

Europass Curriculum Vitae



Personal information

Surname(s) / First name(s)

Pereira Vilaça, Ricardo Manuel

Address(es)

Rua Victor de Sá, nº 15, 3º direito poente, Fraião, 4715-586 Braga, Portugal

Telephone(s)

936797028

Email(s)

rmpvilaca@gmail.com

Nationality(-ies)

Portuguese

Date of birth

14/02/1982

Gender

Male

Education and Training

2007-2012

Ph.D

School

Universidade do Minho

Professional Area

MAP-i Doctoral Programme. MAP-I is a joint venture of Universidade do Minho, Universidade de Aveiro and Universidade do Porto (MAP), three top Portuguese Universities

Thesis

Large-Scale Decentralized Object Store

Supervisor

Rui Oliveira

Brief description

We envision a large scale, dependable, fully decentralized object store. The target is a large computing network composed of tens to hundreds of non dedicated nodes connected through typical local area links. Over this network, we intend to design and prototype a general purpose peer-to-peer object store that seamlessly fragments and replicates application objects aiming at providing a dependable massive storage with in-place processing capabilities. The system should be capable to leverage replication to balance read scalability and fault tolerance.

2000-2005

Computer Science and Systems Engineering

School

Universidade do Minho – Department of Informatics

Professional Area

Specialization in Information Technologies

Student Number

33223

Average Grade

17

1997-2000

High School

School

Escola Secundária de Maximinos

Área

Agrupamento 1- Científico Natural – Geral

Teaching Experience

- 2011-2012 Cosupervisor of master thesis of Luís Pedro Zamith de Passos Machado Ferreira entitled "Bridging the Gap Between SQL and NoSQL".
- 2010-2011 Professor responsible of classroom Operating Systems and Distributed systems in CET de Redes at Instituto Politécnico do Cávado e do Ave (IPCA).
- 2009-2010 Professor of classroom practice of EISD, Distributed Systems, Universidade do Minho
- 2008-2009 Professor of classroom practice of EISD, Distributed Systems, Universidade do Minho
- 2009 Supervisor of a junior research grant, Bolsa de Integração à Investigação, with title "Benchmarking sistemas de armazenamento baseados em objetos"

Research Interests

Distributed Systems, Database Replication, Cloud Computing Storage, Peer-to-Peer

Publications

- 2013 Francisco Maia, Miguel Matos, Ricardo Vilaça, José Pereira, Rui Oliveira, and Etienne Riviere. DATAFLASKS: an epidemic dependable key-value substrate (to appear). In *International Workshop on Dependability of Clouds, Data Centers and Virtual Computing Environments (DCDV 2013)*, June 2013
- 2013 Ricardo Vilaça, Francisco Cruz, José Pereira, and Rui Oliveira. An effective scalable sql engine for nosql databases (to appear). In *Proceedings of the 13th IFIP WG 6.1 international conference on Distributed applications and interoperable systems*, DAIS'13. Springer-Verlag, 2013
- 2013 Leander Bessa Beernaert, Miguel Matos, Pedro Gomes, Ricardo Vilaça, and Rui Oliveira. Evaluating cassandra as a manager of large file sets. In *Third Workshop on Cloud Data and Platforms (CloudDP 13)*, Prague, Czech Republic, 2013
- 2013 Francisco Cruz, Francisco Maia, Miguel Matos, Rui Oliveira, Joao Paulo, José Pereira, and Ricardo Vilaça. MeT: Workload aware elasticity for NoSQL. In *EuroSys*, 2013
- 2012 Leander Bessa Beernaert, Miguel Matos, Ricardo Vilaça, and Rui Oliveira. Automatic elasticity in openstack. In *Workshop on Secure and Dependable Middleware for Cloud Monitoring and Management (SDMCM)*, Montreal, Quebec, Canada, 2012
- 2012 Ricardo Jimenez-Peris, Marta Patiño-Martínez, Iván Brondino, José Pereira, Rui Oliveira, Ricardo Vilaça, Bettina Kemme, and Yousuf Ahmad. Cumulonimbo: Parallel-distributed transactional processing. In *CloudFutures 2012 Workshop*, Berkeley, California, USA, May 2012
- 2011 Ricardo Vilaça, Rui Oliveira, and José Pereira. A correlation-aware data placement strategy for key-value stores. In *Proceedings of the 11th IFIP WG 6.1 international conference on Distributed applications and interoperable systems*, DAIS'11, pages 214–227, Berlin, Heidelberg, 2011. Springer-Verlag
- 2011 Miguel Matos, Ricardo Vilaca, José Pereira, and Rui Oliveira. An epidemic approach to dependable key-value substrates. In *International Workshop on Dependability of Clouds, Data Centers and Virtual Computing Environments (DCDV 2011)*, June 2011

