

Rares Vernica

E-mail: rvernica@gmail.com Phone: +1-949-954-0195

LinkedIn: <https://www.linkedin.com/in/raresv> GitHub: <https://github.com/rvernica>

Work Experience

- Sep. 2011 - present HP Labs, HP Inc., Palo Alto, CA, USA
Senior Research Scientist - Print Adjacencies & 3D Lab
- Research, prototyping, deployment, and program management in large-scale data management, analytics, and content delivery.
- Built efficient data management stack for life science data analysis
 - Delivered cloud-based software infrastructure for new printer line
 - Analyzed customer behavior and built customer model for print subscription service
 - Built scalable image processing pipeline for book publishing. Deployed to final HP customer
 - Managed program, ran pilots, and led customer interactions for education vertical. Enhanced HP presence in the education market
 - Led development of web-based content ingestion and delivery platform
 - Developed automatic layout engine for web content delivery
 - Architected and deployed cloud-based programmable image analytics platform. Improved HP Cloud offering
- Jun. 2005 - Sep. 2011 University of California, Irvine, CA, USA
Graduate Student Researcher & Teaching Assistant - Department of Computer Science
- Research in the area of large-scale data management, in particular, efficient indexing and query processing when the data and the query conditions do not match exactly.
- Developed a highly scalable parallel platform for managing semi-structured data
 - Prototyped data integration, indexing, and search techniques to help finding people during and after a disaster. Deployed during 2010 Haiti earthquake
 - Built library to deal with errors and inconsistencies in information systems. Released as open source package
 - Developed and deployed web search for University of California, Irvine. Used by students and staff on daily basis
- Jun. 2010 - Dec. 2010 IBM Almaden Research Center, IBM, San Jose, CA, USA
Research Intern - Information Management Group
- Developed adaptive techniques for large-scale data processing
 - Extended Hadoop with adaptive data processing framework
 - Conducted performance studies and implemented optimizations for reducing response time
 - Transferred code into IBM's analytics platform

- Jun. 2007 - Microsoft Research, Microsoft, Redmond, CA, USA
 Sep. 2007 *Research Intern - Data Management, Exploration and Mining Group*
- Developed entity categorization techniques for large document collections
 - Conducted performance studies and implemented optimizations for improving result quality and reducing response time
- Jun. 2006 - Yahoo!, Sunnyvale, CA, USA
 Sep. 2006 *Technical Intern - Strategic Data Solutions*
- Developed testing framework for shared-nothing distributed pipeline
 - Implemented a testing framework using the TETware library
- Feb. 2005 - Transylvania University of Brasov, Romania
 Jun. 2005 *Lecturer - Department of Applied Informatics*
- Apr. 2004 - L3S Research Center, Hannover, Germany
 Aug. 2004 *Research Assistant*
- Developed publish/subscribe framework for peer-to-peer distributed systems
 - Implemented in Java using the JXTA library
- Feb. 2003 - Freelance Developer
 Jun. 2005
- Developed on-line games, web application, and modules for on-line scheduling, domain registration, and payment systems
 - Implemented in Java, Python, and PHP
- Jul. 2002 - Finsiel, Bucharest, Romania
 Feb. 2003 *Web Developer - Python, Zope*
- Oct. 2001 - Websoft, Bucharest, Romania
 Jul. 2002 *Web Developer - PHP, JavaScript*

Education

- 2011 University of California, Irvine, USA
 PhD in Computer Science
 Thesis: *Efficient Processing of Set-Similarity Joins on Large Clusters*
 Advisers: Prof. Michael J. Carey and Prof. Chen Li
- 2004 University Politehnica of Bucharest, Romania
 BSc in Computer Science and Engineering

Publications

1. *METIS: A Multi-faceted Hybrid Book Learning Platform*. Lei Liu, **Rares Vernica**, Tamir Hassan, Niranjana D. Venkata, Yang Lei, Jian Fan, Jerry Liu, Steven J. Simske, Shanchuan Wu. ACM SIGWEB, DocEng 2016.
2. *AERO: An Extensible Framework for Adaptive Web Layout Synthesis*. **Rares Vernica**, Niranjana D. Venkata. ACM SIGWEB, DocEng 2015.

