

# MAP-I doctoral program

## Distributed Systems: from principles to the Grid

### Summary

This document presents a Ph.D. level course, as a joint FCUP-DIUM proposal for the MAP-I doctoral program, corresponding to a Curriculum Unit - UCPC - credited with 5 ECTS.

### I. Theme, Justification and Context

This course is intended to Ph.D. level students whose main theme of research is related to parallel and distributed systems in a broad context. As such, the course will focus on communication, naming, synchronisation, data consistency and replication, fault tolerance and security, with a walk through main distributed algorithms found in the literature being them synchronous or asynchronous. The course is structured in 3 main parts. The first part is dedicated to basic concepts of parallel and distributed computing focusing on parallel programming models, parallel architectures, including modern multi-cores and multi-processors, performance metrics, performance analysis, parallel benchmarks, and state-of-the-art in parallel and distributed systems. The second part focuses on distributed algorithms for synchronous and asynchronous systems. In the last part of the course we will study grid computing as an instance of an asynchronous distributed system, and discuss about state-of-the-art solutions and their limitations.

The relevance of the topics related to distributed systems and algorithms is well demonstrated by the number of articles published every year in main stream scientific events and periodicals with high impact rate, such as the Journal of Parallel and Distributed Computing (JPDC), IEEE Transactions on Parallel and Distributed Systems (TPDS), Journal of Grid Computing (JGC), the Journal of High Performance Applications (JHPA), the IEEE International Conference on Parallel and Distributed Systems (ICPADS), the International Parallel and Distributed Processing Symposium (IPDPS), the Grid Computing Conference (GC), the Conference on Cluster and Grid Computing (CCGRID), the Conference on Middleware for Grid Computing (MGC), among many others.

Regarding the first part of the course, the students will have the opportunity to review basic concepts that are relevant to understand the motivation for the problems and solutions presented in parts 2 and 3. The ability to understand and develop algorithms that take advantage of the parallel and distributed architectures that are currently available is a mandatory skill for next generation Ph.D. students.

The second part of the course will cover classical problems in parallel and distributed environments related to communication, naming, synchronisation, data consistency and replication, fault tolerance and security. We will complement the discussion of these problems with a walk through main distributed algorithms found in the literature being them synchronous or asynchronous. For example, we will study scalable synchronisation algorithms, resource allocation, workload management, fault tolerance and data consistency and replication, among others.

The third part of the course will use a case study to demonstrate where to use these algorithms in the context of a grid environment. Most software written to support grid computing and data grids,

do not rely or only partially apply the rich background of distributed systems and algorithms built over the past decades. Our objective in this part of the course is to study the existing solutions and indicate alternative ways of solving grid problems on the light of the distributed systems and algorithms studied in the previous parts.

Experience has shown that the technologies, algorithms and techniques involved in sharing resources across the Internet through grid computing need to cope with extra requirements not foreseen by the distributed algorithms and systems studied, because of geographical distances, security (authentication, authorization), confidentiality, link capacities, and scalability. Our aim in this course is to understand how to extend the classical algorithms to cope with this new paradigm, and to be able to develop new parallel and distributed algorithms.

Grid protocols and technologies are being adopted in a wide variety of academic, government, and industrial environments, and there is a growing body of research-oriented literature in grid computing. As such, there is a need for educational resources on this field. The next generation of scientists and engineers need to be prepared for a technological workplace that is changing the world.

Several courses on distributed algorithms and systems are available throughout the world at graduate level. In fact, the books used as main references for this proposed course are also used in other graduate courses at M.Sc. and Ph.D. levels, depending on the depth level of the topics. We believe that our course has a relevant contribution as it will try to apply and extend classical parallel and distributed algorithms to the context of grid computing, that poses new challenges related to geographical distances, link capacities, authentication and authorization, confidentiality and scalability.

The students attending this course must have a strong background on foundations of Computer Science, specially on Operating Systems, Parallel and Distributed programming and Algorithms and Data Structures.

## **II. Objectives and Learning Outcomes**

By the end of the course students should be able to understand the basic concepts of parallel and distributed algorithms and systems. They should be able to understand a parallel and distributed problem and be able to develop their own parallel and distributed algorithms. They also should be able to criticize peers works and develop their own independent research work.

