

# Ippokratis Pandis

Principal Engineer, Amazon Web Services  
 ippo@amazon.com / ippokratis@apache.com  
<http://www.pandis.net>

## RESEARCH INTERESTS

Data management systems, online transaction processing on modern hardware, analytics on Big Data

## ACADEMIC BACKGROUND

### PhD in Electrical & Computer Engineering,

May 2012

Dissertation: Scalable transaction processing through data-oriented execution

Advisor: Prof. Anastasia Ailamaki

Carnegie Mellon University

### MSc in Information Networking, Dec 2004

Thesis title: Context-aware distributed systems for pervasive and ubiquitous computing

Carnegie Mellon University

### Diploma in Computer Engineering & Informatics, Nov 2002

Thesis title: Development of a taxonomic hypertext management system

University of Patras, Greece

## HONORS & AWARDS

- Best Paper Award Runner Up, Outrageous Ideas and Vision Track, **CIDR**, 2013
- Best Demonstration Award, ACM **SIGMOD** Conference, 2011
- Intel Research Fellowship, Summer 2009
- Best Demonstration Award, IEEE **ICDE**, 2006
- Scholarship by Intracom S.A., 2003-2004

## PROFESSIONAL ACTIVITIES

PC chair: DaMoN 2014, DaMoN 2015, CloudDM 2016, NorCal DB Day 2017

PC area chair: CIKM 2014 – Main-memory, Parallel and Distributed Database Systems (DB Track)

PC: SIGMOD 2018, HPTS 2017, VLDB 2017, CIDR 2017, EuroSys 2017, DaMoN 2017, VLDB 2016 Industrial, HotCloud 2016, SIGMOD 2016, SIGMOD 2016 Industrial, ICDE 2016, DaMoN 2016, IMDM/ADMS 2016, SOCC 2015, IMDM 2015, HPTS 2015, CIDR 2015, VLDB 2015, ICDE 2015, EDBT 2015 Demo, HardBD 2015, HADM 2015, SIGMOD 2014, VLDB 2014, ICDE 2014, EDBT 2014 Demo, IMDM 2014, CloudDB 2014, VLDB 2013, ICDE 2013, CIKM 2013 Poster, IMDM 2013, CloudDB 2013, SSDBM 2013 Demo, VLDB 2012 Demo, DaMoN 2012, DbTest 2011, VLDB 2007, ACM Hypertext 2006

Steering Committee: BPOE series

Reviewer: Communications of the ACM (CACM), The VLDB Journal (VLDBJ), ACM Transactions on Database Systems (TODS), IEEE Transactions on Knowledge and Data Engineering (TKDE), MICRO46

Apache Project Committer: Apache Arrow (<http://arrow.apache.org/>)

## PUBLICATIONS

### Data Management Systems:

- T. Wang, R. Johnson, A. Fekete, and I. Pandis. “Efficiently making (almost) any concurrency control mechanism serializable.” The VLDB Journal (**VLDBJ**), 2017.

