computing models and architectures to support distributed agent systems.

A basis for supporting planning, decision support, and intermediation between user and system levels in distributed grid computing systems (*Semantic Grid*)

How agent abstractions can support reasoning, planning, intelligent decision support and intermediation

- Between the user and a PSE
- To support reactive and autonomic behaviour at application and system levels

Results in more detail

PhD thesis completed: 1

Rui Marques, *Concurrent Tabling: Algorithms and Implementation*, Nov.´07, sups: José Cunha / Terrance Swift, SUNY at Stony Brook.

Problem-Solving Environments and Grid Computing

Abstractions & Models

Group-oriented Abstractions and Models

Service-oriented Mobile Computations

Parallel and Distributed Logic Programming

Software Transactional Memory

Tools for Parallel and Distributed Program Development

Cluster Computing Systems

Software Transactional Memory

[since mid 2006]

- Concurrency control on shared-memory (SMP and multi-core) architectures
  - Optimistic approach
  - Higher level of abstraction and less error-prone
- Current and future research lines
  - Performance optimization
    (1 MSc concluded, 1 undergrad student)
    - Algorithms, monitoring, benchmarking
  - Increase functionality
    (2 MSc ongoing)
    - Development of a Transactional Filesystem
    - Unified model with database transactions
  - Correction testing and debugging
    (1 MSc ongoing)
    - Dynamic instrumentation of STM-based source code
  - Applications of STM
    (1 MSc ongoing)
    - In production environments
    - In speculative execution
      (participation in the Byzantium Project- DSG group)

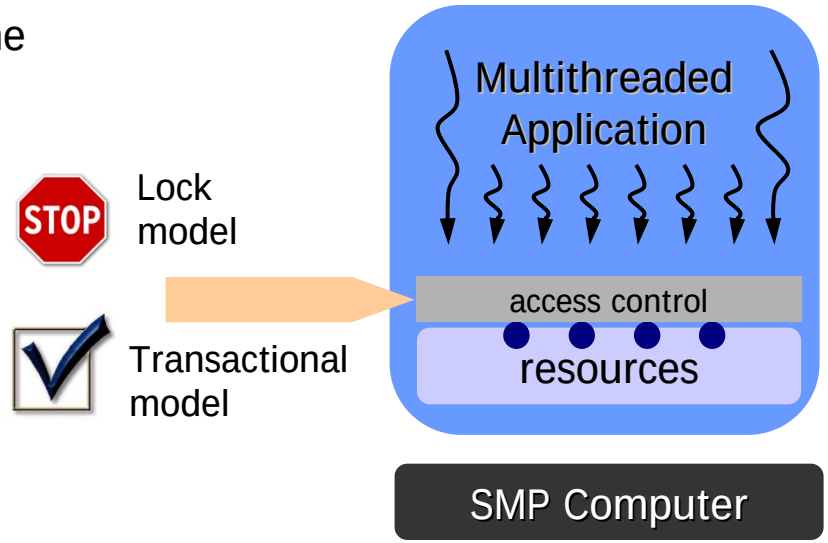

Results (in 2007):

MSc Thesis:	**1** completed, (4 ongoing)
Conf. Papers:	**1**
Prototypes:	**1**
Projects:	**1 (**with DSG, 2008**)**

Results in more detail

MSc thesis completed: 1

Gonçalo Cunha: *Consistent State Software Transactional Memory*, 2007, supervision J. Lourenço.

MSc thesis ongoing: 4

Ricardo Dias, *Cooperative Memory and Database Transactions*, sup: J. Lourenço.

Ana Ferreira, *Testing of Software Transactional Memory Programs*. Sup: J. Lourenço.

Artur Martins, *Transactional File System*, sup:J. Lourenço.

Paulo Matos, *Adaption of an Applicational Server to Support Non-Persistent Transactions*, sup: J. Lourenço.

Publications: [+1]

International Workshop (with peer review)

J. Lourenço and G. Cunha. Testing patterns for software transactional memory engines. In Proceedings of International Symposium on Software Testing and Analysis - Workshop on Parallel Testing and Debugging, pages 36-42. ACM Press, 07 2007.