- 2010 Ricardo Vilaça, Francisco Cruz, and Rui Oliveira. On the expressiveness and trade-offs of large scale tuple stores. In Robert Meersman, Tharam Dillon, and Pilar Herrero, editors, *On the Move to Meaningful Internet Systems, OTM 2010*, volume 6427 of *Lecture Notes in Computer Science*, pages 727–744. Springer Berlin / Heidelberg, 2010
- 2009 Ricardo Vilaca, José Orlando Pereira, Rui Carlos Oliveira, José Enrique Armendariz-Inigo, and José Ramón de Mendivil. On the cost of database clusters reconfiguration. In *SRDS '09: Proceedings of the 2009 28th IEEE International Symposium on Reliable Distributed Systems*, pages 259–267, Washington, DC, USA, September 2009. IEEE Computer Society
- 2009 Ricardo Vilaça and Rui Oliveira. Clouder: a flexible large scale decentralized object store: architecture overview. In *WDDDM '09: Proceedings of the Third Workshop on Dependable Distributed Data Management*, pages 25–28, New York, NY, USA, 2009. ACM
- 2007 Afrânio Correia Jr., José Pereira, Luís Rodrigues, Nuno Carvalho, Ricardo Vilaça, Rui Oliveira, and Susana Guedes. Gorda: An open architecture for database replication. In *IEEE International Symposium on Network Computing and Applications (NCA)*, pages 287–290, 2007

Conference Reviews

JPDC	Elsevier Journal of Parallel and Distributed Computing
SRDS	IEEE Symposium on Reliable Distributed Systems
LADC	Latin-American Symposium on Dependable Computing
DAIS	IFIP International conference on Distributed Applications and Interoperable Systems
MW4SOC	Workshop on Middleware for Service Oriented Computing

Research Projects

Date	April 2012 - Current
Projet	COSS - Gestão Integrada de Serviços Cloud com os Sistemas de Suporte às Operações (Plano de Inovação PT Inovação 2011-2013)
Main activities and responsibilities	Co-supervise the research grant holder that is working on the monitoring and elasticity support for OpenStack and also implement the JCloud driver for the Cloud platform from Portugal Telecom, SmartCloud PT.
Data	February 2011 – Current
Projet	CumuloNimbo: Highly Scalable Transactional Multi-Tier (FP7-257993) - http://www.cumulonimbo.eu/
Partners	UPM, SAP, FORTH, Yahoo, McGill, Universidade do Minho, Flexiant
Main activities and responsibilities	I'm designing and developing the CumuloNimbo's Query Engine layer that is responsible for taking an SQL query and generating an efficient execution plan for the query, to be executed over the tuple store layer. For this, I had identified several fundamental design modifications to an existing SQL query engine (Derby), allowing it to be distributed and efficiently run SQL on top of a scalable data store (HBase) while being able to scale with the data store. I had also done the preliminary evaluation of the prototype using both standard benchmarks for relational databaes (TPC-W e TPC-C) and NoSQL databases (YCSB - https://github.com/brianfrankcooper/YCSB).
Date	March 2011 - Current
Projet	Stratus: Stratus: A Layered Approach to Data Management in the Cloud FCOMP-01-0124-FEDER-015020 - http://stratus.lsd.di.uminho.pt/

