

Dimitris Tsitsigkos

Date of Birth: 18 January 1989

Nationality: Hellenic

Address: 34 Zaxaria Papantwniou,
Kato Patissia, Athens, Greece

Phone: +30 210 8315040

Mobile: +30 6942951698

Email: tsitsigkosdim@gmail.com
dtsitsigkos@imis.athena-innovation.gr
dtsitsigkos@cse.uoi.gr

Webpage: <https://www.imis.athena-innovation.gr/en/people/member/79>

LinkedIn: <https://www.linkedin.com/in/dimitris-tsitsigkos>

Education

July 2019 - now

PhD candidate

Department of Computer Science and Engineering
University of Ioannina

PhD Title: Join Operators for Complex Data

Advisor: Prof. Mamoulis Nikos

November 2012 - September 2016

M. Sc. in Computer Systems Technology

Department of Informatics and Telecommunications
National and Kapodistrian University of Athens
G.P.A: 8,2/10

M.Sc. Thesis: Complex Event Processing(CEP) for Intrusion Detection

Advisor: Prof. Hadjiefthymiades Stathes

September 2006 - June 2012

B.Sc. in Informatics and Telecommunications

Department of Informatics and Telecommunications
National and Kapodistrian University of Athens
G.P.A: 7,7 / 10

B.Sc. Thesis: Clustering Wikipedia resources

Advisor: Prof. Gunopoulos Dimitris

May 2006

High School Graduate: Lycee Leonin of Patisiwn
G.P.A.: 18/20

Research Interests

- Join Operators of omplex data types (spatial, spatio-temporal)
- Real Time Analytics
- Parallel Algorithms
- Distributed Algorithms
- Big Data
- Data Anonymization
- Data Mining

Theses

May 2016 – November 2016

M.Sc. Thesis: *Complex Event Processing(CEP) for Intrusion Detection*

[Pervasive Computing Research Group \(p-comp\)](#)

Abstract: CEP is an application solution to the problem of security intrusions (anomaly-based intrusion detection) by using streams generated by IoT devices relevant to their network properties in order to detect abnormal behavior and notify the user via an alert. In our case, each device participating in a IoT network is handled as a sensor device that generates streams of network measurements using the Simple Network Management Protocol (SNMP). These measurements are provided as input to the Esper Complex Event Processing framework. I implemented three algorithms to perform anomaly detection, one algorithm with dynamic thresholds, Shewhart Controller Algorithm and Adaptive Resonance Theory Algorithm.

Programming languages: **Java**

January 2012 – September 2012

B.Sc. Thesis: *Clustering Wikipedia resources*

[KDD Lab](#)

Abstract: The goal of this thesis was to crawl and cluster Wikipedia pages into “natural” clusters about a specific topic. I explored different approaches to produce such indicative clusters, including keyword-based, link-based, and possible combinations. The final result was to identify topics on which a Wikipedia writer/contributor is an expert on, using external sources or parts of an article that they have authored.

Programming languages: **Java**

Working Experience

December 2012 - now

Software Engineer (full-time)

Institute for the Management of Information Systems (IMIS)
"Athena" Research and Innovation Center, Athens, Greece

April 2017 – November 2017

Software Engineer (full-time)

Hellenic Army Information Technology Support Center (KEPYES), Athens, Greece

As a member of IMIS, I have been involved in the following R&D projects:

October 2020 – now

MORE: Management of Real-time Energy data

Project Overview: MORE will deliver a platform to address the technical challenges in time series and stream management, focusing on the RES industry. Specifically, more's platform will introduce an architecture incorporating edge computing and cloud computing to address responsiveness and the need for sophisticated analytics simultaneously. This architecture will be combined with time series summarization techniques, or as we more accurately term them in MORE, modeling techniques for sensor data. Models are any compressed representations that allow the reconstruction of the original data points of a time series (e.g., a linear function) within a known error-bound (possibly zero). This approach synergizes with the edge computing approach since summarization can be done at the edge, reducing the load in the whole data processing pipeline.

Responsibilities: I implemented a continuous evaluation module of sliding window aggregations on the edge, using Java. Additionally, I implemented all the parallel and distributed versions of pattern extraction methods using Python and [Dask](#) framework.

Project website: <https://www.more2020.eu/>

Programming languages: **Java, Python**

Other technologies: **Dask**

August 2015 – September 2020

Amnesia: A Powerful Data Anonymization Platform

Project Overview: Amnesia is a web-based platform for anonymizing data, including relational, multi-dimensional, and hierarchical data. Currently, the platform supports a variety of anonymization algorithms such as Flush, parallel Flush, k-

anonymity, km-anonymity, and Incognito. Amnesia also supports the definition of custom anonymization rules and hierarchies.

