


2020 BOARD NOMINEE

Fatma Ozcan

Principal Research Staff Member and Senior Manager, Cognitive Solutions
IBM Research



Awards and Honors and Year Received

- ACM Distinguished Member, 2019
- IBM Corporate Award for PureXML Database Technology, 2008
- IBM Extraordinary Technical Accomplishment for „Research Contributions to DB2“, 2010
- IBM Extraordinary Technical Accomplishment for „Standards Leadership: The Evolution from XML to Web Services to SOA“, 2009
- IBM Outstanding Technical Achievement Award for Join Algorithms for Big Data, 2016

Involvement in CRA Activities

- Although I have not participated in CRA workshops and other events directly, I have been following CRA publications, and using them at work as well. For example, I have been using CRA and ACM resources in training and mentoring my colleagues in diversity.

Other Relevant Experience

- ACM SIGMOD Treasurer and Secretary, part of Executive Committee, since 2017
- Board of trustees of the VLDB Endowment, since 2016 (6-year term)
- Strategy leader on data servers at IBM, working on data servers research agenda
- Editor in chief of PVLDB Volume 12, PC Chair of VLDB 2019
- General chair of SIGMOD 2016

Research Interests

Infra-structure and platforms for large-scale data analysis, conversational interfaces to data and NLQ, semi-structured data processing, query processing and optimization

2020 BOARD NOMINEE

Fatma Ozcan

Principal Research Staff Member and Senior Manager, Cognitive Solutions
IBM Research



CRA
Computing Research
Association

Personal Statement

She has over 18 years of experience in industrial research, and delivered core technologies into many IBM products. She currently leads the data servers research strategy for IBM research. She is the treasurer/secretary of ACM SIGMOD, and is on the board of trustees for the VLDB Endowment. She is an ACM Distinguished Member. She has co-authored many scholarly papers and patents. She was the general chair for SIGMOD 2016, and the PC chair of VLDB 2019. She serves on program committees of leading database conferences, and editorial boards of journals, as well as NSF panels.

Fatma ÖZCAN

Address:

1120 Mallard Ridge Loop
San Jose, CA 95120

Cell : (408) 250 4757

Email : fatma.ozc@gmail.com

TECHNICAL INTERESTS:

Infra-structure and platforms for large-scale data analysis, conversational interfaces to data and NLQ, SQL-on Hadoop, semi-structured data processing, query processing and optimization, and query languages.

EDUCATION

- August 2001** *Ph.D. in Computer Science.....University of Maryland, College Park*
Dissertation: Improving the Performance of Heterogeneous Databases and Agents
Advisor: Prof. V.S Subrahmanian
GPA: 3.90/4.00
- August 1999** *MS in Computer Science.....University of Maryland, College Park*
GPA: 3.90 /4.00
- January 1996** *MS in Computer Engineering.....Middle East Technical University, Ankara, Turkey*
Thesis: Dynamic Query Optimization on a Distributed Object Management Platform
Advisor: Prof Asuman Dogac. Doç
GPA: 4.00 / 4.00
- July 1994** *BS in Computer Engineering.....Middle East Technical University, Ankara,, Turkey*
GPA: 3.87 / 4.00 (First in rank)

WORK EXPERIENCE

July 2018	<i>Principal RSM and Senior Manager IBM Almaden Research Center, San Jose, CA</i> Leading a group of researchers for exploring, prototyping, and delivering emerging architectures and solutions for information management systems. Primarily focusing on large scale cloud environments and services, including HTAP systems, conversational interfaces to data and NLQ, scalable machine learning, and big data platforms. She is also a research strategist for the area
------------------	---

