

LUÍS MANUEL OLIVEIRA SOARES

Travessa da Carreira de Cima
n.º 35, Corvite
4805-638 Guimarães
PORTUGAL

mobile: (351) 914479918
home: (351) 253478622
fax: (351) 253604441
e-mail: los@di.uminho.pt
skype id: losoares

languages: Portuguese and English (spoken/written), Spanish (verbal)
Curriculum Vitae revision date: July 2007

nationality: portuguese
date of birth: 11th September 1979

EDUCATION BACKGROUND

- ◇ *Universidade do Minho, Portugal*: **Ph.D. in Informatics.** **2006 - Presently**
Thesis on synchronous clustered database replication based on the interface defined by the *Open Replication of Databases* specification.
- ◇ *Universidade do Minho, Portugal*: **Master of Informatics.** **2003 - 2006**
Thesis on development and testing synchronous database replication.
Cumulative Grade Point Average: 17/20. Thesis rated *Very Good* (top grade).
Level in national or international classification: ISCDE 6.
- ◇ *Universidade do Minho, Portugal*: **Licenciatura em Engenharia de Sistemas e Informática I.T. Branch** (5-year Bachelor of Science in Informatics). **1997 - 2002**
Cumulative Grade Point Average: 15/20.
Level in national or international classification: ISCDE 5A.
- ◇ *Instituto Britânico de Guimarães*: **First Certificate of English** (University of Cambridge ESOL Examinations). Final Grade Point: B. **1990 - 1995**

WORK EXPERIENCE

- ◇ **Software Engineer and Researcher** at Universidade do Minho (Distributed Systems Group) - Braga, Portugal. **2002 - Presently**

I have been responsible to perform web/software development (Plone, GWT and Spring MVC with AppFuse) database replication research and implementation and system administration tasks. Detailed description of my work experience is available below:

1. Involved in the European funded project named GORDA (Open Replication of Databases). I am directly involved in three different project activities: *i*) design and implementation of a general purpose data replicator written in Java (Spring IoC, Appia and jGCS). I am also a co-author of a wide-area inter-cluster replication protocol (WICE) included in this software module; *ii*) implementation of the GORDA API (GAPI) binding to PostgreSQL; *iii*) quality assurance resorting to regression testing based on unit testing (JUnit), evaluation, by means of hybrid simulation mixing real/simulated components under realistic environments, continuous integration, and statistical performance assessment using python and R. I was also responsible to perform source code maintenance (using bazaar, maven2 and trac).
2. Took part in the National funded project named ESCADA (Fault Tolerant Scalable Distributed Databases). Responsible for testing and evaluation of the protocols under study (DBSM and PDBSM) which required instrumentation of the Linux kernel and PostgreSQL internals. Results obtained were published in relevant scientific papers on multi-master synchronous database replication protocols.
3. Within the context of the national program *Porto Digital*, participated in all stages of the Portorama product development life-cycle (currently in beta-testing stage). The product enables the sharing of geo-referenced Oporto photos using P2P, GoogleMaps and GWT. I was responsible to implement the website using Java and AppFuse (Hibernate, Spring MVC, Spring IoC, MySQL, AXIS, Acegi and Jetty). Additionally, I participated in the client side development (using GWT and AJAX).
4. Maintained several sites based on Plone content management system. These included the research group website as well as the GORDA, PSON, ESCADA and StrongRep projects websites. I was responsible to customize the access control engine by modifying ACLs and content workflows. These modifications greatly reduced complexity while sharing content with the public in general as well as sharing content among project partners.

5. Installed and currently support internal servers based on CentOS and Ubuntu. The set of services provided include: plone, qmail, ezmlm, xen, vmware, apache2, jetty, ldap, samba/nfs/OCFSv2, bind, dhcp. I was also responsible to drive software development into using maven2 as the build tool, hudson as continuous integration framework and trac as the issue management system. Together these three tools, which I maintain on a daily basis, with the help of bazaar SCM and proximity (maven2 proxy), increased productivity and reduced software project management complexity.
6. As a teacher assistant, lectured operating systems concepts and design, focusing distributed multi-threaded programming. The technology used in the classes was Linux and C.

