

Grigoris (Gregory) Karvounarakis

Dept. of Computer and Information Science
University of Pennsylvania
3330 Walnut Street
Philadelphia, PA 19104, USA

Phone: 215-681-7228
Fax: 215-898-0587
Email: gkarvoun [at] cis.upenn.edu
Web: <http://www.cis.upenn.edu/~gkarvoun>

RESEARCH INTERESTS

Primary: Data Integration/Exchange, Data Provenance, Query Languages and Web Data Management Systems.

Other interests (or areas I have worked on): Web Services, Digital Libraries and Metadata, Distributed/Peer-to-Peer Systems.

EDUCATION

University of Pennsylvania, Philadelphia, PA, USA **July 2009 (expected)**
PhD, Computer and Information Science, GPA: 3.98/4.0
Dissertation: *Provenance in Collaborative Data Sharing*
Advisors: Zachary G. Ives, Val Tannen

University of Crete, Heraklio, Greece **November 2000**
M.Sc. Computer Science, Avg.grade: 8.8/10 (Excellent)
Thesis: *Querying RDF metadata for Community Web Portals*
Advisors: Vassilis Christophides, Dimitris Plexousakis and Panos Constantopoulos.

University of Crete, Heraklio, Greece **September 1998**
B.Sc. Computer Science, Avg.grade: 8.08/10 (Very good)
Thesis: *Submission and repository management tools for digital libraries using WWW*
Advisor: Sarantos Kapidakis

EMPLOYMENT

University of Pennsylvania, Philadelphia, PA, USA **September 2003 - Present**
Research assistant
Lead developer in the ORCHESTRA project of the Penn Database Research Group, as a part of my PhD dissertation. ORCHESTRA is a Collaborative Data Sharing System (CDSS) that supports update exchange between heterogeneous peers, while allowing disagreement between them, tracks the provenance of propagated data and exploits it for a variety of operations such as trust assessment, incremental maintenance etc.

Bell Labs, Lucent Technologies, Murray Hill, NJ, USA **June 2004 - August 2004**
Research Intern (Mentor: Rick Hull)
Worked on verification of termination of E-service compositions, defined using a powerful model that allows exchange of messages with arbitrary contents through unbounded queues, as well as unbounded spawning of new processes.

Institute of Computer Science (ICS) - FORTH, Heraklio, Greece **March 2003 - August 2003**
Research assistant/R&D engineer, Information Systems Group
Worked on an architecture for a Semantic Web Integration Middleware that translates queries in the RDF query language *RQL*, posed over a mediator, to queries over heterogeneous (XML, RDBMS) sources, and employs the chase/backchase technique to minimize redundant queries.

Greek Air Force
(Mandatory) Military Service

March 2002 - March 2003

Institute of Computer Science (ICS) - FORTH, Heraklio, Greece **September 2001 - March 2002**
Research assistant/R&D engineer, Information Systems Group
Worked on the optimization of the interpreter of *RQL*, including the use of encoding schemes to optimize transitive closure queries on hierarchies and query rewriting.

Bell Labs, Lucent Technologies, Murray Hill, NJ, USA
Research Intern (Mentor: Rick Hull)

June 2001 - August 2001

Worked on an event-driven architecture allowing composition of session-oriented E-services called AZTEC, implementing the workflow definition language XASC.

Institute of Computer Science (ICS) - FORTH, Heraklio, Greece
Research assistant/R&D engineer, Information Systems Group

December 2000 - May 2001

Developed extensions in the design and implementation of *RQL*.

Institute of Computer Science (ICS) - FORTH, Heraklio, Greece
Graduate scholar/research assistant, Information Systems Group

July 1999 - November 2000

Designed and implemented the RDF query language *RQL* over an RDBMS (Postgres), that was used to store RDF metadata and schemas, as a part of my Master's thesis.

University of Crete, Heraklio, Greece
Part-time digital library system administrator

February 2000 - July 2000

Worked on the project OMOSPONDO-SKEPSIS, which allowed distributed indexing/searching of lecture notes and other academic material among several greek universities.

6th TEE/11th Public High Schools, Heraklio, Greece
High school teacher of Informatics

December 1998 - February 2000

Institute of Computer Science (ICS) - FORTH, Heraklio, Greece
Undergraduate/graduate scholar/research assistant, Pleiades (Parallel and Distributed Systems) Group

June 1996 - June 1999

Designed and implemented a set of tools for the submission and repository management of on digital library systems (in particular the NCSTRL project DL system - DIENST) as a part of my diploma thesis.

