Mohammad Sadoghi Hamedani, Ph.D.

CONTACT INFORMATION

Research Staff Member, Database Group
IBM T.J. Watson Research Center
Yorktown Heights, New York, USA

Tel: (914) 319-7937
E-mail: msadoghi@us.ibm.com
http://researcher.ibm.com/person/us-msadoghi

RESEARCH INTERESTS

My research focuses on high-performance and extensible *Big Data Management Systems* in the context of designing novel data structures and (parallel) algorithms and utilizing *modern hardware advancements*, especially many-core processors, hardware accelerators (e.g., FPGAs and GPUs), and storage-class memories (e.g., flash and phase change memory). In particular, I am interested in *rethinking the foundation of relational database system design* for future hardware by reshaping the transaction and storage model to sustain the unprecedented scale of *data proliferation* and *heterogeneity* observed in the *Big Data* era.

EDUCATION

University of Toronto, Toronto, Canada

Ph.D., Computer Science, 2013
- Advisor: Hans-Arno Jacobsen
- **Thesis:** An Efficient, Extensible, Hardware-aware Indexing Kernel.

M.Sc., Computer Science, 2008
- Advisor: Renée J. Miller
- **Thesis:** A Framework for Computing the Overlap and Difference of Relational Schemas.

B.Sc., Computer Science, 2006
- Advisor: Nick Koudas
- **Thesis:** SPIDER: Data Integration and Data Quality.

REFEREED JOURNAL ARTICLES

31. M. Sadoghi, M. Jergler, H.-A. Jacobsen, R. Hull, R. Vaculin. Safe Distribution and Parallel Execution of Data-centric Workflows over the Publish/Subscribe Abstraction. *Transactions on Knowledge and Data Engineering*, ?:?, 2015.
30. M. Sadoghi and H.-A. Jacobsen. Analysis and Optimization for Boolean Expression Indexing. *ACM Trans. Database Syst.*, 38(2), pages 8:1–8:47, 2013.

REFEREED PAPERS IN CONFERENCE PROCEEDINGS

29. M. Najafi, M. Sadoghi, H.-A. Jacobsen. Configurable Hardware-based Streaming Architecture using Online Programmable-Blocks. In *IEEE* 30^{th} *International Conference on Data Engineering (ICDE 2015), Seoul, Korea, April 13-16, 2015*, pages ?–?, 2015.
28. P. Menon, R. Rabl, M. Sadoghi and H.-A. Jacobsen. Optimizing Key-Value Stores for Hybrid Storage Architectures. In 24^{th} International Conference of Center for Advanced Studies on Collaborative Research, IBM CASCON 2014, Toronto, Canada, November 3-5, 2014.
27. M. Sadoghi, M. Canim, B. Bhattacharjee, F. Nagel, K. Ross. Reducing Database Locking Contention Through Multi-version Concurrency. In 40^{th} *International Conference on Very Large Data Bases, PVLDB 2014, Hangzhou, China, September 1-5, 2014.*
26. P. Menon, R. Rabl, M. Sadoghi and H.-A. Jacobsen. Adaptive Parallel Compressed Event Matching. In *IEEE* 30^{th} *International Conference on Data Engineering (ICDE 2014), Chicago, IL, USA, March 31 - April 4, 2012*, pages 364-375, 2014.
25. M. Sadoghi and H.-A. Jacobsen. CaSSanDra: An SSD Boosted Key-Value Store. In *IEEE* 30^{th} *International Conference on Data Engineering (ICDE 2014), Chicago, IL, USA, March 31 - April 4, 2012*, pages 1162-1167, 2014.

