

## Ioannis A. Mytilinis

---

(0030) 6984750374

gmytilinis@gmail.com or gmytil@cslab.ece.ntua.gr

<http://www.cslab.ece.ntua.gr/~gmytil/>

### EDUCATION

**PhD Student** in the National Technical University of Athens (NTUA) (2013-present)

**Thesis:** *Algorithms and Techniques for Large Scale Approximate Query Processing*

**Advisor:** Prof. Nectarios Koziris

**Research Interests:** Stream Processing, Heterogeneous Computing, Approximate Query Processing, Large Scale Data Management, Blockchain

**Diploma** in Electrical and Computer Engineering, NTUA (2012) – Grade: 8.04/10

**Thesis:** *A Distributed Storage System for Social Network Information*

**Advisor:** Prof. Nectarios Koziris

1st ACM Europe Summer School on **Data Science and Big Data** (2017)

Linux Foundation Certified **System Administrator** (2018)

### PUBLICATIONS

I. Mytilinis, D. Tsoumakos and N. Koziris: **Scaling the construction of wavelet synopses for maximum error metrics**. In IEEE Transactions on Knowledge and Data Engineering. doi: 10.1109/TKDE.2018.2867185., 2018

I. Mytilinis, C. Bitsakos, K. Doka, I. Konstantinou and N. Koziris: **The Vision of a Heterogeneous Scheduler**. In proceedings of the 2018 IEEE International Conference on Cloud Computing Technology and Science (CloudCom), Xtreme Cloud workshop 2018.

K. Doka, I. Mytilinis, I. Giannakopoulos, I. Konstantinou, D. Tsitsigkos, M. Terrovitis and N. Koziris: **Exploiting Social Networking and Mobile Data for Crisis Detection and Management**. In Proceedings of the International Conference on Information Systems for Crisis Response and Management in Mediterranean Countries. Springer, Cham 2017.

I. Mytilinis, D. Tsoumakos and N. Koziris: **Distributed Wavelet Thresholding for Maximum Error Metrics**. In proceedings of the 2016 ACM SIGMOD/PODS International Conference on Management of Data (SIGMOD'16), Research track, San Francisco, USA, 2016.

I. Mytilinis, I. Giannakopoulos, I. Konstantinou, K. Doka, D. Tsitsigkos, M. Terrovitis, L. Giampouras and N. Koziris: **MoDisSENSE: A Distributed Spatio-Temporal and Textual Processing Platform for Social Networking Services**. In Proceedings of the 2015 ACM SIGMOD/PODS International Conference on Management of Data, Demo track, Melbourne, Victoria, Australia, 2015

I. Mytilinis, D. Tsoumakos, V. Kantere, A. Nanos and N. Koziris: **I/O Performance Modeling for Big Data Applications over Cloud Infrastructures**. In proceedings of the 2015 IEEE International Conference on Cloud Engineering (IC2E 2015), 9-13 March, Tempe, AZ, USA.

I. Mytilinis, I. Giannakopoulos, I. Konstantinou, K. Doka and N. Koziris: **MoDisSENSE: A Distributed Platform for Social Networking Services over Mobile Devices**. In proceedings of the 2014 IEEE International Conference on Big Data (BigData 2014), Washington DC, USA, October 27-30 2014

I. Konstantinou, D. Tsoumakos, I. Mytilinis and N. Koziris: **DBalancer: Distributed Load Balancing for NoSQL Data-stores**. In proceedings of the 2013

ACM SIGMOD/PODS International Conference on Management of Data, New York, USA, June 22-27 2013.

## AWARDS

**2016 NTUA Thomaidion Award** for outstanding publications in scientific journal or conference proceedings.

## INTERNSHIPS

**Research Intern**, IBM, T.J. Watson Research Center, NY (Summer 2017)

**Research Intern**, IBM, T.J. Watson Research Center, NY (Summer 2016)

## FUNDED RESEARCH PROJECTS

**E2Data: European Energy-efficient Big Data Stacks. (EU-funded project).**

E2Data proposes an end-to-end solution for Big Data deployments that will fully exploit and advance the state-of-the-art in infrastructure services by delivering more performance from fewer resources. The E2Data stack will achieve this by dynamically profiling, compiling and optimising code for execution on chosen devices, such as CPUs, GPUs, FPGAs, and others. By removing the need for developers to craft specific device code in languages like CUDA or OpenCL, E2Data will create substantial savings in developer time, while still exploiting the power of diverse device architectures, such as are currently offered by Microsoft, Amazon and others.

**Data Sovereignty through the use of Blockchain. (GR-funded project).**

The project proposes the combination of Cloud Computing and Blockchain Technology into a platform that enables the on-demand creation of a fully distributed, scalable and secure virtual infrastructure for data storage and processing. One of the main targets of the project is to face the limitations of current blockchain implementations in terms of required storage, latency and throughput by adopting different architectural choices that will render blockchain usable in the context of Cloud Computing.