Students should leave the course with:

- A good understanding of the principles and foundations of parallel and distributed architectures and programming models.
- An appreciation of the various parallel and distributed algorithms found in the literature.
- Confidence in using several of the available algorithms.
- An awareness of good strategies for solving parallel and distributed problems.
- Identify and characterize several types of parallelism and architectures that support them.
- Characterize and develop scalable algorithms, selecting the best parallelization for a particular parallel system.
- Design, implement and optimize concurrent, parallel, distributed and grid enabled applications
- An introduction to the research challenges and the future of Grid computing.
- Awareness of current research issues in: parallel and distributed algorithms and their application to grid systems.

- Insight into the architectural implications of Grid-scale computation.

In addition students will also gain proficiency into the following areas:

- Security: administrators, developers, and grid users all need to understand the use of public key encryption, certificates, and how security is implemented in the Grid.
- Grid Access and Management – administrators and grid users all need to know how to use commands and tools to submit and execute different grid applications, how to locate, access, and manage data repositories and other grid resources, and how to interface to advanced grid computing toolkits.

## IV. Course Contents

### Part I. Parallel and Distributed Systems

#### Review of Parallel and Distributed Architectures

##### Programming Paradigms

- Shared memory
- Message passing
- Distributed shared memory

##### Development of parallel and distributed applications

- Design phases
- Common parallel patterns
- Performance metrics and profiling
- Optimizations techniques
- Mapping and scheduling
- State-of-the-art on parallel and distributed systems and applications

### Part II. Distributed Systems and Algorithms

#### Distributed Systems

- Introduction
- Communication
- Naming
- Synchronisation
- Consistency and Replication
- Fault Tolerance
- Security
- Web-based systems
- Grid-based systems

#### Distributed Algorithms

- Introduction
- Synchronous Network Model
- Synchronous Network Algorithms
  - Leader election
  - Algorithms in general synchronous networks
  - Distributed Consensus with link failures
  - Distributed Consensus with Process Failures
  - The Commit problem
- Asynchronous Algorithms
  - Asynchronous system model
  - Asynchronous shared memory algorithms
    - Asynchronous shared memory model
    - Mutual exclusion
    - Resource Allocation
    - Consensus
    - Atomic Objects
  - Asynchronous Network Algorithms

- Asynchronous Network Model
- Basic Asynchronous Network Algorithms
- Synchronisers
- Shared memory versus Networks
- Logical time
- Global snapshots and stable properties
- Network resource allocation
- Asynchronous Networks with Process failures
- Data Link protocols
- Partially Synchronous algorithms
  - Partially Synchronous system models
  - Mutual exclusion with partially synchrony
  - Consensus with partial synchrony

### **Part III. Case study: Grid Computing**

#### **Overview of Grid Computing (OGSA, OGSI)**

##### **Grid Computing Components**

- Authentication and Authorization (ssl, pk12, x509)
- Resource Discovery and Info Services (GRAM, ldap, BDII)
- Workload Management System (WMS)
- Middleware and Application (globus, Gridbus Broker, OSG, Unicore, gLite etc)
- Virtual Organization management (VOMS)
- Data Access Integration and Management (R-GMA, LFC, AMGA, GFAL)
- Web versus Grid Services

## **V. Methodology**

The overall aim is to give students a broad and well balanced understanding of distributed systems and algorithms that will serve well as a foundation for more specific work or research.

The Course will include lectures and discussions on the principles, technologies, experience and exploitation of distributed systems and algorithms. Students are expected to deliver small assignments along the course, and give short presentations accompanied of a written document. The methodology will also favour critical discussion and reasoning about large-scale distributed system architectures, infrastructures and technologies.

We will use 2 textbooks in this course. The first one, Distributed Systems; Principles and Paradigms, by Andrew Tanenbaum and Maarten van Steen, will be used to cover aspects of the second part of the course. The second one, Distributed Algorithms, by Nancy Lynch, will be used to cover classical synchronous and asynchronous distributed algorithms. As a complement, students will read and discuss papers that will be assigned according to the lecture topic.