Problem-Solving Environments and Grid Computing

Abstractions & Models

Group-oriented Abstractions and Models

Service-oriented Mobile Computations

Parallel and Distributed Logic Programming

Software Transactional Memory

Tools for Parallel and Distributed Program Development

Cluster Computing Systems

[Tools] Distributed Monitoring

- Monitoring of parallel/distributed applications:
  - Application and system behavior
  - Performance and QoS
- Services as basic building blocks
- The system architecture is independent of the application runtime or instrumentation technique used

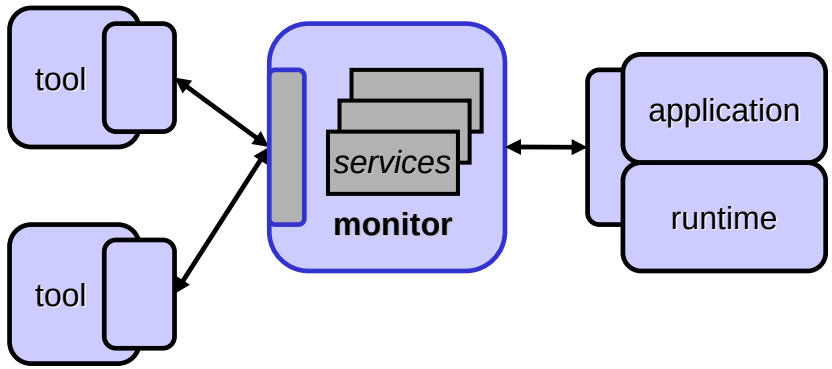

Ongoing work: monitoring distributed applications behavior, performance and resource usage

Case studies:

Scientific workflow applications on Grid systems
Distributed Java applications

Results :

PhD thesis:	**1** (Vítor Duarte 2005)
Prototype:	**1** (DAMS)
Papers:	**[+1** accepted 2007]
Projects:	**3**
Ongoing PhD:	**1** (started 2007)
Ongoing MSc:	**1**

Results in more detail

PhD thesis completed: 1

Vítor Duarte, *An Architecture for Monitoring Parallel and Distributed Computations*, May´05, sup: José Cunha.

PhD thesis ongoing: 1

Liu Wei, *Grid Workload Design for an Intelligent Grid Environment*, PhD Student (GUCAS, Graduate School of Chinese Academy of Sciences), AsiaLink Programme, joint supervision J.Cunha & V.Duarte (UNL), and Tiejian Luo (GUCAS, China): 2007 stay at CITI

MSc thesis ongoing: 1

Paulo Mendes, *Services and Infrastructure Monitoring of Data Centerrs*, M.Sc. Student, sup: V. Duarte.

Publications: [+1]

International Conference (with peer review)

L. Wei and V. Duarte and J. C. Cunha. Towards a Framework for Monitoring Grid Scientific Workflows. IEEE International Conference on Networking, Sensing and Control, 2008. IEEE Computer Society, 2008.

[Tools] Flexible Interface for Distributed Debugging – Library and Engine [until end of 2004]

- Interactive debugging of distributed applications
- Support for multiple debugging tools operating simultaneously over the same target application
  - Allow debugging activity at different abstraction levels
  - Providing multiple distinct views of the target application state