24. M. Sadoghi, K. Ross, M. Canim, B. Bhattacharjee. Making Updates Disk-I/O Friendly Using SSDs. In 39^{th} *International Conference on Very Large Data Bases, PVLDB 2013, Riva del Garda, Trento, Italy, August 26-30, 2013.*

23. K. Zhang, M. Sadoghi, V. Muthusamy, H.-A. Jacobsen. Distributed Ranked Data Dissemination in Social Networks. In *IEEE* 33^{th} *International Conference on Distributed Computing Systems (ICDCS 2013), Philadelphia, PA, USA, 8-11 July, 2013.*

22. H.-A. Jacobsen, K. Mokhtarian, T. Rabl, M. Sadoghi, R. S. Kazemzadeh, Y. Yoon, and K. Zhang. Grand Challenge: The Bluebay Soccer Monitoring Engine. In *The* 7^{th} *ACM International Conference on Distributed Event-Based Systems, DEBS '13, Arlington, TX, USA - June 29 - July 03, 2013*, pages 295–300. ACM, 2013.

21. T. Rabl, M. Sadoghi, H.-A. Jacobsen, S. Gómez-Villamor, V. Muntés-Mulero, and S. Mankowskii. Solving Big Data Challenges for Enterprise Application Performance Management. In 38^{th} *International Conference on Very Large Data Bases, PVLDB 2012, Istanbul, Turkey, August 27-31, 2012*, pages 5(12):1724–1735, 2012.

20. T. Rabl, K. Zhang, M. Sadoghi, N. K. Pandey, A. Nigam, C. Wang, and H.-A. Jacobsen. Solving Manufacturing Equipment Monitoring through Efficient Complex Event Processing: DEBS Grand Challenge. In *Proceedings of the* 6^{th} *ACM International Conference on Distributed Event-Based Systems, DEBS 2012, Berlin, Germany, July 16-20, 2012*, pages 335–340, 2012. – Best DEBS Challenge Award - Public Voting, 2012

19. M. Sadoghi and H.-A. Jacobsen. Relevance Matters: Capitalizing on Less (Top-k Matching in Publish/Subscribe). In *IEEE* 28^{th} *International Conference on Data Engineering (ICDE 2012), Washington, DC, USA (Arlington, Virginia), 1-5 April, 2012*, pages 786–797, 2012.

18. M. Sadoghi and H.-A. Jacobsen. BE-Tree: An Index Structure to Efficiently Match Boolean Expressions over High-dimensional Discrete Space. In *Proceedings of the* 37^{th} *ACM SIGMOD International Conference on Management of Data, SIGMOD 2011, Athens, Greece, June 12-16, 2011*, pages 637–648, 2011. – Winner of EPTS Innovative Principles Award 2011

17. M. Sadoghi, I. Burcea, and H.-A. Jacobsen. GPX-Matcher: A Generic Boolean Predicate-based XPath Expression Matcher. In *EDBT 2011,* 14^{th} *International Conference on Extending Database Technology, Uppsala, Sweden, March 21-24, 2011, Proceedings*, pages 45–56, 2011. – Selected for ACM TODS Special Issue on Best Papers of EDBT 2011

16. A. Farroukh, M. Sadoghi, and H.-A. Jacobsen. Towards Vulnerability-based Intrusion Detection with Event Processing. In *Proceedings of the* 5^{th} *ACM International Conference on Distributed Event-Based Systems, DEBS 2011, New York, NY, USA, July 11-15, 2011*, pages 171–182, 2011.

15. A. Chandel, O. Hassanzadeh, N. Koudas, M. Sadoghi, and D. Srivastava. Benchmarking Declarative Approximate Selection Predicates. In *Proceedings of the* 33^{rd} *ACM SIGMOD International Conference on Management of Data, Beijing, China, June 12-14, 2007*, pages 353–364, 2007.

Refereed Conference Demo Papers

14. M. Jergler, M. Sadoghi and H.-A. Jacobsen. D^2WORM: A Management Infrastructure for Distributed Data-centric Workflows. In *Proceedings of the* 41^{th} *ACM SIGMOD International Conference on Management of Data, SIGMOD 2015, Melbourne, Australia, May 31-4, 2015*, pages ?–?, 2015.