◇ **Software developer and systems administrator** at Bragaredes - I.R.I Telecomunicações, Lda - Braga, Portugal. (Part-time) **2004 - 2006**

This was a part-time collaboration with Bragaredes. Responsibilities were wide-ranged and are described by my work experienced, which is detailed below:

1. Established a VoIP product based on Asterisk software PBX which is currently replacing the traditional PBXs, present core business of the company.
2. Maintained and developed the company content management system, which automated internal content deployment leading to increased business. Reimplemented the company website based using Mambo CMS. The new website, eased interaction between customers and Bragaredes as well as attracted more web visitors.
3. Led software development (python/GTK and MVC paradigm) for retrieving data from several devices like bar-code label scanners and present it to the user.
4. Administered web hosting servers. Common duties involved handling qmail, apache, plesk and backups issues. I was also responsible to provide support for internal servers (debian sarge): tasks included maintenance of the following services: Exim, MailMan, Plone, Mambo, Drupal, Asterisk, Apache, Samba, IPCop firewall and VPNs using IPsec.

◇ **Systems Analyst** at Enabler S.A. - Porto, Portugal. (Concurrently with University degree.) **2002**

This position was concurrent with the Bachelor's classes.

1. Maintained a data warehouse batch process, based on **Oracle/Retek** solutions, by writing *ProC* and *PL/SQL* scripts.

QUALIFICATIONS

◇ 7+ years of strong and extensive Linux/UNIX system engineering and administration proficiency in complex client/server enterprises (using GNU/Linux since 1997 - Red Hat 5.0) | Red Hat based (Red Hat, Fedora Core, Mandrake / Mandriva) and Debian based (Debian, Knoppix, Ubuntu) distributions | x86, PowerPC, & RISC architectures | Tools: syslog, logwatch, ntop, rsync, ethereal, tcpdump, etc. Capable of administering Xen and VMWare virtual machines.

◇ 7+ years proficiency in TCP/IP networking infrastructure: router, switch, and WiFi access point configuration; services administration: DNS, SNMP, FTP, HTTP/S, DHCP, NTP, IMAP/S, POP/S, SMTP, Samba, NFS, LDAP (OpenLDAP and Fedora DS); security: Netfilter/iptables, VPN, etc.

◇ Proficiency in web/software application development: Java (RMI, JDBC, JCA), Python, Perl, groovy, PHP, SQL, XML (XSL), YAML, C, Visual Basic, Prolog, Bash shell, HTML, CGI, UML software modeling and Algebra of programming (implemented in Haskell and VDM), Berkeley Sockets, RPC, Threads. Make use of Eclipse framework for Java/C/Python development. Capable of developing and deploying solutions based on group communication toolkits, specifically: Spread, JGroups and Appia.

◇ Proficient in database, web server and network server administration, particularly Apache web server and PostgreSQL.

◇ Advanced skills in source code management with extensive experience in gnu-arch (tla), Bazaar (baz), Bazaar-ng (bzn). Issue tracking with trac and plone collective-ng. Experience in Maven2 (mvn) build system and administering maven repositories and proxies (using proximity and Jetty).

- ◊ Expert knowledge in testing, evaluation and benchmarking. Familiar with several workloads (especially TPC-C and TPC-W) and extensive skills in unit testing (JUnit, PyUnit) and simulation modeling (SSF, SSFNet, OMNet++, JSim). Distributed systems testing know-how.
- ◊ Implementation and Maintenance of Linux mail servers: Qmail (mainly qmail-toaster) and Exim. Deployment and administration of mailing-lists (using ezmlm and mailman) as well as setting up mail filters (maildrop and procmail), SPAM filters (bogofilter, spamassassin and dspam) and anti-virus (clamav).
- ◊ Experience in engineering and administering Open Source telecommunications switching and PBX powered by Asterisk; with features like conference calling, voicemail, and automatic call distribution, and supporting technologies like VoIP (SIP), ISDN, GSM, and PSTN channels.
- ◊ Implementation and administration of content management systems: Mambo, Plone / Zope, Drupal.
- ◊ Practical application of security technology to business applications, and thorough understanding of cryptography practices including public key/asymmetric and shared key/symmetric cryptography, digital certificates, md5sum/checksums, SSL, and PKI. Tools include PGP/GnuPG, SSH, Java Cryptography Architecture, OpenSSL, Apache certificates, WiFi Security (WEP, WPA), Tripwire, Nmap etc.
- ◊ Oracle PL/SQL and ProC practical experience and training.
- ◊ Excellent verbal and written command of Portuguese and English languages; good verbal ability of Spanish language.
- ◊ Advanced troubleshooting skills in all listed IT proficiencies; versatile learner of technology; and analytical team player.
- ◊ Willing to travel up to 100% and relocate.