	<p>of AI for data, looking into ways of improving data servers using AI, as well as building next generation data platforms to support AI applications. In this role, she coordinates research across IBM labs, particularly in India and Almaden.</p>
<p>April 2016</p>	<p><i>Principal Research Staff Member</i> <i>IBM Almaden Research Center, San Jose, CA</i> Projects: <i>Wildfire, NLQ</i> Significant Accomplishments:</p> <ul style="list-style-type: none"> • <i>She leads the database research group at IBM Research – Almaden.</i> • <i>She leads the design and development of ontology based querying of knowledge bases, and existing data sets. She is leading the design and development of an ontology based query language that abstracts the storage and indexing of data, abstracting the physical layer from the semantic data model. This language is used by NLQ (Natural language query) to connect to multiple back-ends, including relational as well as elastic search. This abstraction and NLQ service is also connected to conversational interfaces.</i> • <i>She leads the research team in the joint work with Watson Healthcare to build a conversational interface, leveraging NLQ stack, for a drug knowledge base</i> • <i>She leads the research team working together with the analytics group to build NLQ over enterprise data warehouses</i> • <i>She co-leads the wildfire project, which aims to build a hybrid transactional and analytical processing engine in Spark.</i> • <i>She is a member of the research strategy committee, responsible for, providing technical leadership on data and query research across IBM research labs</i>
<p>September 2001 to March 2016</p>	<p><i>Research Staff Member</i> <i>IBM Almaden Research Center, San Jose, CA</i> Projects: <i>DB2 XML, SEDA, XAP/BigInsights, Hadoop/Warehouse Integration, SQL on Hadoop data (Big SQL), Next Generation Analytics Platform, Watson Content Services, Next Generation Analytics Platform</i></p> <p>Significant Accomplishments:</p> <ul style="list-style-type: none"> • <i>She has been a manager since 2014, managing a team of 5-7 researchers.</i> • <i>She was one of the main architects for Big SQL 3.0, a standards compliant SQL offering over HDFS data, part of IBM BigInsights</i> • <i>She co-designed and developed the SQL/JSON, SQL PTF (Polymorphic parallel table functions) standards, together with IBM SQL standards representatives, providing technical consultancy</i> • <i>She led, designed, and delivered DB2/Hadoop integration, which includes connectors to communicate computation and data between</i>

	<p><i>IBM Infosphere BigInsights and IBM Infosphere Warehouse</i></p> <ul style="list-style-type: none"> • <i>She designed, architected, and led the development of OLAP-style complex querying over heterogeneous XML data within the SEDA system</i> • <i>She innovated, designed, developed and delivered XQuery semantics and several rewrite transformations into DB2 pureXML version 9 and 9.5. She was one of the main architects of the XQuery and SQL/XML compiler in DB2 pureXML and received a corporate award and an outstanding innovations award for her technical contributions to DB2 pureXML.</i> • <i>She consulted and worked together with IBM representatives to deliver SQL/XML 2005 standard. She was a key contributor to both XQuery and SQL/XML language standards.</i> • <i>She designed and delivered techniques to improve the performance of computing multiple XPath expressions into DB2 pureXML</i> • <i>She led the cross-site (research and SWG) team (~6) that designed and implemented a model to represent XQuery, rewrite optimizations and XPath matching algorithms that determine index eligibility for DB2 pureXML</i>
<p>September 1997 to August 2001</p>	<p><i>Research Assistant, University of Maryland, College Park</i> Project: <i>IMPACT (Interactive Maryland Platform for Agents Collaborating Together)</i></p> <p>Significant Accomplishments:</p> <ul style="list-style-type: none"> • <i>She designed and developed algorithms to detect common sub-expressions among heterogeneous agent programs and to exploit commonalities for improving query performance</i> • <i>She designed and developed algorithms to partition agent activities to enable sharing computations and improving performance</i> • <i>She designed, architected a query optimizer for heterogeneous agent programs</i>
<p>May 1998 to August 1998 and May 1999 to August 1999</p>	<p><i>Research Intern, Almaden Research Center, San Jose California</i> Project: <i>Garlic</i></p> <p>Significant Accomplishments:</p> <ul style="list-style-type: none"> • <i>She designed, and implemented cost models for Garlic query optimizer</i>
<p>September 1994 to August 1997</p>	<p><i>Research Assistant, Middle East Technical University, Ankara, Turkey</i></p>