Proposed schedule (assuming 14 classes of 3 hours each)

#### **Part 1:**

- Aula 1: Introduction  
(scope of the course, methodology, resources, grading, bibliography, intro)

- Aula 2: Review on Parallel and Distributed Programming and Architectures (suggested papers to be discussed next class)

**Part 2: Distributed Systems**

- Aula 3: **Students presentations and discussion**  
Overview of Distributed Systems
- Aula 4: Communication, Naming
- Aula 5: Processes (migration, checkpointing)
- Aula 6: Synchronization  
(suggested papers to be presented)
- Aula 7: **Students presentations and discussion**  
Consistency and replication, fault-tolerance
- Aula 8: Security, file systems
- Aula 9: Web-based systems

**Part 2: Distributed Algorithms**

- Aula 10: Synchronous, Asynchronous and Partially Synchronous algorithms
- Aula 11: Shared memory versus message passing algorithms  
(suggested papers to be presented)

**Part 3: Case Study**

- Aula 12: **Students presentations and discussion**  
Grid architecture
- Aula 13: Grid computing components
- Aula 14: Problems and limitations

## VI. Grading

Students will be evaluated through assignments, short presentations, and paper discussion.

## VII. Bibliographic References

1. **Distributed Algorithms**, by Nancy Lynch, **The Morgan Kaufmann Series in Data Management Systems**.
2. **Distributed Systems: Principles and Paradigms**, by Andrew S. Tanenbaum and Maarten van Steen, Prentice Hall.
3. **List of suggested papers:**
  - o Leslie Lamport, Time, Clocks, and the Ordering of Events in a Distributed System, CACM, 21(7), July, 1978.
  - o Leslie Lamport, On Interprocess Communication. Part I: Basic Formalism. Distributed Computing 1(2): 77-85 (1986)
  - o Leslie Lamport, On Interprocess Communication. Part II: Algorithms. Distributed Computing 1(2): 86-101 (1986)
  - o Leslie Lamport: The mutual exclusion problem: part I - a theory of interprocess communication. J. ACM 33(2): 313-326 (1986)
  - o Leslie Lamport: The mutual exclusion problem: partII - statement and solutions. J. ACM 33(2): 327-348 (1986)

- o J. Mellor-Crummey and M. L. Scott, Algorithms for Scalable Synchronisation on Shared-Memory Multiprocessors. ACM, TOCS, V. 9, n. 1, pp. 21-65, 1991
- o Agustin Caminero, Anthony Sulistio, Blanca Caminero, Carmen Carrion, and Rajkumar Buyya, Extending GridSim with an Architecture for Failure Detection, ICPADS, 2007
- o T. Srinivasan, R. Balakrishnan, S.A. Gangadharan, and V. Hayawardh, A Scalable Parallelization of All-Pairs Shortest Path Algorithm for a High Performance Cluster Environment, ICPADS, 2007
- o Olivier Beaumont, Larry Carter, Jeanne Ferrante, Arnaud Legrand, Loris Marchal, Yves Robert, Centralized versus Distributed Schedulers for Bag-of-Tasks Applications, IEEE Transactions on Parallel and Distributed Systems, pp. 698-709, May, 2008.

### Annex: Team

This course is supported by a team involving researchers from the University of Porto, FCUP (Inês de Castro Dutra), and from the University of Minho, DI-CCTC (António Manuel da Silva Pina). In what follows we give a brief presentation of each researcher, which includes, for each of them, up to 5 key publications related to the scientific area in which this course is proposed.