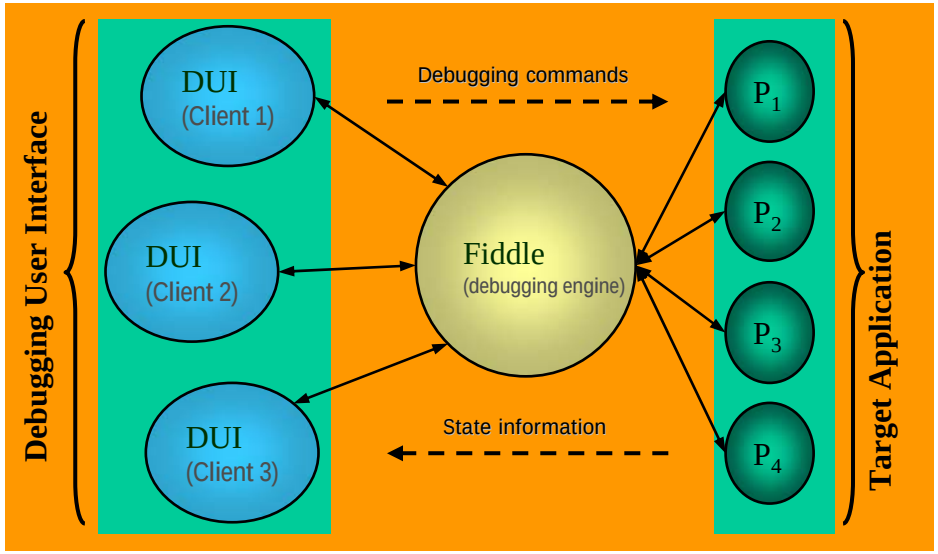

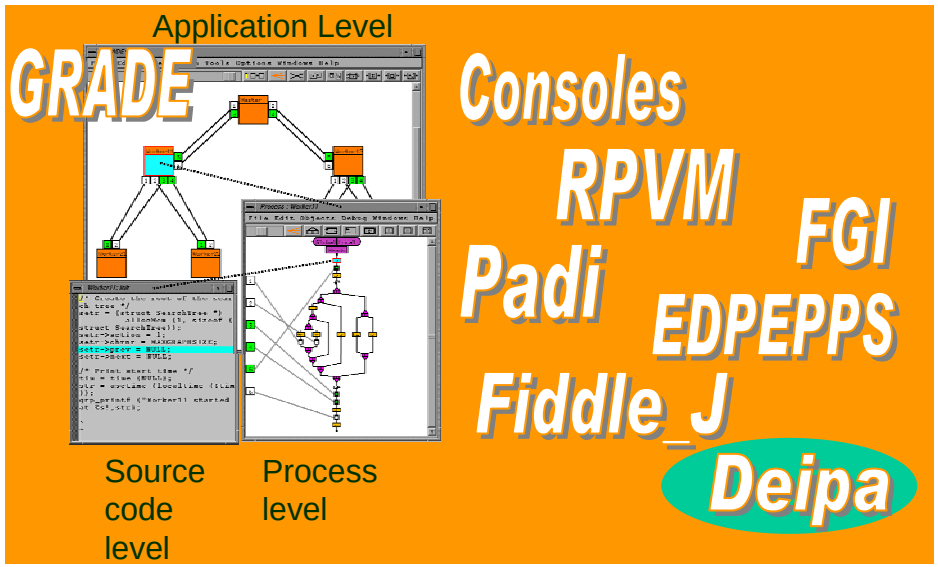

Results:

PhD thesis:	**1** (João Lourenço, Nov´4)
Conf. papers:	**1**
Prototypes:	**1** (FIDDLE)

Results in more detail

PhD thesis completed: 1

João Lourenço, *A Debugging Engine for Parallel and Distributed Programs*, Nov´04, sup: José Cunha.

Publications: 1

International Conference (with peer review)

J. Lourenço, J. C. Cunha, and V. Moreira. Control and debugging of distributed programs using fiddle. In es) Michiel Ronsse, editor, Proc. of the 5th International Workshop on Automated Debugging (AADebug'2003), pages 143-157. arXiv, 2003.

Parallel and Distributed Processing Systems

http://asc.di.fct.unl.pt/pdp

Problem-Solving Environments and Grid Computing

Abstractions & Models

- Group-oriented Abstractions and Models
- Service-oriented Mobile Computations
- Parallel and Distributed Logic Programming
- Software Transactional Memory

Tools for Parallel and Distributed Program Development

Cluster Computing Systems

Cluster Computing Systems

Runtime environments for high-performance computing

Core:

- Lightweight inter-process communication mechanisms
- **File systems for high-performance computing**
- Parallelization of science and engineering codes

Emerging

- File systems for grid systems

Results (in the three core topics):

M.Sc thesis: **1 [+2]**

Conf. Papers: **3**

Project: **1**

PhD thesis: **[+1]**

(Paulo Lopes, to subm. Mar 2008)

Approach:

- pCFS, a parallel file system for clusters with shared disk systems;
- uses cooperative cache and all interconnects available to achieve high performance and reliability.