Main activities and responsibilities	Write of proposal. I have defined a novel correlation-aware placement strategy that considers arbitrary tags on tuples and combines the usage of a Space Filling Curve (SFC) with random partitioning to store and retrieve correlated tuples. I have also designed DataDroplets, a flexible, efficient, dependable and fully decentralized key-value store for the management of data in a Cloud Environment. DataDroplets shifts the current trade-off of Cloud based tuple stores towards the needs of common business users, providing additional consistency guarantees and higher level data processing primitives smoothing the migration path for existing applications. I also developed a working prototype of DataDroplets with implementations of the novel and state of the art data placement strategies for key-value stores.
Date	July 2011 - March 2012
Projet	OSS - Aproveitamento de Serviços Cloud Computing (Plano de Inovação PT Inovação 2011-2013)
Main activities and responsibilities	Co-supervise the research grant holder that is working on the installation and maintenance of OpenStack, a massively scalable cloud operating system (http://www.openstack.org/). Design of an abstract model for the management of computing instances independent of the Cloud provider and development of a Web portal for the supply of resources, both computing and storage, in the Cloud.
Date	April 2010 - June 2011
Projet	CC - Cloud Computing em Data Centers (Plano de Inovação PT Inovação 2010-2012)
Main activities and responsibilities	Studied the viability of using NoSQL databases as complementary solutions to traditional relational database management systems to achieve scalability and elasticity in telecom applications.
Date	November 2005 – March 2008
Projet	GORDA: Open Replication of Databases (FP6-004758) - http://gorda.di.uminho.pt/
Partners	Universidade of Minho, Università della Svizzera Italiana, Faculdade de Ciências da Universidade de Lisboa, INRIA, Continuent, MySQL
Main activities and responsibilities	I participated: in the design of the GAPI (http://gorda.di.uminho.pt/community/gapi/) mainly in the recovery interfaces; design and implementation of the recovery protocols and in the integration with the group communication toolkits in Escada Replication Server (http://escada.sourceforge.net/) and in the design and implementation of the Sequoia (http://sequoia.continuent.org/) GAPI implementation

Program Committee

2013	Workshop on Planetary-Scale Distributed Systems - http://wpsds13.lsd.di.uminho.pt
2012	7th MW4NG Workshop - Middleware for Next Generation Internet Computing - http://www.dedisy.org/mw4ng12/

Event Organization

2013	Organizing chair of the Workshop on Planetary-Scale Distributed Systems - http://wpsds13.lsd.di.uminho.pt
2010	Help on organizing INForum 2010
2010	Organization of the 1 ^o Annual meeting of EDP10C 2010
2010	Founder of EDP10C, Grupo de Portugueses estudantes de Doutoramento em Ciências da Computação
2009	Help on organizing DAIS 2009

2007-2009	Member of the direction of the University of Minho ACM Student SIGPLAN Chapter
2008	Organizing committee, JOIN + JOBSHOP 2008
2007	Organizing Committee, 2nd Summer School on Generative and Transformational Techniques in Software Engineering (GTTSE 2007), 2 - 7 July, 2007, Braga, Portugal, http://wiki.di.uminho.pt/twiki/bin/view/Events/GTTSE2007/WebHome