University of Crete, Heraklio, Greece
Undergraduate scholar/research assistant, Network Administration Group (UCNET)

June 1995 - June 1996

English Language School, Agii Deka, Greece
English language school teacher

June 1995 - September 1995

Publications

REFEREED
CONFERENCE
PUBLICATIONS

1. T. J. Green, G. Karvounarakis, Z. G. Ives, and V. Tannen. Update Exchange with Mappings and Provenance. In *Proceedings of International Conference on Very Large Databases (VLDB)*, Vienna, Austria, 2007. Extended version available as Univ. of Pennsylvania report MS-CIS-07-26.
2. T. J. Green, G. Karvounarakis, and V. Tannen. Provenance Semirings. In *Proceedings of ACM Symposium on Principles of Database Systems*, Beijing, China, 2007. (presented at conference).

3. Y. Theoharis, V. Christophides, and G. Karvounarakis. Benchmarking Database Representations of RDF/S Stores. In *Proceedings of International Semantic Web Conference*, Galway, Ireland, Nov 2005.
4. G. Karvounarakis, S. Alexaki, V. Christophides, D. Plexousakis, and Michel Scholl. RQL: A Declarative Query Language for RDF. In *Proceedings of International World Wide Web Conference*, Honolulu, Hawaii, May 2002. Preliminary version available as ICS-FORTH Technical Report 296, 2000.
5. G. Karvounarakis, V. Christophides, D. Plexousakis, and S. Alexaki. Querying RDF Descriptions for Community Web Portals. Agadir, Morocco, November 2001. (presented at conference).
6. G. Karvounarakis and S. Kapidakis. Submission and Repository Management of Digital Libraries Using WWW. In *TERENA Networking Conference*, Lisbon, Portugal, May 2000.

REFEREED
JOURNAL
PUBLICATIONS

1. Z. G. Ives, T. J. Green, G. Karvounarakis, N. E. Taylor, and V. Tannen. The ORCHESTRA Collaborative Data Sharing System. *SIGMOD Record*, 37(3):26–32, 2008.
2. J. N. Foster and G. Karvounarakis. Provenance and Data Synchronization. *IEEE Data Engineering Bulletin*, 30(4):13–21, 2007.
3. V. Christophides, G. Karvounarakis, D. Plexousakis, M. Scholl, and S. Tourtounis. Optimizing Semantic Web Queries Using Labeling Schemes. *Computer Networks and ISDN Systems*, 1:207–228, 2004.
4. G. Karvounarakis, A. Magkanaraki, S. Alexaki, V. Christophides, D. Plexousakis, M. Scholl, and K. Tolle. Querying the Semantic Web with RQL. *Computer Networks and ISDN Systems*, 42(5):617–640, 2003.
5. V. Christophides, G. Karvounarakis, A. Magkanaraki, D. Plexousakis, and V. Tannen. The ICS-FORTH Semantic Web Integration Middleware (SWIM). *IEEE Data Engineering Bulletin*, 26(4):11–18, 2003.
6. S. Kapidakis G. Karvounarakis. Submission and Repository Management of Digital Libraries Using WWW. *Computer Networks and ISDN Systems*, 34(6):861–872, 2000.

REFEREED
BOOK
CHAPTERS

1. G. Karvounarakis, A. Magkanaraki, S. Alexaki, V. Christophides, D. Plexousakis, M. Scholl, and K. Tolle. RQL: A Functional Query Language for RDF. In P.M.D.Gray, L.Kerschberg, P.J.H.King, and A.Poulovassilis, editors, *Functional Approach to Data Management: Modelling, Analyzing and Integrating Heterogeneous Data*. LNCS Series, Springer-Verlag, 2003.
2. G. Karvounarakis. Datalog. In Ling Liu and M. Tamer Özsu, editors, *Encyclopedia of Database Systems*. Springer-Verlag, 2008. to appear.

REFEREED
WORKSHOP
PUBLICATIONS

1. G. Karvounarakis and Z. G. Ives. Bidirectional Mappings for Data and Update Exchange. In *International Workshop on the the Web and Databases*, Vancouver, BC, Canada, 2008. (presented at workshop) Extended version available as Univ. of Pennsylvania report MS-CIS-08-17.
2. V. Christophides, G. Karvounarakis, I. Koffina, G. Kokkinidis, A. Magkanaraki, D. Plexousakis, G. Serfiotis, and V. Tannen. The ICS-FORTH SWIM: A Powerful Semantic Web Integration Middleware. In *1st International Workshop on Semantic Web and Databases (SWDB)*, Berlin, Germany, September 2003.