Cooperation with IRISA and KerLabs; extension to wide area settings

Participation in CoreGRID WP2 – Knowledge & DataMgt Task since mid-2007

Parallelization of science and engineering codes

Context: Using clusters and pools of PCs for executing parallel versions of science and engineering codes

Application areas: [Multidisciplinary Collaborations]

- Tomography in Materials Science
- Pollutant dispersion in shallow waters in Environmental Science

Approach:

- Shared and distributed memory
- Standard message passing (PVM/MPI) and .NET
- Dedicated frameworks

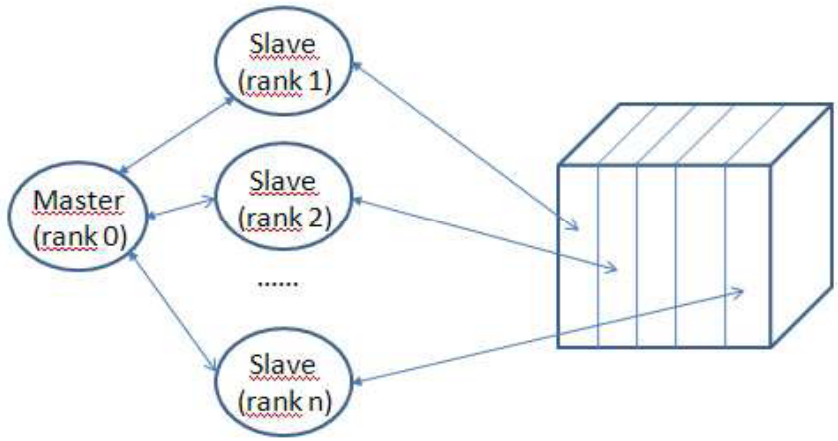

Use of the developments made in the areas of parallel file systems and lightweight communication

Results:

MSc Thesis: [+2]

Ongoing MSc Thesis: 2

PhD Thesis: 1 [starting in 2008]

File Systems for High Performance Computing

Context: Efficient and reliable file access in cluster environments

Approach:

pCFS, a parallel file system for clusters with shared disk systems;

uses cooperative cache and all interconnects available to achieve high performance and reliability.