3. *To Print or Not to Print: Hybrid Learning with METIS Learning Platform*. Joshua Hailpern, **Rares Vernica**, Molly Bullock, Udi Chatow, Jian Fan, Georgia Koutrika, Jerry Liu, Lei Liu, Steven J. Simske, Shanchan Wu. ACM SIGCHI, EICS 2015.
4. *AsterixDB: A Scalable, Open Source BDMS*. Sattam Alsubaiee, Yasser Altowim, Hotham Altwaijry, Alexander Behm, Vinayak R. Borkar, Yingyi Bu, Michael J. Carey, Inci Cetindil, Madhusudan Cheelangi, Khurram Faraaz, Eugenia Gabrielova, Raman Grover, Zachary Heilbron, Young-Seok Kim, Chen Li, Guangqiang Li, Ji Mahn Ok, Nicola Onose, Pouria Pirzadeh, Vassilis J. Tsotras, **Rares Vernica**, Jian Wen, Till Westmann. PVLDB 2014.
5. *Cloud Based Multimedia Analytic Platform*. Peng Wu, **Rares Vernica**, Qian Lin. ACM Multimedia 2013.
6. *Speeding Up Chemical Searches Using the Inverted Index: The Convergence of Chemoinformatics and Text Search Methods*. Ramzi Nasr, **Rares Vernica**, Chen Li, Pierre Baldi. Journal of Chemical Information and Modeling 2012.
7. *ASTERIX: An Open Source System for "Big Data" Management and Analysis*. Sattam Alsubaiee, Yasser Altowim, Hotham Altwaijry, Alexander Behm, Vinayak R. Borkar, Yingyi Bu, Michael J. Carey, Raman Grover, Zachary Heilbron, Young-Seok Kim, Chen Li, Nicola Onose, Pouria Pirzadeh, **Rares Vernica**, Jian Wen. PVLDB 2012.
8. *CIRCUMFLEX: A Scheduling Optimizer for MapReduce Workloads with Shared Scans*. Joel L. Wolf, Andrey Balmin, Deepak Rajan, Kirsten Hildrum, Rohit Khandekar, Sujay Parekh, Kun-Lung Wu, **Rares Vernica**. Operating Systems Review 2012.
9. *On the Optimization of Schedules for MapReduce Workloads in the Presence of Shared Scans*. Joel L. Wolf, Andrey Balmin, Deepak Rajan, Kirsten Hildrum, Rohit Khandekar, Sujay Parekh, Kun-Lung Wu, **Rares Vernica**. VLDB J.2012.
10. *Adaptive MapReduce using Situation-Aware Mappers*. **Rares Vernica**, Andrey Balmin, Kevin S. Beyer, Vuk Ercegovic. EDBT 2012.
11. *Adaptive Processing of User-Defined Aggregates in Jaql*. Andrey Balmin, Vuk Ercegovic, **Rares Vernica**, Kevin S. Beyer. IEEE Data Engineering Bulletin 2011.
12. *Efficient Processing of Set-Similarity Joins on Large Clusters*. **Rares Vernica**. PhD Dissertation, University of California, Irvine 2011.
13. *ASTERIX: Towards a Scalable, Semistructured Data Platform for Evolving World Models*. Alexander Behm, Vinayak R. Borkar, Michael J. Carey, Chen Li, Nicola Onose, **Rares Vernica**, Alin Deutsch, Yannis Papakonstantinou, Vassilis J. Tsotras. Journal of Distributed and Parallel Databases 2011.
14. *Hyracks: A Flexible and Extensible Foundation for Data-Intensive Computing*. Vinayak R. Borkar, Michael J. Carey, Raman Grover, Nicola Onose, **Rares Vernica**. ICDE 2011.
15. *AKYRA: Efficient Keyword-Query Cleaning in Relational Databases*. **Rares Vernica**, Chen Li. Technical Report, University of California, Irvine 2010.
16. *Efficient Parallel Set-Similarity Joins Using MapReduce*. **Rares Vernica**, Michael J. Carey, Chen Li. SIGMOD 2010.
17. *Efficient Top-k Algorithms for Fuzzy Search in String Collections*. **Rares Vernica**, Chen Li. KEYS 2009, Workshop on Keyword Search on Structured Data, collocated with SIGMOD 2009.
18. *Entity Categorization Over Large Document Collections*. Venkatesh Ganti, Arnd Christian König, **Rares Vernica**. KDD 2008.
19. *SEPIA: Estimating Selectivities of Approximate String Predicates in Large Databases*. Liang Jin, Chen Li, **Rares Vernica**. VLDB J. 2008.
20. *Relaxing Join and Selection Queries*. Nick Koudas, Chen Li, Anthony K. H. Tung, **Rares Vernica**. VLDB 2006.

Patents

1. *Print Intent Type* Georgia Koutrika, **Rares Vernica**. April 28, 2015.
2. *Adaptive Parallel Data Processing* Andrey Balmin, Kevin S. Beyer, Vuk Ercegovac, **Rares Vernica**. February 10, 2015.

Certifications

Scrum Product Owner	Scrum Alliance
Scrum Master	Scrum Alliance
SAFe Agilist	Scaled Agile, Inc.
Advanced Leader Bronze	Toastmasters International
Competent Communicator	Toastmasters International

Honors and Awards

Jan. 2013	HP Gold Innovator Award
Jan. 2012	Excellency Award, Student Association, Romania
Sep. 2010	IBM Tuition Award
Sep. 2010	Yahoo! Key Scientific Challenges Award - Web Information Management
Jul. 2009	Microsoft Student Travel Award
2005 - 2006	Chair's Fellowship, Department of Computer Science, UC Irvine
Sep. 2005	California Institute for Telecommunications and Information Technology (Calit ²) Fellowship
Feb. 2005	Second place, Pentalog Programming Contest, Romania
1999 - 2004	Merit-Based Scholarship, University Politehnica of Bucharest, Romania

Professional Services

2016	Judge: HP CodeWars Programming Competition
2015	Program committee member: VLDB
2014	Reviewer: Parallel Computing Journal
2011	External reviewer: VLDB
2010	External reviewer: ICDE, AAI
2009	External reviewer: EDBT, ICDM, IIWeb
2008	External reviewer: VLDB, ICDE, WWW, EDBT, CIKM, ICDL, CoopIS
2007	External reviewer: SIGMOD, KDD
2006	External reviewer: KDD, WAIM, APWeb

Hobbies

Kendo, All United States Kendo Federation, Rank 3-dan
Sport Pilot license