**Inês de Castro Dutra:** is a lecturer at Science Faculty of Porto University, FCUP. She obtained her B.Sc. degree in Computer Science from State University of Rio de Janeiro, in 1985, and her M.Sc. degree in the Systems Engineering and Computer Science department of Federal University of Rio de Janeiro, in 1988. She obtained her Ph.D. degree from Bristol University, in 1995. In 1998, she was a lecturer in the Department of Systems Engineering and Computer Science of COPPE, an institution for postgraduate studies in Engineering, at Federal University of Rio de Janeiro, where she taught courses on Operating Systems, Concurrent Programming and Topics on High Performance Computing, at M.Sc. and Ph.D. levels and Artificial Intelligence and Logic Programming, at undergraduate level. During the periods between October 2001 and December 2002, and between April 2004 and March 2005, she worked as a visiting researcher at University of Wisconsin-Madison, USA, in the department of Biostatistics and Medical Informatics. During these periods, she worked for machine learning projects funded by NSF, DARPA and American Air Force (projects COLLEAGUE, EELD and EAGLE) and started to work with applications that demanded a huge amount of resources. At this time, she had the opportunity to work with the Condor team, and to largely use the Condor resource manager to run her experiments. Having some experience on running and managing thousands of experiments in a grid environment, she supervised a D.Sc. student that worked on a new middleware for managing applications, an M.Sc. student that worked on a tool to manage the submission and monitoring of machine learning applications, and another M.Sc. student that worked on the development of a scheduling algorithm based on reinforcement learning for grid environments. In the past she supervised students that worked on parallel logic programming systems. She was invited for committees of several workshops and conferences., She has published over 40 scientific articles in conferences and journals. Currently she is a member of the EELA-2 (E-science grid facility for Europe and Latin America) project, whose main objective is to promote and maintain the infrastructure of hardware and software between Europe and Latin America produced during the first phase of EELA, to form a human research network and deploy applications on e-Science.

#### Supervision:

Currently, she is supervising three M.Sc. students in the area of grid computing. One of them is proposing and implementing a grid interface to access data from geographically distant databases. Another one is working on the parallelisation of algorithms for the classification of capsule endoscopic images. The third one is working on supporting scheduling of MPI tasks across domains using several different scheduling strategies over the Samba-Grid system developed by Plastino et al. at Federal University Fluminense, Brazil.

- Ph.D. students:
  - Rogerio Lopes Salvini, *Obtaining Efficient Classifiers from Induced Rules*, [Department of Systems Engineering and Computer Science, COPPE, Federal University of Rio de Janeiro](#), 2008;
  - Luciana Itida Ferrari: *Analysis of Immunoblots*, [Department of Systems Engineering and Computer Science, COPPE, Federal University of Rio de Janeiro](#), 2009;
  - Laci Mary Barbosa Manhaes, *Multi-Classification in Inductive Logic Programming*, [Department of Systems Engineering and Computer Science, COPPE, Federal University of Rio de Janeiro](#), 2009;
  - Alberto Arkader Kopiler, *Managing infrastructures using systems inspired on the nature*, [Department of Systems Engineering and Computer Science, COPPE, Federal University of Rio de Janeiro](#), 2009.
  
- M.Sc. students:
  - Andre Oliveira: meta-scheduling of MPI processes, over Samba-Grid, [Department of Systems Engineering and Computer Science, COPPE, Federal University of Rio de Janeiro](#), 2008;
  - Joao Victor de Almeida Pap: Grid querying, [Department of Systems Engineering and Computer Science, COPPE, Federal University of Rio de Janeiro](#), 2008.
  - Filipe Valpereiro: Parallelisation of algorithms for classification of endoscopic capsule images, [Department of Computer Science, Porto University](#), 2008.

#### Research Projects:

Currently she is working in the EELA-2 project (E-science grid facility for Europe and Latin America).

#### Selected Publications:

- P. K. Vargas, I. C. Dutra, V. D. Nascimento, L. A. S. Santos, L. C. Silva, C. F. R. Geyer and B. Schulze, GRAND: Toward Scalability in a Grid Environment, *Concurrency and Computation: Practice and Experience*, November, 2006, Wiley InterScience, doi: 10.1002/cpe.1138;
- J. Davis, D. Page, E. Burnside, I. C. Dutra, R. Ramakrishnam, V. Santos Costa, and J. Shavlik, View Learning for Statistical Relational Learning: with an Application to Mammography, *International Joint Conference on Artificial Intelligence*, Edinburgh, Scotland, July, 2005;
- J. A. L. Sanches, P. K. Vargas, I. C. Dutra, V. Santos Costa, and C. F. R. Geyer, ReGS: user-level Reliability in a Grid Environment, *IEEE International Symposium on Cluster Computing and the Grid*, May 9--12, 2005, Cardiff, UK;
- P. K. Vargas, L. A. S. Santos, C. F. R. Geyer, I. C. Dutra, An Implementation of the GRAND Hierarchical Application Management Model using the ISAM/EXEHDA system, *III Workshop on Computational Grid and Applications*, LNCC, Petrópolis, RJ, January 31st -- February 2nd, 2005;
- I. C. Dutra, D. Page, V. Santos Costa, J. Shavlik and M. Waddell, Toward Automatic Management of Embarrassingly Parallel Applications, *Euro-Par'03*, Klagenfurt, Austria, August 26 - 29, pp. 509--516, 2003.