13. M. Najafi, M. Sadoghi, H.-A. Jacobsen. Flexible Query Processor on FPGAs. In 39^{th} *International Conference on Very Large Data Bases, PVLDB 2013, Riva del Garda, Trento, Italy, August 26-30, 2013*, pages 1310-1313, 2015.

12. T. Rabl, M. Sadoghi, K. Zhang, and H.-A. Jacobsen. MADES - A Multi-layered, Adaptive, Distributed Event Store. In *The* 7^{th} *ACM International Conference on Distributed Event-Based Systems, DEBS '13, Arlington, TX, USA - June 29 - July 03, 2013*, pages 343–344. ACM, 2013.

11. M. Sadoghi, R. Javed, N. Tarafdar, H. Singh, R. Palaniappan, and H.-A. Jacobsen. Multi-query Stream Processing on FPGAs. In *IEEE 28th International Conference on Data Engineering (ICDE 2012), Washington, DC, USA (Arlington, Virginia), 1-5 April, 2012*, pages 1229–1232, 2012.

10. V. Muthusamy, Y. Yoon, M. Sadoghi, and H.-A. Jacobsen. eqosystem: Supporting Fluid Distributed Service-oriented Workflows. In *Proceedings of the 5th ACM International Conference on Distributed Event-Based Systems, DEBS 2011, New York, NY, USA, July 11-15, 2011*, pages 381–382, 2011.

9. M. Sadoghi, H. Singh, and H.-A. Jacobsen. fpga-topss: Line-speed Event Processing on FPGAs. In *Proceedings of the 5th ACM International Conference on Distributed Event-Based Systems, DEBS 2011, New York, NY, USA, July 11-15, 2011*, pages 373–374, 2011.

8. M. Sadoghi, H.-A. Jacobsen, M. Labrecque, W. Shum, and H. Singh. Efficient Event Processing through Reconfigurable Hardware for Algorithmic Trading. In 36th *International Conference on Very Large Data Bases, PVLDB 2010, Singapore, Singapore, September 13-17, 2010*, pages 3(2):1525–1528, 2010.

Refereed Workshop Papers

7. M. Sadoghi, K. Ross, M. Canim, B. Bhattacharjee. Making Updates Disk-I/O Friendly Using SSDs. In 5th *Annual Non-Volatile Memories Workshop 2014, San Diego, California, USA, March 9-11, 2014.*

6. M. Sadoghi. Towards an Extensible Efficient Event Processing Kernel. In *Proceedings of the 38th ACM SIGMOD/PODS PhD Symposium 2012, Scottsdale, AZ, USA, May 20, 2012*, pages 3–8, 2012.

5. M. Sadoghi, H. Singh, and H.-A. Jacobsen. Towards Highly Parallel Event Processing through Reconfigurable Hardware. In *Proceedings of the 7th International Workshop on Data Management on New Hardware (Collocated with ACM SIGMOD), DaMoN 2011, Athens, Greece, June 13, 2011*, pages 27–32, 2011.

4. O. Hassanzadeh, M. Sadoghi, and R. J. Miller. Accuracy of Approximate String Joins using Grams. In *Proceedings of the 5th International Workshop on Quality in Databases, QDB 2007, at the VLDB 2007 conference, Vienna, Austria, September 23, 2007*, pages 11–18, 2007.

Other Publications

3. M. Sadoghi, M. Labrecque, H. Singh, W. Shum, and H.-A. Jacobsen. fpga-ToPSS for High-throughput and Low-Latency Event Processing in Algorithmic Trading. *IBM 20th International Conference of Center for Advanced Studies on Collaborative Research 2010*

2. M. Sadoghi, A Framework for Computing the Overlap and Difference of Relational Schemas. *Technical Report CSRG-609, University of Toronto 2008*

1. M. Sadoghi, and N. Koudas. SPIDER: Data Quality & Data Cleaning Project. *In 15th International Conference of Center for Advanced Studies on Collaborative Research, IBM CASCON 2005*