**A Scalable Analytics Platform (ASAP) (EU-funded project).** The ASAP FP7 research project develops a dynamic open-source execution framework for scalable data analytics. The underlying idea is that no single execution model is suitable for all types of tasks, and no single data model (and store) is suitable for all types of data. Complex analytical tasks over multi-engine environments therefore require integrated profiling, modeling, planning and scheduling functions. The project has four goals: (i) A general-purpose task-parallel programming model and a runtime system to execute it in the cloud. The runtime will incorporate and advance state-of-the-art task-parallel programming models features: irregular general-purpose computations, resource elasticity, synchronization, data-transfer, locality and scheduling abstraction, ability to handle large sets of irregular distributed data and fault-tolerance. (ii) A modeling framework that constantly evaluates the cost, quality and performance of data and computational resources in order to decide on the most advantageous store, indexing and execution pattern available. (iii) A unique adaptation methodology that will enable the analytics expert to amend the task she has submitted at an initial or later stage. (iv) A state-of-the-art visualization engine that will enable the analytics expert to obtain accurate, intuitive results of the analytics tasks she has initiated in real-time.

**CLARIN-EL: Common Language Resources and Technology Infrastructure (GR-funded project).**

CLARIN-EL is the Greek counterpart of the CLARIN project, a pan-european effort for the collection and the distribution to the research community of language resources (text/speech/multimodal corpora, lexica, terminological glossaries etc.) in all languages and the relevant language processing tools (morphological/syntactic analysers, parsers, taggers, statistical tools etc.) through a

web-based Research Infrastructure.

**CELAR: Automatic, multi-grained elasticity-provisioning for the Cloud (EU-funded project).** Auto-scaling resources is one of the top obstacles and opportunities for Cloud Computing: consumers can minimize the execution time of their tasks without exceeding a given budget; cloud providers maximize their financial gain while keeping their customers satisfied and minimizing administrative costs. Many systems claim to offer adaptive elasticity, yet the “throttling” is usually performed manually, requiring the user to figure out the proper scaling conditions. In order to harvest the benefits of elastic provisioning, it is imperative that it be performed in an automated, fully customizable manner. CELAR delivers a fully automated and highly customizable system for elastic provisioning of resources in cloud computing platforms.

**MoDisSENSE: A Distributed Platform for the Development of Social Networking Services over Mobile Devices (GR-funded project).** MoDisSENSE enriches social networking services by exploiting the continuous data flow from the daily use of mobile phones. This flow includes data from user visited locations, contacts, calls and calendar combined with data acquired from the user’s social network (list of friends, profile and preferences). The project combines these heterogeneous data sources (geographic and social log files, user profiles and preferences and context information) and offers innovative services based on advanced searches and combined queries that exploit all these aforementioned sources by utilizing state-of-the-art distributed data processing techniques. Furthermore, the project deals with the development and deployment of services that exploit spatiotemporal data generated by user paths.

## TEACHING

### Teaching Assistant:

**Introduction to Programming** (2011-2012): Undergraduate program, 1st Semester, School of Electrical and Computer Engineering, National Technical University of Athens.

**Operating Systems** (2012-2017): Undergraduate program, 7-th Semester, School of Electrical and Computer Engineering, National Technical University of Athens.

**Advanced Database Systems** (2015, 2016, 2018): Undergraduate program, 9-th Semester, School of Electrical and Computer Engineering, National Technical University of Athens.

**Advanced Database Systems** (2019): Master’s program of Data Science and Machine Learning, School of Electrical and Computer Engineering, National Technical University of Athens.

### Lectures

**MapReduce, Distributed Systems** (2018): Undergraduate program, 9-th Semester, School of Electrical and Computer Engineering, National Technical University of Athens.

**Stream Processing, Big Data Analytics** (2019): Master’s program, Department of Informatics, Ionian University.

## ACADEMIC ACTIVITIES

### PC MEMBER:

International Conference on Machine Learning and Data Science (ICMLDS 2018)

### Reviewer:

IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing (CC-Grid 2019)

Conference on Information and Knowledge Management (CIKM 2018)

IEEE International Conference on Big Data (2018)

IEEE International Conference on Cloud Computing Technology and Science (Cloud-Com 2018)

International Conference on Smart Cities and Green ICT Systems (SMARTGREENS 2018)

IEEE International Conference on Data Engineering (ICDE 2017)

Conference on Information and Knowledge Management (CIKM 2017)

IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing (CC-Grid 2017)

IEEE International Conference on Big Data (2017)

International Conference on Computer Communications and Networks (ICCCN 2017)

International Conference on Machine Learning and Data Science (ICMLDS 2017)

AlgoCloud (2017)

IEEE NCA (2017)

Algorithms and Systems for MapReduce and Beyond (BeyondMR 2017)

IEEE International Conference on Big Data (2016)

International Conference on Service Oriented Computed (RMSOC workshop 2015)

IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing (CC-Grid 2015)

## TECHNICAL SKILLS

**Languages & Software:** JAVA, C, C++, Python, Scala, Ocaml, SQL, UNIX-Linux bash scripting, Docker, Kubernetes, eucatoools, REST, Git, Maven

**Databases:** PostgreSQL, MySQL, HBase, Cassandra, MonetDB, MongoDB

**Distributed Programming Frameworks:** Hadoop, Spark, Flink, Hama, Giraph, Mahout

## COMMUNITY SERVICE

Ministry of National Defense-Hellenic Navy – **Military service** (2008-2009)

## FOREIGN LANGUAGES

**English:** Certificate of Proficiency, IELTS: 7.5

**French:** Sorbonne II (Assez Bien)