**António Manuel Silva Pina:** is a lecturer at the Department of Computer Science of University of Minho, and a researcher member of CCTC. His scientific research activities are centred on the areas of Parallel Computing modelling and programming, cluster and grids platforms development and deployment.

### Supervision

- Ph.D. students:
  - o Albano Alves, Rocmeu: orientação ao recurso na modelação de aplicações paralelas e exploração cooperativa de clusters multi-SAN, Eng /G/GC-6300, Universidade do Minho, Dezembro 2004.
  - o José Carlos Rufino, Domus: Tabelas de Hash Distribuídas em Clusters de Nós de Computação Heterogéneos, Universidade do Minho, Abril 2008;
  - o José Luís Exposto, Partição multi-objectivo para descarga eficiente de recursos na WWW, Universidade do Minho, (waiting public defense);
  - o Vítor Oliveira, Civil Protection applications in a Grid supported environment, Thesis proposal and planning, MAP-I 2008.
- MSc students
  - o Albano Serrano, Infra-estrutura para a Computação Multi-cluster em Ambiente Grid, Universidade do Porto, Jan. 2007
  - o Cecília Moreira, CoRes – Computação orientada ao Recurso – uma Especificação, Universidade do Minho 2001;
  - o António Tavares, Realização de um Modelo de Computação Baseado em Agentes, Universidade do Minho, 1997.

### Research Projects:

- o EELA-2, E-science grid facility for Europe and Latin America: Deployment of e-Infrastructures for scientific communities, INFRA-2007-1.2.3: e-Science Grid infrastructure, JRU Portugal, Principal Investigator of UMinho, 2008;
- o CYCLOPS Cyber-Infrastructure for Civil protection Operative Procedures, 6 FP, Research Infrastructure-Communication Network Development, proposal nº 031874, Principal Investigator of UMinho, 2007;
- o AspectGrid: Pluggable Grid Aspects for Scientific Applications, member of the research team, GRID/GRI/81880/2006;
- o CROSS-Fire-Collaborative Resources Online to Support Simulations on Forest Fires : a Grid Platform to Integrate Geo-referenced Web Services for Real-Time Management, member of the research team, GRID/GRI/81795/2006;
- o IGIDE-Interactive Global illumination within Dynamic Environments, member of the research team, PTDC/EIA/65965/2006;
- o SeARCH: Serviços e Investigação em Computação Avançada com Clusters HTC/HPC, member of the research team, Refª: CONC-REEQ/443/2001;
- o “SIRe: A Scalable Information Retrieval environment”, Principal Investigator, POSI/2001/CHS/41739.

### Selected Publications (cluster and Grid computing):

- o A. Pina, B. Oliveira, A. Serrano, V. Oliveira, “[EGEE Site Deployment & Management Using theRocks toolkit](#)”, 2<sup>nd</sup> Iberian Grid Infrastructure Conference, Porto, May,2008;
- o A. Alves, A. Pina, J. Rufino, J. Exposto, “[Deploying Applications in Multi-SAN SMP Clusters](#)”, full version, Special Issue of International Journal of Computational Science and

Engineering (LJCSE). ISSN (print):1742-7185, ISSN (on-line):1742-7193. InderScience Publishers;

- o J. Rufino, A. Pina, A. Alves, J. Exposto, [pDomus: a prototype for Cluster-oriented Distributed Hash Tables](#), 15th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP 2007), Naples, Italy, February 2007;
- o J. Exposto, J. Macedo, A. Pina, A. Alves, J. Rufino, "Geographical Partition for Distributed Web Crawling", 2nd International Workshop on Geographic Information Retrieval (GIR 2005), ACM Press, pp. 55-60, Bremen, 2005;
- o Alves, A. Pina, "Bridging the gap between cluster and grid computing", 6th International Conference on Parallel Processing and Applied Mathematics (PPAM 2005), LNCS 3911, Springer, Poznan, 2005;
- o Alves, A. Pina, J. Rufino, J. Exposto, "[meu: unifying application modeling and cluster exploitation](#)", 16th Symposium on Computer Architecture and High Performance Computing (SBAC-PAD 2004), IEEE Computer Society, 132-139, Foz do Iguacu, Brazil, 2004.