Event Participation

Middleware 2012	ACM/IFIP/USENIX 13th International Conference on Middleware - http://middleware2012.cs.mcgill.ca/
SEEW 2012	Science and Engineering Entrepreneurship Workshop Learning by Doing - http://seew2012.inescporto.pt/
SOSP 2011	ACM Symposium on Operating Systems Principles - http://sosp2011.gsd.inesc-id.pt/
DisCoTec 2011	IFIP International Conference on Distributed Applications and Interoperable Systems. - http://discotec.ru.is/dais/main
OTM 2010	OnTheMove federated conferences and workshops - http://www.onthemove-conferences.org/
INForum 2010	Simpósio de Informática - http://inforum.org.pt/INForum2010
SRDS 2009	IEEE Symposium on Reliable Distributed Systems - http://www.cse.buffalo.edu/srds2009/
DSN 2009	IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) - http://2009.dsn.org/
EuroSys 2009	ACM EuroSys Conference http://eurosys2009.informatik.uni-erlangen.de
HTDC 2009	Winter School: Hot Topics in Distributed Computing - http://proton.inrialpes.fr/~quema/htdc2009/
JOIN 2008	Jornadas de Informática e JOBSshop 2008 http://di.uminho.pt/jornadas/join08
JOIN 2004	3D Animation, Scripting Languages, Search on the Internet, Artificial Intelligence, UMTS
Workshop WCT08	Oracle 9i Application Server and Oracle Discoverer.
JOIN 2003	Open Source, Data Mining, Programming platforms and Paradigms, Mobile Computing.
Jornadas da LESI 2002	Systems, Security and Mobility

Work Experience

Date	November 2005 – March 2008
Employer	Computer Science and Technology Center (CCTC) – University of Minho
Occupation	Software Engineer and Researcher
Main activities and responsibilities	I worked in the scope of the IST FP6 project "GORDA - Open Replication of Databases" on the design, development and assessment of online recovery protocols for replicated databases and contributing to the middleware implementation of the GORDA project.
Date	March 2005 - October 2005
Employer	PT Inovação
Occupation	Internship
Type of business or sector	Telecommunications and Information Technology

Main activities and responsibilities

I worked in the research project shIPnet (Service Handling on IP networks). This project aims at defining the new standard architectural framework for supporting the PT Inovação family of products, designed to meet the challenges and needs of next generation networks (NGN) in the context of convergence of fixed and mobile systems. The architecture is based on open standards from the 3GPP IP Multimédia Subsystem (IMS). I researched on the IMS standards and design and develop a prototype for the implementation of the X-CSCF allowing to tests basic VoIP functionalities.

Software Projects

Date
Title
Description
Main activities and responsibilities
Technology

2013
MeT - <https://github.com/fmaia/met>
MeT is a Cloud-enabled framework that can be used alone or in conjunction with OpenStack for the automatic and heterogeneous reconfiguration of HBase. MeT is an workload aware system that provides automatic elasticity for the HBase NoSQL database. MeT not only adds and removes nodes automatically and according to system load, but also reconfigures them according to the observed workloads. As a result, it achieves a significant increase in overall system performance
Openstack and HBase management.
Java, Python

Date
Title
Description
Technology

2013
Escada TPC-C - <https://github.com/rmpvilaca/EscadaTPC-C>
Escada TPC-C is an easy to use JDBC benchmark that closely resembles the TPC-C standard for OLTP. Database's supported include PostgreSQL, MySQL, and Derby.
Java

Date
Title
Description
Technology

2011
BenchmarkSQLHBase - <https://github.com/rmpvilaca/BenchmarkSQLHBase>
A fork of the BenchMarkSQL, <http://sourceforge.net/projects/benchmarksql/>, to run atop Apache Derby and with a console version supporting multiple clients with disjoint warehouses
Java

Date
Title
Description
Main activities and responsibilities
Technology

2010
UBlog Benchmark - <https://github.com/rmpvilaca/UBlog-Benchmark>
UBlog is a performance evaluation toolkit that mimics the usage of the Twitter social network. The workload may be used with both key-value stores and relational databases.
Design and development of the benchmark.
Java

Date
Title
Description

2005-2008
Escada Replicator - <http://escada.sourceforge.net/>
ESCADA is a portable and flexible database replication system that accommodates multiple reliability, performance, and applicability tradeoffs: Target scenarios range from shared-nothing clusters to remote sites for disaster prevention. Currently supported database management systems include Apache Derby, PostgreSQL, and MySQL.