- K. Kim, T. Wang, R. Johnson and I. Pandis. “ERMIA: Fast memory-optimized database system for heterogeneous workloads.” In Proc. of ACM **SIGMOD**, 2016.
- D. Porobic, I. Pandis, M. Branco, P. Tozun and A. Ailamaki. “Characterization of the Impact of Hardware Islands on OLTP.” The VLDB Journal (**VLDBJ**), Special Issue on "Data Management on Modern Hardware", 2015.
- T. Wang, R. Johnson, A. Fekete, and I. Pandis. “The Serial Safety Net: Efficient Concurrency Control on Modern Hardware.” In Proc. of the 11th International Workshop on Data Management on New Hardware (**DaMoN**), 2015.
- M. Kornacker, A. Behm, V. Bittorf, T. Bobrovitsky, C. Ching, A. Choi, J. Erickson, M. Grund, D. Hecht, M. Jacobs, I. Joshi, L. Kuff, D. Kumar, A. Leblang, N. Li, I. Pandis, H. Robinson, D. Rorke, S. Rus, J. Russell, D. Tsirogiannis, S. Wanderman-Milne, and M. Yoder. “Impala: A Modern, Open-Source SQL Engine for Hadoop.” In Proc. of **CIDR**, 2015.
- G. Attaluri, R. Barber, N. Chainani, S. Lightstone, G. Lohman, I. Pandis, V. Raman, D. Sharpe, and R. Sidle. “Memory-Efficient Hash Joins.” Proc. of the VLDB Endowment (**PVLDB**), 8(4), 2015.
- R. Johnson, K. Kim, T. Wang, and I. Pandis. “Robust concurrency control in main-memory DBMS: What main memory giveth, the application taketh away.” In Proc. of the International Workshop on In-Memory Data Management and Analytics (IMDM), 2014.
- J.-G. Lee, G. Attaluri, R. Barber, N. Chainani, O. Draese, F. Ho, S. Idreos, M.-S. Kim, S. Lightstone, G. Lohman, K. Morfonios, K. Murthy, I. Pandis, L. Qiao, V. Raman, V. KulandaiSamy, R. Sidle, K. Stolze, and L. Zhang. “Joins on Encoded and Partitioned Data.” Proc. of the VLDB Endowment (**PVLDB**), 2014.
- R. Johnson, I. Pandis, and A. Ailamaki. “Eliminating unscalable communication in transaction processing.” The VLDB Journal (**VLDBJ**), Vol. 23(1), 2014.
- R. Barber, G. Lohman, R. Mueller, I. Pandis, V. Raman, and W. Wilcke. “Go Server Go: Parallel Computing with Moving Servers.” In Proc. of the ACM Symposium on Cloud Computing (**SoCC**), 2013.
- V. Raman, G. Attaluri, R. Barber, N. Chainani, D. Kalmuk, V. KulandaiSamy, J. Leenstra, S. Lightstone, S. Liu, G. M. Lohman, T. Malkemus, R. Mueller, I. Pandis, B. Schiefer, D. Sharpe, R. Sidle, A. Storm, and L. Zhang. “DB2 with BLU Acceleration: So much more than just a column store.” Proc. of the VLDB Endowment (**PVLDB**), 2013.
- O. Hassanzadeh, A. Kementsietsidis, B. Kimelfeld, R. Krishnamurthy, F. Özcan, and I. Pandis. “Next Generation Data Analytics at IBM Research.” Proc. of the VLDB Endowment (**PVLDB**), 2013.
- A. Ailamaki, R. Johnson, I. Pandis and P. Tözün. “Toward Scalable Transaction Processing - Evolution of Shore-MT.” Proc. of the VLDB Endowment (**PVLDB**), 2013.
- P. Tözün, I. Pandis, C. Kaynak, D. Jevdjic and A. Ailamaki. “From A to E: Analyzing TPC’s OLTP Benchmarks - The obsolete, the ubiquitous, the unexplored.” In Proc. of the International Conference on Extending Database Technology (**EDBT**), 2013.
- Y. Li, I. Pandis, R. Mueller, V. Raman, and G. Lohman. “NUMA-aware algorithms: the case of data shuffling.” In Proc. of **CIDR**, 2013.
- R. Johnson, and I. Pandis “The bionic DBMS is coming, but what will it look like?” In Proc. of **CIDR**, 2013. **Best Paper Award runner up**, Outrageous Ideas and Vision Track.
- P. Tözün, I. Pandis, R. Johnson, and A. Ailamaki. “Scalable and Dynamically Balanced Shared-Everything OLTP with Physiological Partitioning.” The VLDB Journal (**VLDBJ**), Vol. 22(2), 2013.
- D. Porobic, I. Pandis, M. Branco, P. Tözün, and A. Ailamaki. “OLTP on Hardware Islands.” Proc. of the VLDB Endowment (**PVLDB**), Vol. 5(12), 2012.
- R. Barber, P. Bendel, M. Czech, O. Draese, F. Ho, N. Hrle, S. Idreos, M.-S. Kim, O. Koeth, J.-G. Lee, T. Li, G. Lohman, K. Morfonios, R. Mueller, K. Murthy, I. Pandis, L. Qiao, V. Raman, R. Sidle, K. Stolze, S. Szabo. Business Analytics in (a) Blink. Data Engineering Bulletin (**DEBULL**), Special Issue on Column Store Systems, Vol. 35 (1), 2012.
- R. Johnson, I. Pandis, R. Stoica, M. Athanassoulis, and A. Ailamaki. “Scalability of write-ahead logging on multicore and multisoocket hardware.” The VLDB Journal (**VLDBJ**), Vol. 21 (2), 2012.