#### **Selected Technical / Scientific Activities** (cluster and grid)

- o Member of the Scientific Committee of the "2st IBERIAN INFRASTRUCTURE FOR DISTRIBUTED COMPUTING CONFERENCE, Porto, May 2008;
- o Member of the Scientific Committee of the "1st IBERIAN INFRASTRUCTURE FOR DISTRIBUTED COMPUTING CONFERENCE, Santiago de Compostela, Spain, May 2007;
- o Co-supervision technique, scientific and administration of the communication, /computation, and storage of SeARCH, Universidade Minho, since 2005;
- o Review of scientific papers submitted to the "11th International Euro-Par 2005, Parallel Processing, Lisboa 2005.

**Fernando Manuel Augusto da Silva:** is an associate professor at the Computer Science Department of the Faculty of Science of the University of Porto and currently is the coordinator of the new research unit CRACS (Center for Research in Advanced Computing Systems). He holds a Ph.D. In Computer Science from the University of Manchester (1993) and the Habilitation in Informatics from the Universidade Nova de Lisboa (2007). He has been teaching for more than 20 years, both undergraduate and postgraduate courses on Programming, Operating Systems, Data Structures and Algorithms, Databases and Parallel and Distributed Computing. His scientific research activities are centred in the areas of Parallel and Distributed Computing, Programming Languages and Information Retrieval.

#### **Students Supervision**

Ph.D.:

- (Ongoing) Pedro Ribeiro (PhD thesis on Parallel Algorithms for Motif Finding in Brain Networks).
- (Ongoing) Rolando da Silva Martins, (PhD thesis on Real Time P2P Middleware).
- (12/06) Hervé Paulino, An infrastructure for mobile service-oriented computing encoded e on a process calculus. Advisor: Luís Lopes. Currently: Assistant Professor at CS Department at UNL in Lisbon.
- (10/06) Nuno A. Fonseca, Parallelism in Inductive Logic Programming, Co-advised by R. Camacho. Currently a Post-Doc at IBMC-Porto.
- (6/04) Alvaro Reis Figueira, Design and Implementation of a Distributed System with Mobility Based on a Process Calculus. Co-advised by Luís Lopes. Currently: Assistant Professor at FCUP.
- (12/01) Ricardo Nuno Lopes, An Implementation of the Extended Andorra Model. Advisor: V. Santos Costa.

- (11/01) Ricardo Jorge Rocha, On Applying Or-Parallelism and Tabling to Logic Programs, Co-advised by V. Santos Costa. Currently: Assistant Professor at FCUP.
- (11/01) Manuel Eduardo Correia, On the Implementation of And/Or Parallel Logic Programming Systems, Co-advised by V. Santos Costa. Currently: Assistant Professor at FCUP.
- (12/99) Luís Lopes, On the Design and Implementation of a Virtual Machine for Process Calculi. Co-advised by Vasco Vasconcelos. Currently: Associated Professor at FCUP.

#### Msc:

- (Ongoing) Brás Monteiro (MSc thesis on Parallel Prolog on Multicores).
- (4/05) João Paulo Magalhães, A Distributed File System for a P2P Architecture, Co-advised by Luís Lopes.  
(7/03) Rolando Martins, A Distributed Execution of Prolog based on Stack Splitting, Co-advised by Ricardo Rocha.
- (6/03) Licínio Oliveira, P3 : Parallel Peer-to-Peer. Co-advised by Luís Lopes.
- (3/03) Rogério Ferreira, A Methodology for data Sharing Among Heterogeneous Systems using XML, Advisor: José Paulo Leal.
- (1/02) Nuno Manuel Reis Amado, Parallel Algorithms to Induce Decision Trees, Co-advised by João Gama.
- (10/01) Mário João Gonçalves Antunes, Centralized Administration of Groups of Unix Systems, Co-advised by Manuel Eduardo Correia.
- (10/98) Hervé Miguel Cordeiro Paulino, Design and Implementation of the pSystem for Distributed Memory Architectures.
- (10/96) Ricardo Jorge Rocha, An Or-Parallel Prolog System based on Environment Copying.
- (1/95) Luís Miguel Barros Lopes, A Parallel Programming Environment to Evaluate Scheduling Heuristics.