Filed Patents

18. M. Sadoghi, M. Canim, B. Bhattacharjee, K. Ross. AN EFFICIENT METHOD FOR TIME TO LIVE AND ACCESS FREQUENCY APPROXIMATION FOR REMOTE DIRECT MEMORY ACCESS OVER DISTRIBUTED SHARED MEMORY. *2015*

17. R. Barber, B. Bhattacharjee, G. Lohman, C. Mohan, I. PANDIS, V. Raman, M. Sadoghi, Sidle, A. Storm. IN-PLACE UPDATES WITH CONCURRENT READS IN A DECOMPOSED STATE. *2015*

16. M. Sadoghi, R. Barber, B. Bhattacharjee, G. Lohman, C. Mohan, I. PANDIS, V. Raman, R. Sidle, A. Storm. AN EFFICIENT METHOD TO SUPPORT UPDATABLE COLUMN STORE WITH FLEXIBLE SCHEMA THROUGH BATCHED, APPEND-ONLY, RANGE-PARTITIONING SCHEME. *2015*

15. M. Sadoghi, Y. Chang, T. Malkemus. AN EFFICIENT METHOD FOR DATA SKIPPING AND COMPRESSION THROUGH VERTICAL BYTE PARTITIONING OF DATA. *2015*
14. M. Sadoghi, M. Canim, B. Bhattacharjee, B. Makni, K. Ross. ADAPTIVE CONCURRENCY CONTROL USING HARDWARE TRANSACTIONAL MEMORY AND LOCKING MECHANISM. *2015*
13. M. Sadoghi, M. Canim, B. Bhattacharjee. ACCELERATING MULTIVERSION CONCURRENCY CONTROL USING HARDWARE TRANSACTIONAL MEMORY. *2015*
12. R. Barber, B. Bhattacharjee, G. Lohman, C. Mohan, I. PANDIS, V. Raman, M. Sadoghi, Sidle, A. Storm. POOLING WORK ACROSS MULTIPLE TRANSACTIONS FOR REDUCING CONTENTION IN OPERATIONAL ANALYTICS SYSTEMS. *2015*
11. R. Barber, B. Bhattacharjee, G. Lohman, C. Mohan, I. PANDIS, V. Raman, M. Sadoghi, Sidle, A. Storm. A EFFICIENT METHOD TO DO INSERTS AND POINT QUERIES IN A COLUMN STORE. *2015*
10. M. Sadoghi, M. Canim, B. Bhattacharjee. IN-MEMORY LATCH-FREE INDEX STRUCTURE. *2015*
9. M. Sadoghi, M. Canim, B. Bhattacharjee. RENDEZVOUS-BASED OPTIMISTIC CONCURRENCY CONTROL. *2015*
8. M. Sadoghi, M. Canim, B. Bhattacharjee, K. Ross, A. Storm. SUPPORTING TRANSIENT SNAPSHOT WITH COORDINATED/UNCOORDINATED COMMIT PROTOCOL. *2014*
7. M. Sadoghi, Y. Chang, O. HASSANZADEH, T. Malkemus. DATABASE QUERY PROCESSING USING HORIZONTAL DATA RECORD ALIGNMENT OF MULTI-COLUMN RANGE SUMMARIES. *2014*
6. M. Canim., M. Sadoghi, B. Bhattacharjee, K. Ross. PREPLAYING TRANSACTIONS THAT MIX HOT AND COLD DATA. *2014*
5. M. Sadoghi, A-H. Jacobsen, T. Rabl. LOG DATA STORE THAT STORES DATA ACROSS A PLURALITY OF STORAGE DEVICES USING NON-DISJOINT LAYERS. *2014*
4. M. Sadoghi, B. Bhattacharjee, K. Ross, M. Canim., F. Nagel. REDUCING DATABASE LOCKING CONTENTION USING MULTI-VERSION DATA RECORD CONCURRENCY CONTROL. *2013*
3. M. Sadoghi, B. Bhattacharjee, K. Ross, M. Canim. DEFERRING DATA RECORD CHANGES USING QUERY REWRITING. *2013*
2. M. Sadoghi, B. Bhattacharjee, K. Ross, M. Canim. COMPRESSING A MULTI-VERSION DATABASE. *2013*
1. M. Sadoghi, B. Bhattacharjee, K. Ross, M. Canim. MANAGING MULTI-VERSION DATABASES. *2013*