3. V. Christophides, R. Hull, G. Karvounarakis, A. Kumar, G. Tong, and M. Xiong. Beyond Discrete E-Services: Composing Session-oriented Services in Telecommunications. In *2nd Workshop on Technologies for E-Services (TES)*, Rome, Italy, September 2001.
4. J. Robie, L. M. Garshol, S. Newcomb, M. Biezinski, M. Fuchs, L. Miller, D. Brickley, V. Christophides, and G. Karvounarakis. The Syntactic Web: Syntax and Semantics on the Web. *Markup Languages: Theory and Practice*, 3(4):411–440, 2001.
5. S. Alexaki, V. Christophides, G. Karvounarakis, and D. Plexousakis. On Storing Voluminous RDF Descriptions: The Case of Web Portal Catalogs. In *International Workshop on the the Web and Databases*, Santa Barbara, California, May 2001. (presented at workshop).
6. S. Alexaki, V. Christophides, G. Karvounarakis, D. Plexousakis, and K. Tolle. The ICS-FORTH RDFSuite: Managing Voluminous RDF Description Bases. In *2nd International Workshop on the Semantic Web (SemWeb)*, pages 1–13, Hong Kong, China, May 2001.
7. V. Christophides, D. Plexousakis, G. Karvounarakis, and S. Alexaki. Declarative Languages for Querying Portal Catalogs. In *DELOS Workshop: Information Seeking, Searching and Querying in Digital Libraries*, pages 115–120, Zurich, Switzerland, December 2000.
8. S. Alexaki, V. Christophides, G. Karvounarakis, D. Plexousakis, K. Tolle, Bernd Amann, Irini Fundulaki, Michel Scholl, and Anne-Marie Vercoustre. Managing RDF Metadata for Community Webs. In *International Workshop on the World Wide Web and Conceptual Modeling (WCM'00)*, pages 140–151, Salt Lake City, Utah, October 2000. (presented at workshop).

MAGAZINE
PUBLICATIONS

1. S. Alexaki, N. Athanasis, V. Christophides, G. Karvounarakis, A. Magkanaraki, D. Plexousakis, and K. Tolle. The ICS-FORTH RDFSuite: High-level Scalable Tools for the Semantic Web. *ERCIM News*, 51, 2002.
2. G. Karvounarakis, V. Christophides, and D. Plexousakis. RQL: A Declarative Query Language for RDF. *D-Lib Magazine*, 7(12), 2001.

OTHER
PUBLICATIONS
(DEMON-
STRATIONS,
POSTERS,
THESES AND
TECHNICAL
REPORTS)

1. G. Karvounarakis and V. Tannen. Conjunctive Queries and Mappings with Unequalities. Technical Report MS-CIS-08-37, University of Pennsylvania, 2008.
2. T. J. Green, N. Taylor, G. Karvounarakis, O. Biton, Z. Ives, and V. Tannen. ORCHESTRA: Facilitating Collaborative Data Sharing. In *Demonstration Sessions of ACM SIGMOD Conference on Management of Data and of Data Integration in the Life Sciences (DILS) Workshop*, 2007. (presented at both).
3. S. Alexaki, N. Athanasi, V. Christophides, G. Karvounarakis, A. Magkanaraki, D. Plexousakis, and K. Tolle. The ICS-FORTH RDFSuite: High-level Scalable Tools for the Semantic Web. In *Poster Session of the Eleventh International World Wide Web Conference (WWW'02)*, Honolulu, Hawaii, May 2002.
4. A. Magkanaraki, G. Karvounarakis, V. Christophides, D. Plexousakis, and T. Anh. Ontology Storage and Querying. Technical Report 308, ICS-FORTH, 2002.
5. G. Karvounarakis. Querying RDF Metadata for Community Web Portals. Master thesis (available in Greek only), University of Crete, 2000.
6. S. Kapidakis and G. Karvounarakis. Submission and Repository Management Tools for Digital Libraries, with WWW Interface. In *Demonstration Session of the 3rd European Conference on Research and Advanced Technology for Digital Libraries*, page 495, Paris, France, September 1999. (presented at conference).
7. G. Karvounarakis. Submission and Repository Management of Digital Libraries Using WWW. Diploma thesis (available in Greek only), University of Crete, 1998.