#### Selected Research Projects:

- 1/96 - 1/98, Principal Investigator in Design and Implementation of a low cost scalable parallel machine, with Prof. Paul Watson, University of Newcastle, funded by JNICT and British Council.
- 1/96 - 1/99, Team member in MELODIA: Design and Implementation of Advanced Models for the Execution of Prolog Programs, funded by JNICT (PBIC/C/TIT/2495/95).
- 3/97 - 12/00, Principal Investigator in Dolphin: Parallel Computing Environments for High Performance Systems, funded by the Portuguese Foundation for Science and Technology (FCT) (contract PRAXIS/2/2.1/TIT/1577/95), 75,000 Euros.  
11/99 - 7/01, LIACC Team coordinator in DICOMO: Distribution, Concurrency, and Mobility, funded by FCT (contract PRAXIS/P/EEI/12059/98), 26,500 Euros.
- 10/01 - 10/04, Team member in MIMO - Models and Infra-structures for Mobile Computing, funded by FCT (contract POSI/CHS/39789/2001), 35,000 Euros.
- 5/02 - 10/05, Principal Investigator in APRIL: Applications using Parallelism in Induction and Logic, funded by FCT (contract POSI/SRI/40749/2001), 66,000 Euros.
- 10/07 - 10/10, Team member in STAMPA - Sophisticated Tabling Mechanisms for Prolog and their Applications, funded by FCT (contract PTDC/EIA/67738/2006), 150,000 Euros.

#### Selected Publications (parallelism):

Fernando Silva, Gaspar Barreira, and Lúcia Ribeiro (editors). IBERGRID'2008, Proceedings of the 2nd Iberian Grid Infrastructure Conference, Porto, Portugal. University of Porto, May 2008.

- Ricardo Rocha, Fernando Silva, and Vítor Santos Costa. On Applying Or-Parallelism and Tabling to Logic Programs. *Journal of Theory and Practice of Logic Programming*, 5(1&2):161-205, January 2005.
- Nuno Fonseca, Fernando Silva, Vítor Santos Costa, and Rui Camacho. A Pipelined Data Parallel Algorithm for ILP. In *In the Proceedings of 2005 IEEE International Conference on Cluster Computing*, Boston, USA. IEEE, September 2005.
- Nuno Fonseca, Fernando Silva, and Rui Camacho. Strategies to Parallelize ILP Systems. In *Proceedings of the 15th International Conference on Inductive Logic Programming (ILP 2005)*, Bonn, Germany, Springer-Verlag, LNCS 3625, pages 136-153. August 2005.
- R. Rocha, F. Silva, and V. Santos Costa. Speculative Computations in Or-Parallel Tabled Logic Programs. In *Proceedings of the 20<sup>th</sup> International Conference on Logic Programming, ICLP'2004*, Saint-Malo, France, Springer-Verlag, LNCS 3132, pages 254-268, September 2004.
- Ricardo Rocha, Fernando Silva, and Rolando Martins. YapDSS: an Or-Parallel Prolog System for Scalable Beowulf Clusters. In *Proceedings of the 11th Portuguese Conference on Artificial Intelligence, EPIA 2003*, Beja, Portugal, Springer-Verlag, LNAI 2902, pages 136-150, December 2003.
- Luís Lopes, Vasco T. Vasconcelos, and Fernando Silva. Fine-Grained Multithreading with Process Calculi. *IEEE Transactions on Computers*, Special Issue on the Parallel Architectures and Compilation Techniques Conference, 50(8):852-862, 2001.
- Fernando Silva and Paul Watson. Or-Parallel Prolog on a Distributed Memory Architecture. *Journal of Logic Programming*, 43(2):173-186, May 2000.
- Luís Lopes and Fernando Silva. Thread- and Process-Based Implementations of the pSystem Parallel Programming Environment. *Software Practice and Experience*, 27(3):329-351, March 1997.