Presentations

Invited Talks

4. The Role of Modern Hardware in Big Data Processing. Samsung. Suwon, Korea. *April 2015*
3. The Role of Modern Hardware in Big Data Processing. Teradata Aster. California, USA. *October 2012*
2. The Role of Modern Hardware in Big Data Processing. IBM Thomas J. Watson Research Center. New York, USA. *September 2012*
1. BE-Tree: Boolean Expression Tree. IBM Thomas J. Watson Research Center. New York, USA. *July 2011.*

Invited Lectures

1. *Data Warehousing and Data Mining.* University of Connecticut. *UCONN STAMFORD EMBA Program 2011*

Honors and Awards

IBM 4^{th} Invention-Plateau Award	**2015**
IBM 3^{rd} Invention-Plateau Award	**2015**
IBM 2^{nd} Invention-Plateau Award	**2014**
IBM Outstanding Contribution Award	**2014**
IBM 1^{st} Invention-Plateau Award	**2013**
Best DEBS Challenge Award - Public Voting	**2012**
EPTS Innovative Principles Award (BE-Tree: Boolean Expression-Tree [SIGMOD'11])	**2011**
NSERC Postgraduate Scholarships (PGS D) – Comparable to NSF Fellowship	**2009 - 2011**
NSERC Postgraduate Scholarships (PGS M) – Comparable to NSF Fellowship	**2007 - 2008**
Ontario Graduate Scholarship (OGS)	**2006 - 2007**
Natural Sciences and Engineering Research Council (NSERC-USRA)	**2004 - 2006**
Golden Key International Honour Society	**2002 - 2006**

Professional Service

Organizing Committee

- ACM International Conference on Distributed Event-Based Systems (DEBS) **2015**
  (Publicity Co-Chairs)
- IBM Workshop on Big Data Analytics **2013**
  (Organizing Committee)

Program Committee

- International Conference on Distributed Computing Systems (ICDCS) **2015**
  (Research Track)
- International ACM/IFIP/USENIX Middleware Conference (Middleware) **2015**
  (Industry Track)
- International Conference of the Center for Advanced Studies (IBM CASCON) **2015**
  (Research Track)
- International Conference on Service Oriented Computing (ICSOC) **2015**
  (Research Track)
- International Workshop on Accelerating Data Management Systems Using Modern Processor and Storage Architectures (ADMS) **2015**
  (Research Track)
- CSI International Symposium on Computer Architecture & Digital Systems (CADS) **2015**
  (Research Track)
- ACM International Conference on Distributed Event-Based Systems (DEBS) **2014**
  (Research Track)
- International Conference of the Center for Advanced Studies (IBM CASCON) **2014**
  (Research Track)
- European Conference on Object-Oriented Programming (ECOOP) **2013**
  (Artifact Evaluation Track)
- International Conference on Service Oriented Computing (ICSOC) **2013**
  (Demonstration Track)
- Big Data Benchmarking (WBDB) **2013**
  (Research Track)

Invited Reviewer

- Elsevier Information Systems Journal (IS) **2015**
- Transactions on Knowledge and Data Engineering (TKDE) **2015**
- The Computer Journal (Oxford Press) **2015**

