PUBLICATIONS

1. P. Nikitopoulos, G. A. Sfyris, A. Vlachou, C. Doulkeridis, O. Telelis: **Pruning Techniques for Parallel Processing of Reverse Top-k Queries.** *In Distributed and Parallel Databases Journal, Springer, 39(1): 169-186, 2021*

2. A. Vlachou, C. Doulkeridis, N. Koutroumanis, D. Poulopoulos, K. Norvag. **The SPADES Framework for Scalable Management of Spatio-textual Data**. *In Proceedings of 24th Pan-Hellenic Conference on Informatics (PCI'20), Athens, Greece, November 2020*

3. S. Maropaki, S. Chester, C. Doulkeridis, K. Norvag: **Diversifying Top-k Point-of-Interest Queries via Collective Social Reach**. *In Proceedings of 29th ACM Conference on Information and Knowledge Management (CIKM'20), October 19-23, 2020*

4. N. Koutroumanis, P. Nikitopoulos, A. Vlachou, C. Doulkeridis: **NoDA: Unified NoSQL Data Access Operators for Mobility Data**. *In Proceedings of the 16th International Symposium on Spatial and Temporal Databases (SSTD'19), Vienna, Austria, August 2019*

5. G. M. Santipantakis, A. Glenis, C. Doulkeridis, A. Vlachou, G. A. Vouros: **stLD: Towards a Spatio-temporal Link Discovery Framework.** *In Proceedings of the International Workshop on Semantic Big Data (SBD'19) (workshop held in conjunction with SIGMOD'19), Amsterdam, The Netherlands, July 2019, pp.4:1-4:6*

6. P. Nikitopoulos, G. A. Sfyris, A. Vlachou, C. Doulkeridis, O. Telelis: **Parallel and Distributed Processing of Reverse Top-k Queries.** *In Proceedings of the 35th IEEE International Conference on Data Engineering (ICDE'19), Macau SAR, China, April 2019, pp.1586-1589*

7. A. Vlachou, C. Doulkeridis, A. Glenis, G. M. Santipantakis, G. A. Vouros: **Efficient Spatio-temporal RDF Query Processing in Large Dynamic Knowledge Bases**. *In Proceedings of the 34th ACM/SIGAPP Symposium On Applied Computing (SAC'19), Limassol, Cyprus, April 2019, pp.439-447*

PEOPLE

Akrivi Vlachou - Principal Investigator, Associate Professor
Christos Doulkeridis – Scientific Host, Associate Professor
Christos Kalyvas – Phd Candidate
Dimitris Poulopoulos - Phd Candidate
Kjetil Nørvåg – Professor

Alexandros Fakis - Researcher
Spyros Kasdaglis - Researcher
Konstantinos Nestorakis - Researcher
Konstantinos Platis - Researcher
Antonis Psarros - Researcher
Ioannis Tsirovasilis – Researcher

George S. Theodoropoulos - Researcher
Dimitrios Zografakis – Researcher
Panagiotis Tampakis – Researcher
Aikaterini Ntzelepi - Researcher

SPADES
SPAtiotextual Data Exploration at Scale

SPADES
SPAtiotextual Data Exploration at Scale

https://www.ds.unipi.gr/spades/

GET IN TOUCH

Grigoriou Lambraki 126
Piraeus 18532 Greece
(+30) 210 – 4142545
avlachou@unipi.gr
https://www.ds.unipi.gr/spades/

H.F.R.I.
Hellenic Foundation for
Research & Innovation

GSRI
GENERAL SECRETARIAT FOR
RESEARCH AND INNOVATION

Department of Digital Systems
University of Piraeus

Objectives

- Novel rank-aware query types complying with the paradigm of spatial-keyword search that cover a wide variety of information needs, targeting at the mobile user.

- Provision of more expressive querying mechanisms for points of interest that combine spatial information and textual relevance with temporal information and user preferences.

- Advanced distributed indexing structures capable to support complex spatial-keyword queries effectively, by means of harnessing the merits of spatial data structures and text indexes.

- Novel partitioning mechanisms and load balancing techniques for spatio-textual queries, aiming at efficient parallel data processing.

Support of advanced spatio-textual query types

Novel distributed index structures for spatio-textual search

Abstractions for parallel spatio-textual data processing

Impact

SPADES aims to address the limitations of spatio-textual data analysis and processing when applied in the context of Big Spatial Data, as witnessed by the lack of existing systems and techniques for this purpose. Achievement of this goal constitutes a substantial step forward in dealing with challenges emerging from management of Big Spatial Data. At a practical level, the research outcome will benefit applications such as spatio-textual search and retrieval, mining of spatio-textual data, next generation location-based services, and tourism-oriented applications to name a few. By exploiting SPADES the analysis of massive spatio-textual datasets (typically encountered in the aforementioned domains and especially in social networks) is going to be accelerated significantly. In consequence, applications will be able to query and analyse larger quantities of spatio-textual data in shorter time, thus speeding up the process of making new scientific discoveries.

Results

The results of the project are expected to be exploited by the local tourism business that could benefit from the provision of innovative location-based services over vast quantities of data to tourists. Such services entail queries that are processing-intensive and typically run for minutes rather than seconds, and often produce results that are not truly useful to the end user. SPADES promises to facilitate the access to location-based information, a task of particular importance to tourists. In this respect, the research results of SPADES are expected to benefit also society at large.

Highlight

Big Data Development Today

The Vision of NoDA in Track&Know

The NoSQL landscape
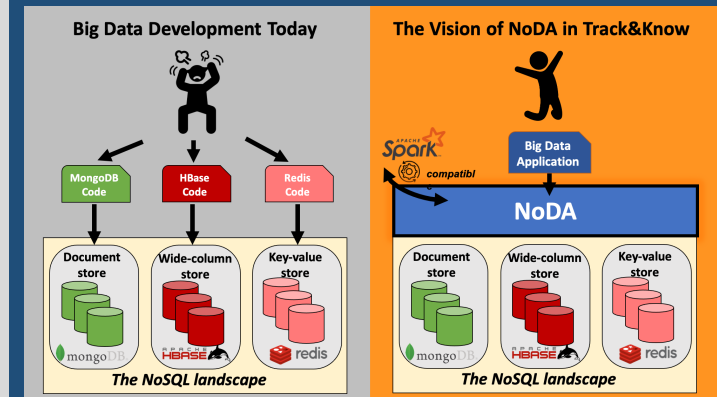
NoDA

- An abstraction layer that enables access to NoSQL stores by offering a unified view

- NoDA "hides" the query language of each underlying NoSQL store by providing access in a *simple* and *unified* manner

- It leverages upon the native API libraries of the NoSQL Systems

- NoDA focuses on mobility data and it offers the following operations:

  - Filtering (Boolean, Comparison and Geographical)
  - Aggregations
  - Sorting

- NoDA is compatible with big data frameworks such as Apache Spark