

CURRICULUM VITAE

DR. SPYROS ZOUPANOS

PERSONAL INFORMATION

Date of birth: 3rd of December, 1982

Nationality: Greek

Spoken languages: English, French, Greek

Current position: Researcher and software engineer at École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

Tel: +41 (0) 77 420 49 46

E-mail: spyros@zoupanos.net

Homepage: <http://spyros.zoupanos.net>



I am a researcher, software engineer and data management professional, currently working at École Polytechnique Fédérale de Lausanne (EPFL) in Switzerland as a scientific collaborator and lead developer of an Open Science platform for managing simulations. I have obtained a Ph.D. from INRIA, France in efficient peer-to-peer data management as well as a B.Sc. and an M.Sc. in Computer Science from the National University of Athens, Greece. I spent two years at the Max-Planck-Institut für Informatik in Germany as a post-doctoral fellow pursuing big data analytics problems. I have also worked at Quartet Financial Systems in Paris, France on the development of an in-memory multidimensional database and at Max-Planck-Institut für Physik in Germany where I managed and implemented the single sign-on project of the Institute.

CURRENT POSITION

6/2015-present Researcher and software engineer at EPFL, Lausanne, Switzerland.

We build a novel Open Science platform for managing simulations, generating data and disseminating them. The platform is composed of AiiDA (<http://www.aiida.net>), which is used as the core engine to automate simulations and track provenance, and Materials Cloud (<https://www.materialscloud.org>) which is the dissemination portal and cloud simulation platform. My responsibility as one of the lead developers is to contribute to the data management, scalability, and software engineering problems that arise.

Technologies used: Python 2.7/3.5, SQLAlchemy, Django, PostgreSQL, Web Services, Javascript.

RECENT POSITIONS

7/2014-6/2015 Software engineer and manager of the single sign-on project at Max-Planck-Institut für Physik, Munich, Germany.

Single sign-on infrastructure allows users to authenticate in a secure way using the same credentials in various systems enabling them to use a variety of services. I analysed the institute's needs, I designed and I performed the migration to the new authentication mechanism.

Technologies used: Kerberos, LDAP, NFS, Active Directory, Samba, Linux, Windows.

10/2012-7/2014 Senior developer at Quartet Financial Systems, Paris, France.

ActivePivot enables instant analysis of large volumes of complex and dynamic data. It is a powerful in-memory database platform that continuously processes new events and passes updated results instantly to users or other downstream services.

My responsibility was to provide efficient solutions to open problems, to redesign parts of the system to match new updated needs and to participate in the analysis and the documentation of complex parts of the platform. Moreover, I also made the connection between academia and industry by attending venues, by following the state of the art and by participating in the European project CoherentPaaS. I also worked on adapting and configuring ActivePivot to Quartet's clients needs.

Technologies used: Java 8 core and Web services.

- 9/2010-9/2012 Post-doctoral fellow at Max-Planck-Institut für Informatik, Saarbrücken, Germany.
The post-doc focused on knowledge extraction from large document collections. We analysed journal articles by extracting useful information categorised in certain relations. We also worked on the extraction of frequent word sequences (with gap) from big document sets. We developed and optimized distributed (MapReduce) algorithms focusing on good partitioning of the given collection and problem. Research results presented in top-tier international conferences .
Technologies used: Java 7, MapReduce, Hadoop, HBase, HDFS.
- 12/2009-9/2010 Post-doctoral collaborator of the Leo Team at INRIA-Saclay-Ile de France, Orsay, France.
After the end of my Ph.D. thesis, I continued working in the same team on remaining open problems.
- 9/2009-9/2010 Teaching assistant at University Paris 9 (Dauphine). For more information, please see below.
-