PUBLICATIONS

- ◊ *A Pragmatic Protocol for Database Replication in Interconnected Clusters*, In Proc. of The 12th IEEE International Symposium Pacific Rim Dependable Computing, (PRDC'06), Riverside - USA, December 2006.
- ◊ *Evaluation of Group-based Database Replication Using Centralized Simulation*, Master Thesis, Universidade do Minho, Braga - Portugal, July 2006.
- ◊ *Group-based Replication of On-line Transaction Processing Servers*, In Proc. of the 2nd Latin-American Symposium on Dependable Computing, (LADC'2005), Salvador da Bahia - Brazil, October 2005.
- ◊ *Experimental Performability Evaluation of Middleware for Large-Scale Distributed Systems*, In Proc. of the 7th Intl. Workshop on Performability Modeling of Computer and Communication Systems (QEST-PMCCS'2005), Torino - Italy, September 2005.
- ◊ *Testing the Dependability and Performance of Group Communication Based Database Replication Protocols*, In IEEE Intl. Conf. on Dependable Systems and Networks - Performance and Dependability Symposium (DSN-PDS'2005), Yokohama - Japan, June 2005.
- ◊ *Revisiting Epsilon Serializability to improve the Database State Machine*, In Proc. of the Workshop on Dependable Distributed Data Management (SRDS-WDDDM'2004), Florianópolis - Brazil, October 2004.
- ◊ *Evaluating database replication in ESCADA*, In Proc. of the Workshop on Dependable Distributed Data Management (SRDS-WDDDM'2004), Florianópolis - Brazil, October 2004.
- ◊ *Avaliação de um SGBD Replicado Usando Simulação de Redes*, Conferência de Redes e Computadores (CRC'2003), Bragança - Portugal, October 2003.

PRESENTATIONS

- ◊ EuroSys, March 2007, Lisbon - Portugal. Open Replication of Databases Poster presentation (Poster Session).
- ◊ Master Thesis presentation, July 2006, Braga - Portugal.
- ◊ International Symposium on Computer Performance, Modeling, Measurements and Evaluation, October 2005, Juan-les-Pins - France. Work in progress presentation (WIP session).
- ◊ 7th International Workshop on Performability Modeling of Computer and Communication Systems (PMCCS-7), September 2005, Torino - Italy. Article presentation (session 6).

◇ Conferência de Redes e Computadores (CRC'2003), September 2003, Bragança - Portugal. Article presentation (session 3).

REVIEWING WORK

- ◇ Reviewer for the *27th International Conference on Distributed Computing Systems*. 2007
- ◇ Reviewer for the *26th International Conference on Distributed Computing Systems*. 2006
- ◇ Reviewer for the *XXXII Conferência Latinoamericana de Informática* 2006

OTHER ACTIVITIES

- ◇ Vice-chair of the University of Minho ACM Sigplan Student's Chapter. 2007
- ◇ Teaching Assistant, *Operating Systems II*, University of Minho. 2004 - 2005
- ◇ Teaching Assistant, *Operating Systems I*, University of Minho. 2004 - 2005
- ◇ Teaching Assistant, *Operating Systems I*, University of Minho. 2003 - 2004

NOTABLE PROJECTS

◇ *Universidade do Minho - Braga, Portugal.*

Portorama: 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. You may find portorama at <http://portorama.portodigital.pt> (it is in beta-stage).

Database Replication Evaluation: Making use of discrete event simulation, a library of network models and my own library of database models (based on queuing theory), evaluated three database replication protocols based on group communication. This work resulted in my Master thesis and is very valuable to the understanding of these novel database replication protocols.

Additional details at: <http://gsd.di.uminho.pt/members/los/abstracts/masters> .

ESCADA Replicator: participated in the design and implementation of a suite of data replication protocols. These are freely available at (<http://gorda.di.uminho.pt/community>) under an open source license. The replicator is written in Java and makes use of group communication (with Appia toolkit), dependency injection (Spring IoC), and GORDA API. This replicator provides a mean of replicating data among databases that implement GORDA API (currently there are GAPI bindings to PostgreSQL and Derby).

AWARDS AND HONOURS

- ◇ Student merit award at Universidade do Minho. 2001
- ◇ Student merit award at Universidade do Minho. 2000

References

Available upon request.