- I. Pandis, P. Tözün, R. Johnson, and A. Ailamaki. “PLP: Page Latch-free Shared-everything OLTP.” Proc. of the VLDB Endowment (**PVLDB**), Vol. 4 (10), 2011.
- R. Barber, P. Bendel, M. Czech, O. Draese, F. Ho, N. Hrle, S. Idreos, M.-S. Kim, O. Koeth, J.-G. Lee, T. T. Li, G. Lohman, K. Morfonios, R. Mueller, K. Murthy, I. Pandis, L. Qiao, V. Raman, S. Szabo, R. Sidle, and K. Stolze. “Blink: Not Your Father's Database!” In Proc. of **BIRTE**, 2011.
- I. Pandis, P. Tözün, M. Branco, D. Karampinas, D. Porobic, R. Johnson, and A. Ailamaki. “A Data-oriented Transaction Execution Engine and Supporting Tools.” In Proc. of ACM SIGMOD, 2011. **Best Demonstration Award**.
- S. Chen, A. Ailamaki, M. Athanassoulis, P. B. Gibbons, R. Johnson, I. Pandis, and R. Stoica. “TPC-E vs. TPC-C: Characterizing the New TPC-E Benchmark via an I/O Comparison Study.” SIGMOD Record, Vol. 39 (4), 2010.
- R. Johnson, I. Pandis, R. Stoica, M. Athanassoulis, and A. Ailamaki. “Aether: A scalable approach to logging.” Proc. of the VLDB Endowment (**PVLDB**), 2010. Invited for publication at the Special Issue of VLDB Journal for **VLDB 2010 Best Papers**.
- I. Pandis, R. Johnson, N. Hardavellas, and A. Ailamaki. “Data-Oriented Transaction Execution.” Proc. of the VLDB Endowment (**PVLDB**), Vol. 3 (1), 2010.
- R. Johnson, I. Pandis, and A. Ailamaki. “Improving OLTP Scalability using Speculative Lock Inheritance.” Proc. of the VLDB Endowment (**PVLDB**), 2009.
- R. Johnson, I. Pandis, N. Hardavellas, A. Ailamaki, and B. Falsafi. “Shore-MT: A Scalable Storage Manager for the Multicore Era.” In Proc. of the International Conference on Extending Database Technology (**EDBT**), 2009.
- R. Johnson, I. Pandis, and A. Ailamaki. “Critical Sections: Re-emerging Scalability Concerns for Database Storage Engines.” In Proc. of the International Workshop on Data Management on New Hardware (**DaMoN**), 2008.
- A. Ailamaki, and I. Pandis. “Query Processor”. Encyclopedia of Database Systems, Ling Liu and M. Tamer Özsu eds., 2008.
- N. Hardavellas, and I. Pandis. “Operator-Level Parallelism”, “Execution Skew”, “Inter-query Parallelism”, “Intra-query Parallelism”, “Stop-and-Go Operator”. Encyclopedia of Database Systems, Ling Liu and M. Tamer Özsu eds., 2008.
- R. Johnson, N. Hardavellas, I. Pandis, N. Mancheril, S. Harizopoulos, K. Sabirli, A. Ailamaki, and B. Falsafi. “To Share Or Not To Share?” In Proc. of the International Conference on Very Large Data Bases (**VLDB**), 2007.
- N. Hardavellas, I. Pandis, R. Johnson, N. Mancheril, S. Harizopoulos, A. Ailamaki, and B. Falsafi. “An Analysis of Database System Performance on Chip Multiprocessors.” In Proc. of the Hellenic Data Management Symposium (HDMS), 2007.
- N. Hardavellas, I. Pandis, R. Johnson, N. Mancheril, A. Ailamaki, and B. Falsafi. “Database Servers on Chip Multiprocessors: Limitations and Opportunities.” In Proc. of the Conference on Innovative Data Systems Research (**CIDR**), 2007.
- K. Gao, S. Harizopoulos, I. Pandis, V. Shkapenyuk, and A. Ailamaki. “Simultaneous Pipelining in QPipe: Exploiting Work Sharing Opportunities Across Queries.” In Proc. of the IEEE International Conference on Data Engineering (**ICDE**), 2006. **Best Demonstration Award**.
- I. Pandis. “Scalable transaction processing through data-oriented execution.” PhD Thesis, May 2011.

