Nikolaos Papailiou

Contact Information

Nationality:	Greek
Date of Birth:	18/03/1989
Address:	Computing Systems Laboratory, National Technical University of Athens, Room 21.34, Zografou Campus, 15773 Athens, Greece
E-mail:	npapa@cslab.ece.ntua.gr
Website:	http://www.cslab.ece.ntua.gr/ npapa
Google Scholar:	http://scholar.google.gr/citations?user=UAfopP4AAAAJ
DBLP:	http://dblp.uni-trier.de/pers/hd/p/Papailiou:Nikolaos

Research Interests

I am interested in data management, distributed systems and cloud computing. I have worked with state-of-the-art distributed frameworks and developed distributed data management systems and algorithms. My recent interests include distributed graph databases, large-scale data-analytics, application elasticity using cloud resources and uncertain graph algorithms.

Education

National Technical University of Athens, Greece *October 2011-November 2016*
Computer Science and Engineering Ph.D. student,
under the supervision of prof. Nectarios Koziris
and prof. Dimitrios Tsoumakos.

Hong Kong University of Science and Technology *September-December 2015*
Internship on uncertain graph related research projects,
under the supervision of prof. Dimitris Papadias.

National Technical University of Athens, Greece *2006 - October 2011*
Electrical and Computer Engineering (5 year diploma)
Graduated with general degree: 8.8 / 10.0

2nd public High-school of Tripoli, Tripolis, Greece. *June 2006*
Graduated with general degree: 19.6 / 20.0

Professional Experience

EU funded project ASAP *2014-present*
Adaptive, scalable analytic workflow execution platform. ASAP handles the planning and execution of workflows containing alternative execution plans and operators executed over various engines (https://github.com/project-asap).

EU funded project CELAR *2012 - 2015*
Cloud computing framework that handles the automatic, multi-grained elasticity provisioning of generic applications in the cloud (https://github.com/CELAR).

EU funded project Arcomem *2011 - 2014*
A targeted web archiving platform, enriching the archives with social web information. Development of H_2RDF(https://code.google.com/p/h2rdf), a fully distributed RDF data store that combines HBase and MapReduce.

Awards and Honors

IKY fellowships of excellence for postgraduate studies in Greece - SIEMENS program

ACM SIGMOD/PODS 2014 student travel award.

IEEE BIGDATA 2013 student travel award.

Selected Publications

Nikolaos Papailiou, Dimitrios Tsoumakos, Panagiotis Karras, Nectarios Koziris:
"Graph-Aware, Workload-Adaptive SPARQL Query Caching." In Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, 2015.

Dimitrios Sarlis, **Nikolaos Papailiou**, Ioannis Konstantinou, Georgios Smaragdakis, Nectarios Koziris:
"Datix: A System for Scalable Network Analytics." ACM SIGCOMM Computer Communication Review 45(5): 21-28 (2015)

Katerina Doka, **Nikolaos Papailiou**, Dimitrios Tsoumakos, Christos Mantas, Nectarios Koziris:
"IReS: Intelligent, Multi-Engine Resource Scheduler for Big Data Analytics Workflows." In Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (Demo track), Melbourne, Victoria, Australia, 2015.

Nikolaos Papailiou, Dimitrios Tsoumakos, Ioannis Konstantinou, Panagiotis Karras, Nectarios Koziris:
"H2RDF+: an efficient data management system for big RDF graphs." In Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data (Demo Track), Snowbird, Utah, USA, 2014.

Nikolaos Papailiou, Ioannis Konstantinou, Dimitrios Tsoumakos, Panagiotis Karras, Nectarios Koziris:
"H2RDF+: High-performance distributed joins over large-scale RDF graphs." In proceedings of the 2013 IEEE International Conference on Big Data (BigData 2013), Santa Clara, CA, USA, 2013.

Nikolaos Papailiou, Ioannis Konstantinou, Dimitrios Tsoumakos, Nectarios Koziris:
"H2RDF: adaptive query processing on RDF data in the cloud." In Proceedings of the 21th International Conference on World Wide Web (WWW demo track), Lyon, France, 2012.

Technical Skills

- Programming languages:
  - JAVA: Development of several open-source, production projects (H2RDF, CELAR, ASAP).
  - C++: Extended several production C++ projects, constructed a compiler for a programming language.
  - Python, Scala, Matlab, R, Haskell, ML, Prolog: Used in several graduate and undergraduate courses.
- Distributed systems
  - Hadoop, HBase: Development of a RDF database on top of Hadoop and HBase.
  - Yarn: Construction of a workflow execution engine that utilizes YARN containers.
  - Spark: Profiling of various MLlib and SparkSQL operators.
  - MPI: Experimented with several MPI-based RDF main-memory databases.
  - HIVE, SparkSQL, Cassandra, ElasticSearch: Performance experimentations.
- Relational databases: MySQL, PostgreSQL
- Web development: HTML, PHP, CSS, JavaScript, Jetty, D3.js
- Operating systems: Ubuntu, Debian, CentOS, bash scripting
- Virtual machines, containers: OpenStack, KVM, Docker
- Code Integration: Git, Maven, ant, make, cmake, Jenkins, Nexus

Languages

- Greek: Mother language
- English: Fluent user
- French: Independent user