EDUCATION

- 10/2006–12/2009 Ph.D. student at INRIA-Futurs (Gemo Team), Orsay, France.
Thesis title: *Efficient peer-to-peer data management*, Supervisors: *Prof. Serge Abiteboul and Prof. Ioana Manolescu*.
The ActiveXML language enables the declarative specification of distributed applications integrating XML data and Web services. During the first half of my thesis, I worked on specifying and prototyping an AXML optimizer. In the sequel, I worked on the efficient rewriting of XQuery queries using stored views at a structured P2P network. We developed a full application able to index views, rewrite queries based on the available views and execute the rewritten queries. Large scale experiments were performed using 250 computers. Research results presented and prototypes demonstrated in top-tier international conferences and national summer schools.
Technologies used: Java core and Web services, XML, XQuery, XPath, XSLT, bash scripting.
- 10/2004–10/2006 Postgraduate Student (M.Sc. in Advanced Information Systems) at the Department of Computer Science and Telecommunications, National University of Athens, Greece. GPA: 8.66/10
Thesis title: *Optimizing queries using materialized views*, Supervisor: *Prof. Alex Delis*.
The thesis focused on building a tool illustrating the performance gains enabled by the usage of materialized views in a relational database management system by estimating query costs and by visualising relational plans.
Technologies used: C++, QT for the GUI.
- 10/2000–10/2004 Undergraduate Student at the Department of Computer Science and Telecommunications, National University of Athens, Greece. GPA: 7.79/10
Thesis title: *Performance Measuring Tool for Database Systems*, Supervisor: *Prof. Yannis Ioannidis*.
We developed a web-based application enabling students to test the performance of a relational database without the need to know SQL.
Technologies used: HTML, Javascript, Oracle database.
-

PROGRAMMING AND SYSTEMS EXPERIENCE

- Programming skills: Java 8 (core and Web services), Python (experience with frameworks: SQLAlchemy and Django) C++, C, Bash, XSLT, XPath, XQuery, JavaScript.
 - Databases: PostgreSQL, Oracle, Microsoft Analysis Services, ActivePivot.
-

TEACHING EXPERIENCE

- 1/2010–2/2010 Teaching XML, XSLT, XPath, XQuery in the postgraduate programme of Univ. Paris 9.
- 1/2010–5/2010 Teaching assistant in the postgraduate course *Relational databases (Bases de données relationnelles)*, Univ. Paris 9.
- 12/2009–3/2010 Teaching assistant in the postgraduate course *Semistructured data and XML (Données semi-structurées et XML)*, Univ. Paris 11.

10/2009–1/2010	Teaching assistant in the undergraduate course <i>Java programming</i> , Univ. Paris 9.
1/2009	Teaching assistant in the postgraduate course <i>Semistructured data and XML (Données semi-structurées et XML)</i> , Univ. Paris 11.
10/2008–11/2008	Teaching assistant in the postgraduate course <i>Information management on the web (Gestion de l'Information sur Internet)</i> , Univ. Paris 11.
02/2006–7/2006	Teaching assistant in the undergraduate course <i>Implementation of Database Systems</i> , Univ. of Athens.
02/2006–7/2006	Computer lab assistant in the Department of computer science computer lab, Univ. of Athens.
10/2004–2/2005	Teaching assistant in the undergraduate course <i>Theory of Computation</i> , Univ. of Athens.
2002	Teaching assistant in the training programme of computer science for high school teachers.

PUBLICATIONS

Conferences

- S. Zoupanos, L. Kahle, S. P. Huber, M. Uhrin, N. Mounet, R. A. Häuselmann, S. Kumbhar, L. Talirz, A. Cepellotti, F. Gargiulo, A. Merkys, B. Kozinsky, N. Marzari, G. Pizzi. *AiiDA: A simulation platform with full provenance support and flexible workflows*. PASC (The Platform for Advanced Scientific Computing Conference co-sponsored by ACM and CSCS) 2018 - poster (not in the proceedings).
- I. Miliaraki, K. Berberich, R. Gemulla, S. Zoupanos. *Mind the Gap: Large-Scale Frequent Sequence Mining*. SIGMOD (ACM's Special Interest Group on Management Of Data Conference) 2013.
- K. Karanasos, A. Katsifodimos, I. Manolescu, S. Zoupanos. *ViP2P: Efficient XML Management in DHT Networks*. ICWE (International Conference on Web Engineering) 2012 - short paper.
- Y. Wang, B. Yang, S. Zoupanos, M. Spaniol, G. Weikum. *Scalable Spatio-temporal Knowledge Harvesting*. WWW (World Wide Web Conference) 2011 - poster.
- I. Manolescu, K. Karanasos, V. Vassalos, S. Zoupanos. *Efficient XQuery rewriting using Multiple Views*. ICDE (IEEE's International Conference on Data Engineering) 2011.
- K. Karanasos, S. Zoupanos. *Viewing a World of Annotations through AnnoVIP*. ICDE (IEEE's International Conference on Data Engineering) 2010 and BDA (Bases de Données Avancées) 2010 - software demonstration.
- I. Manolescu, S. Zoupanos. *Materialized views for P2P XML warehousing*. BDA (Bases de Données Avancées) 2009.
- S. Abiteboul, T. Allard, P. Chatalic, G. Gardarin, A. Ghitescu, F. Goasdoué, I. Manolescu, B. Nguyen, M. Ouazara, A. Somani, N. Travers, G. Vasile, S. Zoupanos. *WebContent: Efficient P2P Warehousing of Web Data*. VLDB (Very Large Data Base Endowment Conference) 2008 - software demonstration.
- S. Abiteboul, I. Manolescu, S. Zoupanos. *OptimAX: Optimizing Distributed ActiveXML Applications*. ICWE (International Conference on Web Engineering) 2008 and BDA (Bases de Données Avancées) 2008.
- S. Abiteboul, I. Manolescu, S. Zoupanos. *OptimAX: efficient support for data-intensive mash-ups*. ICDE (IEEE's International Conference on Data) 2008 - software demonstration.
- S. Abiteboul, I. Manolescu, S. Zoupanos. *OptimAX: optimizing distributed continuous queries*. BDA (Bases de Données Avancées) 2007 - software demonstration.