Main activities and responsibilities	Participated in the design and implementation of the recovery protocols and in the integration with the group communication toolkits.
Technology	Java, Spring, and jGCS
Date	2005-2008
Title	GAPI - https://launchpad.net/gapi
Description	The GORDA Architecture and Programming Interface (GAPI) allows portable replication protocols to be easily implemented and reused with different database management systems. It is designed to be provided either natively by the database server or by a server wrapper. The included functionality supports a variety of replication protocols, both master-slave or multi-master, synchronous or asynchronous. It should also be useful to build other data management tools. The GAPI uses reflection, as abstract transaction processing concepts appear as objects in the target data model and programming language, and is based on widely used and tested design patterns and APIs (e.g. JDBC).
Main activities and responsibilities	Participated in the design of the GAPI mainly in the recovery interfaces and In the design and implementation of the Sequoia (http://sequoia.continuent.org/) GAPI implementation.
Technology	Java
Date	2006
Title	Portorama - http://portorama.portodigital.pt
Description	Portorama is a platform for sharing photos of the city of Porto. Portorama is meant to foster a community willing to create a rich, comprehensive and truly diverse photo collection of Porto. Portorama enables you to georeference and arbitrarily tag your photos of Porto and share them across the community through a peer-to-peer network. The sets of shared photos are made available by the Portorama community and searchable by their localization, title, author and assigned tags. All photos are stored by the community in a fully decentralized manner and under the Creative Commons license.
Main activities and responsibilities	Participated in the design and implementation of UI and integration with GoogleMaps.
Technology	P2P, GoogleMaps, GWT and AppFuse

Awards

2005	Governo Civil de Braga and Câmara Municipal de Braga Award (Sociedade Martins Sarmento)
2005	Engenheiro António de Almeida Award
2005	Best Engineering Student Award (Ordem dos Engenheiros)
2005	Bolsa de Estudo por Mérito
2005	Associação Industrial do Minho Award
2005	University of Minho Merit Award
2004	University of Minho Merit Award
2004	20th place at MIUP 2004 Portuguese inter-University Programming Contest
2003	University of Minho Merit Award
2003	3rd place on MIUPinho 2003 local programming contest at Department of Informatics
2002	University of Minho Merit Award
2000	Escola Secundária de Maximinos Merit Award
1999	Escola Secundária de Maximinos Merit Award
1998	Escola Secundária de Maximinos Merit Award

Technical skill and competences

Programming Languages	C,C++,C#, Java, Visual Basic, HTML, CSS, Prolog, Haskell, Python, XML and Ruby
DataBases	SQL, Informix, MySQL, Derby, PostegreSQL, and Oracle9.0i
Web Programming	Ajax, GWT, Flex, and Ruby Rails
Model-oriented specification	UML and VDM-SL
Operating Systems	Linux, Unix, Windows 9x/NT/2000/XP and Mac OS
Cryptography	Java CE, IAIK CE provider for Java, and OpenSSL
Network protocols	IP, TCP, UDP, DNS, SNMP, Mail, and Multicast
Distributed systems	Java RMI, CORBA, Distributed Objects, Concurrent Programming, Dependability, Replication, Consensus, Group Communication, Database Replication, Peer-to-peer systems
Multimédia	OpenGL

Mother tongue(s)

*Self-assessment
European level^(*)*

English

Portuguese

Understanding		Speaking		Writing
Listening	Reading	Spoken interaction	Spoken production	
B2 Independent user	C1 Proficient user	B2 Independent user	C1 Proficient user	C1 Proficient user

^(*) Common European Framework of Reference (CEF) level

Additional Information

2007-Current
2008-2009

Member of the social solidarity that helps homeless, Projecto SA
Member of Braga's local coordination team of the social solidarity project Rabo de Peixe Sabe Sonhar