- Transactions on Knowledge and Data Engineering (TKDE) **2014**
- Transactions on Computers (TC) **2014**
- International Conference on Extending Database Technology (EDBT) **2014**
- Journal of Computer Science and Technology (JCST) **2014**
- International Workshop on Accelerating Data Management Systems Using Modern Processor and Storage Architectures (ADMS) **2014**
- International Conference on Information Technology for Organization Development (IT4OD) **2014**
- Science of Computer Programming Journal (SCP) **2013**
- IEEE Transactions on Services Computing (TSC) **2012**
- Journal of Communications and Networks (JCN) **2012**

External Reviewer

- International Conference on Distributed Computing Systems (ICDCS) **2013**
- International Conference on Very Large Databases (VLDB) **2013**
- Transactions on Database Systems (TODS) **2012**
- International Middleware Conference (ACM/IFIP/USENIX Middleware) **2012**
- IEICE Transactions **2012**
- International Conference on Distributed Computing Systems (ICDCS) **2012**
- International Conference on Very Large Databases (VLDB) **2012**
- IEEE Transactions on Parallel and Distributed Systems (TPDS) **2011**
- International Journal of Computing **2011**
- International Middleware Conference (ACM/IFIP/USENIX Middleware) **2011**
- International Conference on Distributed Computing Systems (ICDCS) **2011**
- International Conference on Distributed Event-Based Systems (ACM DEBS) **2011**
- International Conference on Distributed Computing Systems (ICDCS) **2010**
- International Conference of the Center for Advanced Studies (IBM CASCON) **2010**

Professional Research Experience

IBM T.J. Watson Research Center, New York, USA

Research Intern (Advisor: Kenneth A. Ross) **May 2012 - August 2012**

- Enhancing database management systems using SSDs (i.e., reducing index maintenance in operational data stores). The list of patents filed:
  - Method and system for accelerating the speed of insert, updates and deletes in multi-version databases
  - Efficient storage and retrieval in multi-version databases using solid-state storage

IBM T.J. Watson Research Center, New York, USA

Research Intern (Advisor: Richard Hull) **June 2011 - August 2011**

- Formalized the semantics of distributed and de-coupled execution of IBM Business Entity with Life-cycle (BEL-GSM) through publish/subscribe primitives
- Extended Padres, an enterprise-grade event management platform (an open-source research project), to execute IBM BEL-GSM model

Academic Experience

University Collaborations at IBM T.J. Watson Research Center, New York, USA

Mentoring

- Bassem Makni (PhD Candidate at Rensselaer Polytechnic Institute (RPI), Research Intern) **2014**
  Focus: Optimistic Concurrency Control using Hardware Transactional Memory (patents pending).

- Ehab Abdelhamid (PhD Candidate at King Abdullah University of Science and Technology (KAUST), Research Intern) **2014**
  Focus: Incremental Frequent Subgraph Mining on Large Dynamic Graphs (patents pending).

- Fabian Nagel (PhD Candidate at University of Edinburgh, Research Intern) **2013**
  Focus: Optimistic Concurrency Control (patents pending, published in PVLDB 2014).

- Tilmann Rabl (Post-Doctoral Fellow at University of Toronto, IBM CAS Fellowship) **2013-present**
  Focus: Write-intensive data structures (patents pending, published in IBM ICDE 2014 and CASCON 2014).

Technische Universitat Munchen, Munich, Germany

Advising (Co-advisor Hans-Arno Jacobsen)

- Mohammadreza Najafi (PhD) **2012-present**
  Focus: Flexible Query Processor on FPGAs (published in PVLDB 2013, ICDE 2015).

- Martin Jergler (PhD) **2012-present**
  Focus: Data-centric Workflow Execution on Publish/Subscribe Abstraction (published in SIGMOD 2015).

University of Toronto, Toronto, Canada

Advising (Co-advisor Hans-Arno Jacobsen)

- Kaiwen Zhang (PhD) **2012-present**
  Focus: Distributed Ranked Data Dissemination in Social Networks (published in ICDCS 2013).

- Prashanth Menon (MSc) **2013-present**
  Focus: CaSSanDra: An SSD Boosted Key-Value Store (published in ICDE 2014).