PROFESSIONAL
ACTIVITIES

PC Member in:

- Conference on Information and Knowledge Management (CIKM) 2008
- Semantic Desktop and Social Semantic Collaboration Workshop 2006

Journal referee for:

- International Journal on Semantic Web and Information Systems (JSWIS) 2008
- Journal of Web Semantics (JWS) 2008
- Data and Knowledge Engineering (DKE) 2008

External reviewer for:

- ACM SIGMOD Conference on Management of Data 2007, 2008
- ACM Symposium on Principles of Database Systems (PODS) 2005, 2008
- International Conference on Very Large Databases (VLDB) 2003, 2006
- International Conference on Data Engineering (ICDE) 2009
- International Conference on Extending Database Technology (EDBT) 2006
- International Semantic Web Conference (ISWC) 2003, 2006
- International Database Engineering and Applications Symposium (IDEAS) 2008
- International Workshop on the Web and Databases (WebDB) 2006
- European Conference on Research and Advanced Technology for Digital Libraries (ECDL) 2000
- Semantic Web Workshop 2000
- Hellenic Conference on Artificial Intelligence (SETN) 2001

Volunteered/assisted the organization of:

- Northeast Database and Information Retrieval (DB/IR) day, October 2005, University of Pennsylvania, Philadelphia, PA, USA
- Semantic Web Workshop, September 2000, Lisbon, Portugal
- European Conference on Research and Advanced Technology for Digital Libraries (ECDL), September 1998, Heraklion, Greece.

Other activities:

- December 2006: Attended the Bertinoro PhD School on Data and Service Integration.
- July 2003: Attended the Onassis Foundation Lectures in Computer Science, ICS-FORTH, Heraklio, Greece. A part of the lecture "The Semantic Web: Myths and Reality" given by my advisors Drs. Vassilis Christophides and Dimitris Plexousakis, was based on my M.Sc. thesis.
- March 2000 Invited to give a seminar on Web technologies and digital library systems by the National Documentation Center, National Hellenic Research Foundation.

SKILLS

Databases/Query Languages: IBM DB2, PostgreSQL, MS Access, SQL, Datalog, OQL, XQuery.

Programming languages: Java, C++, C, CaML, Basic, Pascal, Assembly, lisp, prolog.

Markup languages/Data Models: XML, RDF, HTML, CSS, XSL, VRML.

Scripting tools: Unix shell, Perl, JavaScript, CGI.

Other tools/platforms: Tomcat/servlets, CORBA, Mobile Agent Systems (Grasshopper/EasyAgent), compiler generator tools ((f)lex, yacc/javacc, bison/jflex, Cup), wrapper generators (w4f), Parallel programming languages/environments (PVM, MPI, CC++, HP parallel C compiler), cvs, svn

Languages: Greek, English (Certificate of Proficiency, Cambridge), German (Zertifikat aus Deutsch als Fremdsprache)

TEACHING
EXPERIENCE

Apart from the teaching jobs mentioned in the employment section, I have served as a teaching assistant in the following classes:

University of Pennsylvania, Theory of Computation (Undergraduate) <i>Instructor: Prof. Val Tannen</i>	Summer 2005
University of Pennsylvania, Algorithms (Graduate) <i>Instructor: Prof. Sampath Kannan</i>	Spring 2005
University of Pennsylvania, Database Systems (Undergraduate) <i>Instructor: Prof. Val Tannen</i>	Fall 2004
University of Crete, Compilers (Undergraduate) <i>Instructor: Prof. Dimitris Plexousakis</i>	Spring 2000
University of Crete, Machine Learning (Graduate) <i>Instructor: Prof. Giorgos Potamias</i>	Fall 1999
University of Crete, Software Design Lab (Undergraduate) <i>Instructor: Prof. Spyros Lalis</i>	Spring 1999
University of Crete, Probability Theory (Undergraduate) <i>Instructor: Prof. Catherine Houstis</i>	Fall 1998

REFERENCES

Zachary G. Ives
Assistant Professor
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA, USA.
zives [at] cis.upenn.edu

Val Tannen
Professor
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA, USA.
val [at] cis.upenn.edu

Vassilis Christophides
Associate Professor
Computer Science Department
University of Crete &
Researcher, ICS-FORTH
Heraklion, Crete, Greece.
christop [at] ics.forth.gr

Rick Hull
Research Staff Member and Manager
IBM T.J. Watson Research Center
Yorktown Heights, NY, USA
rhull [at] acm.org

Peter Buneman
Professor
Laboratory for Foundations of Computer Science
School of Informatics
Edinburgh, Scotland
opb [at] inf.ed.ac.uk

PERSONAL

Date of Birth: June 9, 1976
Nationality: Greek
Marital status: Single