Responsibilities: I was the main developer of this project.

Project website: <https://amnesia.openaire.eu/>

Programming languages: **JavaScript, Java**

Other technologies/tools: **Spring**

December 2012 - July 2015

MoDisSENSE: A Distributed Spatio-Temporal and Textual Processing Platform for Social Networking Services.

Project Overview: MoDisSENSE is an open-source distributed platform that provides personalized search for points of interest (POIs) and trending events based on the user's social graph, i.e., by combining spatio-textual user-generated data (e.g. GPS traces, check-ins, user profiles, graph of friendship relations, user posts in social networks, etc.).

Responsibilities: I implemented a distributed version of the DBSCAN algorithm for POI discovery, and a distributed algorithm that collects GPS traces and reconstructs the trajectories of the end-users. Both algorithms were implemented using Hadoop, Hbase and postGIS. I also developed all web services that are relevant to POI discovery and suggestion.

Project website: <http://www.modissense.gr/>

Programming languages: **Java**

Other technologies/tools: **Hadoop, Hbase, PostGIS**

Technical Skills

- Programming Languages : C, C++, Java, Python
- Web Development : HTML, CSS, PHP, Javascript, JSP, JSF
- Operating Systems : Linux Ubuntu, Kubuntu, Microsoft Windows
- Graphics : Java Swing
- Databases : MySQL, PostgreSQL, PostGIS, Oracle(limited working experience)
- NoSQL Systems: Hadoop, Hbase
- Shell Programming : shell(bourne/bash/tcsh-Unix,awk)
- Frameworks: Spring

Teaching Experience

- **Fall Semesters 2013/2014/2015**
Teaching Assistant - B.Sc. course "Operating Systems"
Prof. Stathes Hadjiefthymiades
Department of Informatics and Telecommunications,
National Kapodistrian University of Athens, Greece

I was responsible for the programming exercises during the semester and the examination period.

- **Fall Semesters 2019/2020/2021/2022**

Teaching Assistant – B.Sc. course "Introduction to Programming"

Prof. Nikos Mamoulis

Department of Computer Science & Engineering,

University of Ioannina, Greece

My primary responsibility within the course is to provide assistance during the laboratory sessions. I was also responsible for the programming exercises during the semester and the examination period.

- **Spring Semester 2019/2020/2021/2023**

Teaching Assistant – B.Sc. course "Complex Data Management"

Prof. Nikos Mamoulis

Department of Computer Science & Engineering,

University of Ioannina, Greece

I was responsible for the programming exercises during the semester and the examination period.

Publications

- D. Tsitsigkos, P. Bouros, K. Lampropoulos, N. Mamoulis, and M. Terrovitis, "**Two-layer Space-oriented Partitioning for Non-point Data**," IEEE Transactions on Knowledge and Data Engineering (TKDE), to appear.
 - Programming Language: **C++**.
- P. Bouros, N. Mamoulis, D. Tsitsigkos, and M. Terrovitis, "**In-Memory Interval Joins**," The **VLDB Journal**, 30(4): 667-691, July 2021.
 - Programming Language: **C++**.
- D. Tsitsigkos, K. Lampropoulos, P. Bouros, N. Mamoulis, and M. Terrovitis, "**A Two-layer Partitioning for Non-point Spatial Data**," 37th International Conference on Data Engineering (ICDE), Chania, Greece, April 2021.
 - Programming Language: **C++**.
- P. Bouros, K. Lampropoulos, D. Tsitsigkos, N. Mamoulis, and M. Terrovitis, "**Band Joins for Interval Data**," 23rd International Conference on Extending Database Technology (EDBT), Copenhagen, Denmark, March 2020 (short paper).
 - Programming Language: **C++**.
- D. Tsitsigkos, P. Bouros, N. Mamoulis, and M. Terrovitis, "**Parallel In-Memory Evaluation of Spatial Joins**," 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL/GIS), Chicago, IL, November 2019 (short paper).
 - Programming Language: **C++**.

- Katerina Doka, Ioannis Mytilinis, Ioannis Giannakopoulos, Ioannis Konstantinou, Dimitrios Tsitsigkos, Manolis Terrovitis^{IP}, Nectarios Koziris: **Exploiting Social Networking and Mobile Data for Crisis Detection and Management. ISCRAM-med 2017: 28-40.**
 - Programming Language: **Java.**
- 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 (SIGMOD'15), Demo track, Melbourne, Victoria, Australia, 2015.
 - Programming Language: **Java.**

Spoken Languages

- Greek (mother tongue)
- English