- Rija Javed (MSc), Naif Tarafdar (BSc), and Rohan Palaniappan (BSc) **2011-2012**
  Focus: Multi-Query Stream Processing on FPGAs (published in ICDE 2012).

- Harsh Singh (MSc) **2010-2012**
  Focus:

  i Towards Highly Parallel Event Processing through Reconfigurable Hardware (published in DaMoN 2011).

  ii Efficient Event Processing through Reconfigurable Hardware for Algorithmic Trading (published in PVLDB 2010).

- Amer Farroukh (MSc) **2010-2011**
  Focus: Towards Vulnerability-Based Intrusion Detection with Event Processing (published in DEBS 2011).

Instructor
- APS105 - Computer Fundamentals **2012**

Teaching Assistant
- *Graduate Courses*
  - CSC2531 Advanced Topics in Data Management Systems **2011**
- *Undergraduate Courses*
  - CSC443 Database System Technology **2007**
  - CSC343 Introduction to Databases **2007, 2008, 2011, 2012**
  - CSC309 Web Programming **2008**
  - ECE297 Communication and Design **2010, 2011, 2012**
  - CSC263 Data Structures and Analysis **2010**
  - CSC236 Introduction to the theory of Computation **2010**
  - CSC209 Software Tools and Systems Programming **2009**
  - MAT135 Calculus II for Biological Sciences **2005**
  - MAT130 Calculus I **2005**
  - MAT127 Introduction to Optimization **2006**
  - CSC165 Mathematical Expression and Reasoning for Computer Science **2009, 2010**
  - CSC120 Introduction to Computer Science **2008**
  - CSC108 Introduction to Computer Programming **2003, 2004, 2007**
  - HLP101 Help Center TA (focused on all undergraduate computer science courses) **2010**

Professional Experience

ING Canada, Toronto, Canada

Senior Developer/Architect **July 2008 - April 2009**
- Prototyped, for business as a proof of concept/demo, a web based document management client (a thin client) to replace the vendor supplied thick client solution.
- Designed an ECM Java framework as part of a pillar strengthening strategy to become a center of excellence (COE) for ECM within the organization and provide eForms services. This enabled automatic/manual population of forms and storage within the content manager system. Also provided distribution services via fax, print and email for the pre-existing content and the generated eForms document.

IBM, Toronto, Canada

DB2 Tools Developer **May 2005 - September 2005**
- Implemented new functionalities for DB2 Buffer Pool Management and Schema components
- Investigated and implemented a fix for the DB2 Administrative Tools Center defects
- Maintained and administrated DB2 build using OnDemand regression environment

Compiler Specialist **January 2005 - April 2005**

- Designed and implemented Scenario Builder Tool using Object Oriented Perl which fully integrated into the current C/C++ Compiler Validation environment and offered an automatic mechanism to implement the test scenarios
- Analysed GCC testing environment and proposed and implemented a solution to automatically migrate IBM & GNU GCC environment which significantly enhanced the testing coverage.
- Analysed the requirements and designed the testing methodology for the new C/C++ features. Investigated and identified C/C++ Compiler bugs and provided in-depth defect report which offered possible solution to the problem

Tallán, New York, USA

Software Developer **May 2004 - September 2004**

- Developed the foundation and core functionality such as security, logging, emailing and database access components for the new intranet system
- Analyzed in-depth the current solution and gathered necessary requirements to design and develop both front-end and business tier of the Human Resources module
- Constructed an abstract auto-generated data model for the entire intranet project which extensively relied on Object Oriented design and had similar architecture to strongly typed DataSet in ADO.NET
- Designed and built data tier for the security and Human Resources module in SQL Server 2000
- Evaluated technology such as Active Directory for the security model of the intranet
- Integrated the custom application into Microsoft SharePoint Portal Server 2003