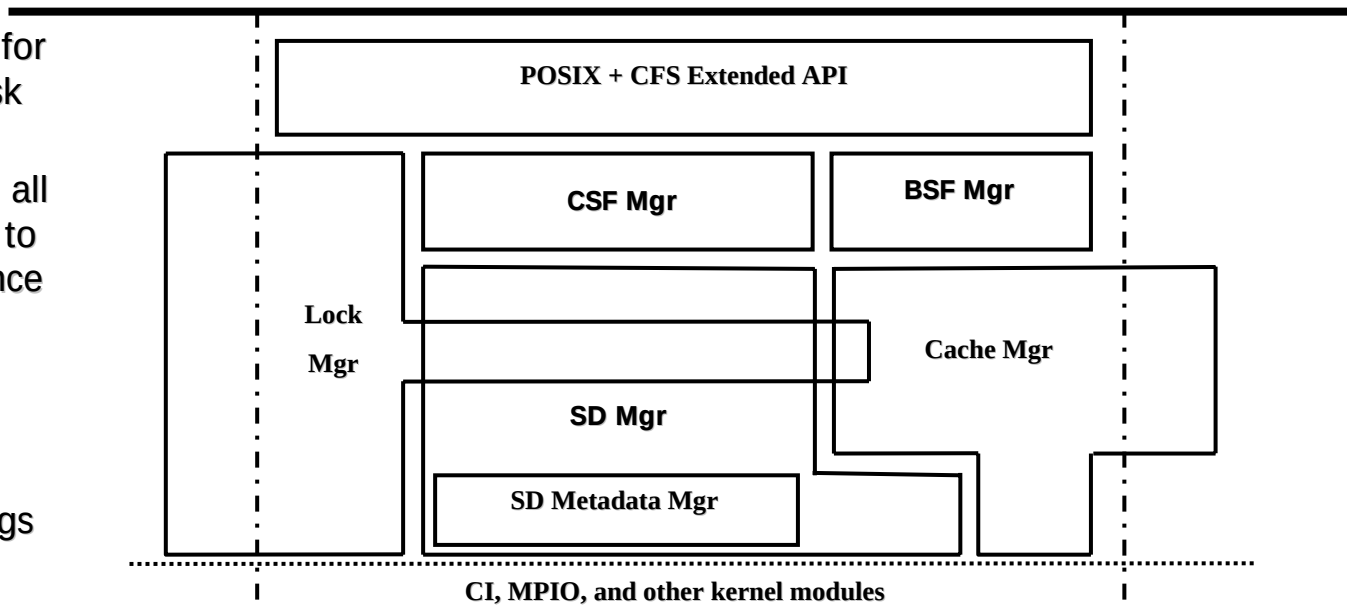

Cooperation with IRISA-Rennes
Christine Morin and KerLabs;
extension to wide area settings

Participation in CoreGRID

WP2 – Knowledge & DataMgt Task
since mid-2007

Paulo A. Lopes, Pedro D. Medeiros, Cooperative Caching in the pCFS Parallel Cluster File System, HPDC 2006

Paulo A. Lopes, Pedro D. Medeiros, pCFS: A Parallel Cluster File System, ParCo 2005

Lightweight communication mechanisms

Context: Reducing latency in communication between processes in distributed memory multiprocessors

Approach: Using dedicated hardware (eg Infiniband) and/or specialized libraries to reduce the overhead associated with the sending and receiving of messages; influence of communication performance in CPU load

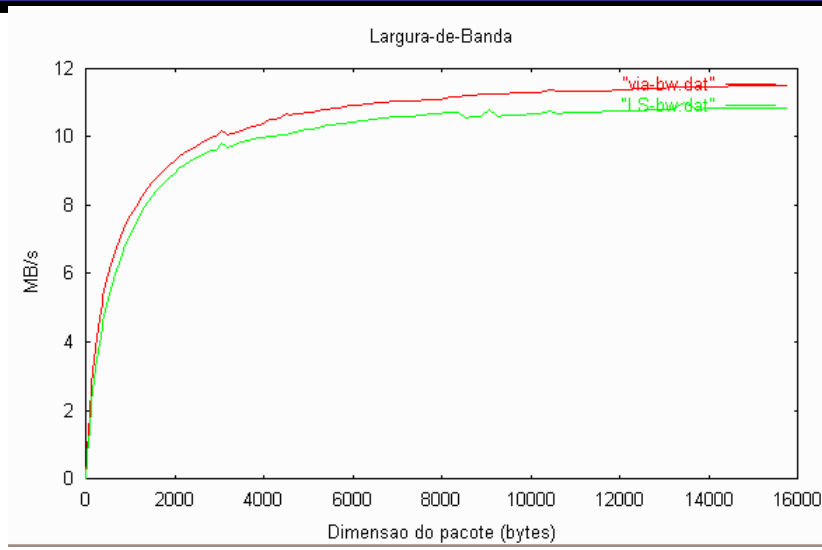

Some results published in two publications about pCFS

Results:

MSc thesis **1** (R. Espenica, The PVM System over the VIA architecture, 2003)

Participation in FCT-funded project TARDE (2002-04)

Background work: R. Espenica, P. Medeiros, Porting PVM to the VIA architecture using a fast communication library, in EuroPVM/MPI 2002

Results in more detail

PhD thesis ongoing: 1

Paulo Lopes, *A Shared-Disk Infrastructure for Parallel File Access*, FCT/UNL, thesis to be submitted March 2008, sup: Pedro Medeiros.

PhD thesis planned: 1

Nuno Oliveira, start as PhD Student 2008, sup: Pedro Medeiros.

MSc thesis completed: 1 [+2]

Roberto Espenica: *The PVM system over the VIA architecture*, 2003, sup: P. Medeiros.

Pedro Garrinhas: *A Framework for Interconnecting Mobile Devices for Parallel Processing Platforms,* 2007, sup: P. Medeiros.

Paulo Quaresma: *Parallelization of Tomographic Image Analysis*, 2007, sup: P. Medeiros.