#### Middleware for Ubiquitous Computing:

- J. Soldatos, I. Pandis, K. Stamatis, L. Polymenakos, and J. L. Crowley. “Agent Based Middleware Infrastructure for Autonomous Context-Aware Ubiquitous Computing Services.” Journal of Computer Communications, Special Issue on Emerging Middleware for Next Generation Networks, Vol. 30(3), 2007.

- J. Soldatos, K. Stamatis, S. Azodolmolky, I. Pandis and L. Polymenakos. “Semantic web Technologies for Ubiquitous Computing Resource Management in Smart Spaces.” *International Journal of Web Engineering and Technology (IJWET)*, Vol. 3(4), 2007.
- I. Pandis, J. Soldatos, A. Paar, J. Reuter, M. Carras, and L. Polymenakos. “An Ontology-based Framework for Dynamic Resource Management in Ubiquitous Computing Environments.” In *Proc. of the International Conference on Embedded Software and Systems (ICESS)*, 2005.
- I. Pandis. “Context Aware Distributed Systems for Pervasive and Ubiquitous Computing.” MSIN Thesis, 2004.

### Hypermedia Systems:

- N. Karousos, I. Pandis, M. Tzagarakis, and M. Vaitis. “Supporting the Provision of Specialized Taxonomic Hypermedia Services to Web Applications.” In *Proc. of the HT06WS*, 2006.
- I. Pandis, N. Karousos, and T. Tiropanis. “Semantically Annotated Hypermedia Services.” In *Proc. of the ACM Hypertext Conference (HT)*, 2005.
- N. Karousos, and I. Pandis. “Developer Support in Open Hypermedia Systems: Towards a Hypermedia Service Discovery Mechanism.” In *Proc. of the Metainformatics International Symposium (MIS)*, 2003.
- N. Karousos, M. Tzagarakis, and I. Pandis. “Increasing the Usage of Open Hypermedia Systems: A Developer-Side Approach.” In *Proc. of the ACM Hypertext Conference (HT)*, 2003.
- N. Karousos, I. Pandis, S. Reich, and M. Tzagarakis. “Offering Hypermedia Services to the WWW: a step-by-step Approach for Developers.” In *Proc. of the International World Wide Web Conference (WWW)*, 2003.
- N. Karousos, I. Panaretou, I. Pandis, and M. Tzagarakis. “Babylon Bookmarks: A Taxonomic Approach to the Management of WWW Bookmarks.” In *Proc. of the Metainformatics International Symposium (MIS)*, 2002.
- I. Pandis. “Development of Taxonomic Hypertext Management System.” Diploma Thesis, September 2002.

### TUTORIALS

- A. Boehm, J. Dittrich, N. Mukherjee, I. Pandis, and R. Sen. “Operational Analytics Data Management Systems.” *VLDB 2016*.
- D. Abadi, S. Babu, F. Ozcan, and I. Pandis. “SQL-on-Hadoop Systems.” *VLDB 2015*.