Alliance Atlantis Communications, Toronto, Canada

Software Developer **January 2003 - April 2003**

- Thoroughly analyzed the current web-based project management and proposed a new solution
- Developed Events & Scheduled Messenger applications using ASP.NET and SQL Server 2000
- Used VBScript to synchronize events and meetings between Microsoft Outlook and a web-based application
- Implemented a Windows Services to manage sending emails and scheduling tasks in .NET Platform

Certification

IBM Training

IBM Certified Database Administrator - DB2 UDB V8.1 for Linux, UNIX and Windows **2005**
IBM Certified Application Developer DB2 Universal Database V8.1 Family **2005**
IBM Certified Solution Designer - DB2 Business Intelligence V8 **2005**

Microsoft Training

Microsoft Certified Application Developer (MCAD) **2004**
Certified Web Applications Developer - Microsoft Visual C# .NET **2004**
Certified Windows-based Applications Developer - Microsoft Visual C# .NET **2004**
Certified XML Web Services and Server Components Developer - Microsoft Visual C# **2004**
Certified Databases Designer - Microsoft SQL Server 2000 Enterprise Edition **2004**

Technical Skills

- **Programming Languages**
  ASP .NET, C, C++, C#, Java, Perl, Python, Scheme, Prolog, and OpenMP
- **Java Technologies**
  J2EE (JMS, JSP, JSF Servlets, JNDI, and JDBC), Web Services (WSDL and SOAP), Hibernate, multi-threading, Internationalization, RMI, and JavaBeans
- **Internet Technologies**
  J2EE, WSDL, XML, XSLT, Apache Struts, Spring, Spring MVC, Spring Web Flow 2, CGI, HTML, DHTML, AJAX, JavaScript, MooTools, CSS, and .NET framework
- **Application/Web Servers**
  WebSphere Application Server, Apache Tomcat, Apache HTTP Server, and Microsoft IIS
- **Communications**
  J2EE, Web Services, Java RMI, TCP/IP, and n-tier Client/Server Architecture
- **Database (and NoSQL) Technologies**
  IBM DB2, DB2 BLU, Orcale, Microsoft SQL Server, Microsoft Hekaton, SAP HANA, MySQL, PostgresSQL, PointBase, H-Store, C-Store, HBase, Cassandra, VoltDB, Redis, Voldemort, LevelDB, NuoDB, Pregel (Giraph), Neo4j, DB2RDF, MapReduce (Hadoop), and Spark
- **Content Management Systems**
  IBM DB2 Content Manager
- **Operating Systems**
  SUN Solaris, IBM AIX, Linux, Mac OS X, MS Windows 7/Vista/2003 Server/XP/2000/NT/98/95/3.1, and MS-DOS
- **IDEs**
  Eclipse, Rational Application Developer (i.e. RAD) and IntelliJ IDEA, Visual Studio .NET
- **Source Control**
  ClearCase & ClearQuest, SourceSafe, SubVersion, CVS, and CMVS
- **Miscellaneous**
  Log4J, Apache Ant, Make, and JUnit

References

Hans-Arno Jacobsen (Research)
Professor
Dept. of Computer Science
Dept. of Electrical and Computer Engineering
University of Toronto, Canada
Email: jacobsen@eecg.toronto.edu

Fahiem Bacchus (Research)
Professor
Dept. of Computer Science
University of Toronto, Canada
Email: fbacchus@cs.toronto.edu

Ryan Johnson (Research)
Assistant Professor
Dept. of Computer Science
University of Toronto, Canada
Email: ryan.johnson@cs.toronto.edu

Kenneth A. Ross (Research)
Professor
Dept. of Computer Science
Columbia University, United States
IBM T.J. Watson, United States
Email: kar@cs.columbia.edu

Rick Hull (Research)
Research Staff Member and Manager
IBM T.J. Watson, United States
Email: hull@us.ibm.com

Yuan-Chi Chang (Research)
Research Staff Member and Manager
IBM T.J. Watson, United States
Email: yuanchi@us.ibm.com