MSc thesis ongoing: 5

Oliveiros Cristina, *Towards Integrating a Shared-Disk File System and a Cluster Operating System*, sup. Medeiros.

João Duarte, *Implementing Internet Protocols in a Dedicated Device under a Real-Time Operating System*, sup: P. Medeiros.

Roberto Moretti, *Monitoring Power Consumption in Large Data Centers*, sup: P. Medeiros.

Pedro Andréz, *Tomography Data Analysis in a Non Dedicated PC Cluster*, sup: P. Medeiros.

Tiago Cadavez, *Parallelization of Tomographic Data Analysis Using OpenMP,* sup: P. Medeiros.

Publications: 2

International Conferences (with peer review)

Paulo Lopes and Pedro Medeiros. Cooperative caching in the pcfs parallel cluster file system. In IEEE, editor, Proceedings of the 15th IEEE International Symposium on High Performance Distributed Computing, HPDC-15, pages 347-348. IEEE Computer Society, 06 2006.

Paulo Lopes and Pedro Medeiros. pcfs: A parallel cluster file system. In Proceedings of the ParCO 2005 Conference, Málaga, Spain, Sep 13-16 2005. Universidad de Málaga, 09 2005.

Projects

- Participation in National Funded Projects
  - InStory - Interactive Personalized Mobile Storytelling, 2004-2006, with IMG CITI Stream
    [Applications of Groups]
  - MIMO - Models and Infra-structures for Mobile Computing, 2002-2004, with Un. Porto
    [Mob Language]
  - TARDE – Tabulation and Revision in a Distributed Prolog Environment, 2002/04, with CENTRIA
    [Multi-Threaded XSB]
- Participation in International Network
  - Agentcities.NET, 2002-2003 [Pattern/Operators for Grid]
- International Projects
  - Software Environments for Grid Computing (Bilateral Cooperation UNL-U.Cardiff), 2003-...
    [PSEs and Grid: Patterns & Workflows]
  - Curriculum Development for High-Performance and Grid Computing (EU AsiaLink), 2005-2008
    with Delft Un. Tech., Tsinghua Un., and GUCAS (China)
    [Grid Computing] [Education & Research: Joint Curriculum & PhD Thesis co-advising]

- IBM SUR – Shared University Research Program for Parallel and Distributed Computing at CITI

  2003 - SUR Grant (Intel Xeon based)

  2007 - SUR Grant (Cell based, to be delivered 10th Jan 08) (jointly with CENTRIA)

Ongoing and New Projects

- International Networks: [2007-]
  - CoreGRID Network of Excelence – as Associate Member
    WP2 (Knowledge and Data Management) + WP3 (Programming) [Cluster & Grid]
  - CYTED-Grid Network (Spain, Portugal, and South America)
- Participation in National Funded Projects: [2008-]
  - CALLAS - Calculi and Languages for Sensor Networks, 2008-2010 with Un. Porto and U. Lisboa [Mob Language]
  - Byzantium - Eficient Byzantine fault-tolerant database replication, 2008-2010 with DSG CITI Stream [Transaction Memory][Monitoring]
  - Portugal-Spain Bilateral Cooperation UNL – UABarcelona, 2007-2008 [Performance and Monitoring Tools]
- Ongoing non-funded Collaboration Projects at FCT/UNL
  - Simulation of pollutants in shallow waters, with Environmental Sciences [Cluster & Grid Computing]
  - Tomography, with Material Sciences [Cluster & Grid Computing]
  - Dispersion of contaminants in sub-soil, with Geological Sciences [Cluster & Grid Computing]
  - Joint multidisciplinary Projects with other Research Centers of FCT/UNL

4 Post-doctorate positions at CITI (Portuguese FCT Program): [2008-2012]

2: BioInfo Project: Concurrency and Parallelism in Bioinformatics, with PLM CITI Stream **[Open]**

1: BioInfo Project: Grid Computing for Bioinformatics **[Open]**

1: GeoInfo Project: HPC and Grid Computing , namely with Geoscience Center of FCT/UNL **[Closed]**

Strategic partnerships

International:

- Distributed Collaborative Group of Cardiff University (UK) with Prof. Omer Rana. [PSEs and Grid Computing]
- IRISA-Rennes, France, with the Kerrighed Group led by Dr. Christine Morin [Parallel File Systems]
- Computer Architecture and Operating Systems Department, Universidade Autonoma de Barcelona, Spain, led by Prof. Emilio Luque [Parallel and Distributed Computing]
- IBM Research Labs at Haifa, Israel, with Dr. Shmuel Ur [Testing of concurrent programs]
- SUNY at Stony Brook, Profs Terrance Swift and David S. Warren [Concurrent Tabling and the XSB System]
- Delft University of Technology, NL, Profs Henk Sips and Hai-Xiang Lin [High Performance and Grid Computing]
- Tsinghua University, China, Prof. Weimin Zheng [High Performance and Grid Computing]
- Graduate School of Chinese Academy of Sciences, Prof Tiejan Luo [High Performance and Grid Computing]
- University of Texas at Austin, US, Profs. Keshav Pingali and Vijay Garg [Distributed and Grid Computing]

National:

- Other Research Centers FCT/UNL (Environmental Sc, Materials Sc, Geologic Sc) [PSE, Cluster & Grid Comp.]
- Un. Minho, Un. Coimbra and Un. Porto [High Performance and Grid Computing] [Mobile Computing]
- Un. Lisboa [Mobile Computing]

International Colaboratory for Emerging Technologies
CoLab Un. Texas at Austin – Portugal

Protocol funded by Portuguese Ministry of Science & Technology

Joint Research and Advanced Education Programs in Advanced Computing

PhD Program Computer Science and Informatics with focus on High Performance, Grid Computing and Computational Science and Engineering

involving 8 Portuguese universities and two research associated laboratories

José Cunha acted as **National Co-Director of CoLab** and the **National Director of the Advanced Computing Area**, from **Jan to Sept´07** launching CoLab coordinating office at FCT/UNL Campus.

Pedro Medeiros is **Co-Director of Advanced Computing Area** and **in the Coordination office of CoLab**

Conclusions

Global Scientific Results 2003-2006 [+2007]

Self Evaluation 2003-2007

Goals and Requirements for 2007-2010

Global Scientific Results: 2003-2006 [+2007]

- Publications:
  - Book edition: 1 [+0]
  - Journal Special Issue: 0 [+1]
  - Proceedings edition: 2 [+0]
  - Internat. Journals: 3 [+2]
  - Book Chapters: 3 [+0]
  - Internat. Conferences: 15 [+6]

76% with foreign co-authors

- Participation in Projects:
  - Current: 2 national; 1 international +2 networks + 2 bilateral coops.
    3 multidisciplinary collaborations, local to FCT / UNL
    2 (multidisciplinary) Post-doctorate FCT Program
    1 IBM SUR Grant (2007 with CENTRIA))
  - Completed: 3 national; 1 network; 1 IBM SUR Grant (2003)
- Dissertations:
  - PhD: 4 completed [+2 in 2007] +1 awaiting defense (+2 ongoing)
  - MSc: 1 completed [+3 in 2007] (+15 ongoing)
- International Scientific Events:
  - Organisation: 2 international conferences on site
  - Program committees: > 60

Self Evaluation: 2003-2007

- Positive:
  - Relevance of core and emerging themes
  - Completed PhDs
  - Scientific events
  - International cooperations
  - Multidisciplinary collaborations
- To improve:
  - Publications
  - International projects
  - More regular flow of PhD students

Goals and requirements for 2007-2010

Strategic Goals:

- To consolidate the developing and emerging research themes
- To sustain a regular flow of post-graduate students and further attract post-doctorate researchers
- Increase participation in international projects and consolidate international cooperations
- **Develop High-Performance, Parallel and Distributed Computing for Computational Science and Engineering in domains such as:**
  - Materials Science
  - Geological Science
  - BioInformatics

Special requirements:

- **PhD Student Scholarships**
- **Equipment / Resources**

  A shared memory multiprocessor with at least 16 nodes (e.g., four quad-cores) for computational support, namely for the Software Transactional Memory and for Parallel Tabling in XSB research themes