Workshops

- S. Zoupanos, M. Uhrin, S. P. Huber, R. A. Häuselmann, L. Kahle, N. Mounet, D. Gresch, S. Kumbhar, T. Müller, L. Talirz, J. Boullier, A. Cepellotti, F. Gargiulo, A. Merkys, B. Kozinsky, N. Marzari, G. Pizzi. *The high-throughput AiiDA platform: Recent developments*. COMDI (MARVEL NCCR's international workshop on conceptual and methodological advances underlying the computational design and discovery of novel materials) 2018 - poster.
- B. Nguyen, S. Zoupanos. *The WebContent XML Store*. Sources Ouvertes et Services 2010 (workshop in conjunction with RFIA 2010) - short paper.
- I. Manolescu, S. Zoupanos. *XML materialized views in P2P networks*. Database Technologies for Handling XML Information on the Web (DataX) 2009 (workshop in conjunction with EDBT 2009 & ICDT 2009)

Thesis

- Ph.D. thesis: S. Zoupanos. *Efficient peer-to-peer data management*.
- M.Sc. thesis: A Delis, S. Zoupanos. *Optimizing queries using materialized views*.
- B.Sc. thesis: Y. Ioannidis, A. Bakogiannis, S. Zoupanos. *Performance Measuring Tool for Database Systems*.

Other

- K. Karanasos, A. Katsifodimos, I. Manolescu, S. Zoupanos. *The ViP2P Platform: XML Views in P2P*. Inria research report, November 2011.
 - *Interdeposit certification for ViP2P from the Agency for the Protection of Programs (APP)*. June 2011
 - I. Manolescu, L. Afanasiev, A. Arion, J. Dittrich, S. Manegold, N. Polyzotis, K. Schnaitter, P. Senellart, S. Zoupanos, D. Shasha. *The Repeatability Experiment of SIGMOD 2008*. SIGMOD (ACM's Special Interest Group on Management Of Data) Record - Volume 37, Number 1, March 2008.
-

OTHER ACTIVITIES

- Participated in the Swiss national project MARVEL (<http://nccr-marvel.ch>). This included: participating in the development of a platform within the project's domain, preparing reports and participating in the project's workshops.
 - Participated in the European project CoherentPaaS (<http://coherentpaas.eu>). This included: preparing & reviewing reports and participating in the project's workshops.
 - Participated in the European project LAWA (<http://www.lawa-project.eu>). This included: preparing reports, representing Max-Planck to conference calls, participating in workshops and doing research within the project's domain.
 - Participated in the project Views in P2P (<http://vip2p.saclay.inria.fr>) within Codex. Views in P2P is a project focused on the management of materialized XML views in P2P networks. Codex is an ANR grant DEFIS 2009-2012.
 - Demonstration of the system ViP2P in the BDA (Bases de Données Avancées) 2010 summer school.
 - TOIT (ACM's Transactions on Internet Technology) Journal reviewer. SIGMOD (ACM's Special Interest Group on Management Of Data Conference) 2009, WWW (World Wide Web Conference) 2009, ICDE (IEEE's International Conference on Data Engineering) 2008, BDA (Bases de Données Avancées) 2007 external reviewer.
 - SIGMOD (ACM's Special Interest Group on Management Of Data) 2008 experimental repeatability committee member.
 - Web master for the International Workshop on Programming Language Techniques for XML (PLAN-X) 2008 and the International Workshop on the Web and Databases (WebDB) 2007.
 - Attended 8th EDBT Summer School Database Technologies for Novel Applications.
-

LANGUAGE SKILLS

- Greek (native tongue), English (proficient), French (advanced), German (beginner).