### PANELS

- I. Pandis (coordinator). P. Boncz, R. Johnson, C. Kozyrakis, A. Pavlo, and E. Sedlar. “Database Machines 2.0: Doomed to fail (again) or Not?” *DaMoN 2014*.
- N. Tatbul, F. Ozcan (coordinators). D. Abadi, C. Mohan, I. Pandis, K. Ramasamy, and J. Wiener. “Are We Experiencing a Big Data Bubble?” *SIGMOD 2014*.
- M. Hsu (coordinators). A. Buchmann, S. Filkenstein C. Freytag, C. Mohan, I. Pandis, and T. B. Pedersen. “Real Time Analytics on Big Data.” *BIRTE 2013*.

### WORKING EXPERIENCE

10/2015 – Present **Amazon Web Services**

**Palo Alto, CA**

Principal Engineer at Amazon Web Services (AWS) working on Amazon Redshift. Spending most of the team on designing and implementing Redshift’s query execution engine. Architected and led the Amazon Redshift Spectrum feature that enables running Amazon Redshift SQL queries against exabytes of data in Amazon S3.

4/2014 – 10/2015    **Cloudera**    **San Francisco, CA**

Software Engineer on the Cloudera Impala team. I worked mostly on Impala's backend. Part of the team that implemented the partitioned and spillable hash joins and aggregations. Spent a lot of time making sure that queries are able to run successfully under memory pressure. Re-implemented Impala's hash joins, with an efficient linear probing implementation, providing up to 20% performance improvement on TPC-H and TPC-DS.

8/2011 – 4/2014    **IBM Almaden Research Center**    **San Jose, CA**

Research Staff Member of the Advanced Database Solutions (K55G) group. Actively involved on the BLU project. BLU is the column store (in) DB2 and was released as part of DB2 for Linux, Unix, Windows (LUW) v10.5. DB2 BLU accelerates read-mostly Business Intelligence queries against column-organized tables by 10 to 100 times and improves compression by one to two orders of magnitude, compared to traditional row-organized tables and execution.

1/2010 – 7/2011    **École Polytechnique Fédérale de Lausanne**    **Lausanne, Switzerland**

Member of the Data-Intensive Applications and Systems (DIAS) laboratory.

9/2005 – 5/2011    **Carnegie Mellon University**    **Pittsburgh, PA**

Research Assistant working on the StagedDB/CMP project, advised by Professor Anastasia Ailamaki. The StagedDB/CMP project studies the behavior of DBMSs in modern hardware environments and tries to enhance their performance by proposing novel designs in both software and hardware.

6/2009 – 9/2009    **Intel Research**    **Pittsburgh, PA**

Summer Fellow at the Intel Research Pittsburgh, PA lablet. Mentored by Dr. Shimin Chen.

6/2007 – 9/2007    **Microsoft Research**    **Redmond, WA**

Summer Intern at the Database Research Group, Redmond, WA. Mentored by Dr. Jingren Zhou.

1/2005 – 8/2005    **Athens Information Technology**    **Athens, Greece**

Member of the CHIL project development team. Mentored by Prof. Lazaros Polymenakos.

8/2002 - 8/2003    **Argo Systems**    **Athens, Greece**

Software developer at Argo Systems O.E., a software development company specialized in the development of shipping management systems.

## TEACHING

Spring 2007    Carnegie Mellon University. Teaching Assistant in *18-842: Distributed Systems*.  
Instructor: Prof. Gregory Ganger (ganger@ece.cmu.edu).

Fall 2006    Carnegie Mellon University. Teaching Assistant in *15-415: Database Applications*.  
Instructor: Prof. Christos Faloutsos (christos@cs.cmu.edu).

## LANGUAGES

Greek    Native

English    Fluent – Certificate of Proficiency, TOEFL: 287/300

German    Intermediate – Zertificat Deutsch als Fremdsprache (Grundstufe Sprachdiplom)

Last update: May 2017