<p>September 1994 to October 1996</p>	<p><i>Part time Researcher at TUBITAK (Scientific and Technical Research Council of Turkey) Software Research and Development Center, Ankara, Turkey</i></p> <p>Projects: <i>MIND (METU INteroperable Database)</i></p> <p>Significant Accomplishments:</p> <ul style="list-style-type: none"> • <i>She designed, architected, and implemented a dynamic query optimizer for MIND, which is a federated database system. She proposed an optimizer which decides on join orders dynamically at run-time as results to partial plans become available.</i> • <i>She implemented query optimization techniques for multi-databases</i> • <i>She designed and implemented a region-based query optimizer using Cascades query optimizer framework.</i>
--	--

HONORS and AWARDS

- ACM Distinguished Member, 2019
- IBM Technical Accomplishment for “Research Contributions to IBM Db2 EventStore”, 2019
- IBM Outstanding Technical Achievement Award for NLQ for ABFRL retail client, 2018
- IBM Technical Accomplishment for “NLQ for ABFRL retail client”, 2018
- IBM Outstanding Technical Achievement Award for Join Algorithms for Big Data, 2016
- IBM Technical Accomplishment for “Join Algorithms for Big Data”, 2015
- IBM Outstanding Technical Achievement Award for Big SQL 3.0, 2014
- IBM Extraordinary Technical Accomplishment for “Research Contributions to DB2”, 2010
- IBM Extraordinary Technical Accomplishment for “Standards Leadership: The Evolution from XML to Web Services to SOA”, 2009
- IBM Corporate Award for PureXML Database Technology, 2008
- IBM Outstanding Innovation Award for Technical Contributions to DB2XML, 2007
- IBM Research Division Award for Technical Contributions to the Standardization of XML Query Languages, 2005
- Graduate Fellowship from University of Maryland, College Park, 1997-1999
- Graduated FIRST in rank (out of 90 students) from the Department of Computer Engineering, Middle East Technical University, Ankara, 1994
- Graduated FIRST in rank from Malatya Science High School (one of the six science high schools in Turkey, accepting students who are in the first 600 out of 20000), 1990
- FIRST in the biology examinations of TUBITAK (Scientific and Technical Research Council of Turkiye), 1989

PATENTS

1. N. Seemann, **F. Özcan**, K. S. Beyer and R. J. Cochrane "LET to FOR to REG query transformation", Patent no US 7,165,063 B2, issued in January 2007
2. A. Balmin, **F. Özcan**, K. S. Beyer, R. J. Cochrane and H. Pirahesh, "XPath Containment for Index Selection and AST Matching", Patent no US 7,315,852, issued in January 2008
3. E. Ting and **F. Özcan**, "Early Out Multi-Pass Navigation for XPath/XQuery", disclosure number, September 2006
4. A. Balmin, **F. Özcan** and T. M. Tran, "Optimization of Extensible Markup Language Path Language (XPath) Expressions in a Database Management System Configured to Accept Extensible Markup Language (XML) Queries", US 7865502, issued January 2011
5. Q. Li, **F. Özcan**, L. Colby, H. Pirahesh, Z. Vagena and E. Shekita, "Join Algorithms over Full-text Indexes", ARC8-2008-0125, filed July 2009
6. M. Eltabakh, **F. Özcan**, Y. Sismanis, P. Haas, H. Pirahesh and J. Vondrak, "Split Elimination in Hadoop", filed in 2014
7. A. Balmin, **F. Özcan**, J. Traub, and R. Goncalves, "Adaptive Fragment Assignment for Processing File Data in a Relational Database", filed in 2014
8. K. Karanasos, A. Balmin, M. Kutsch, **F. Özcan**, V. Ercegovac, C. Xia, and J. Jackson, "Dynamically Optimizing Queries for Large Scale Data Platforms", filed in 2014
9. S. Doraiswamy, **F. Özcan**, H. Pirahesh, and M. Hörsken, "Parallel Data Streaming between Cloud-based Applications and Massively Parallel Systems", filed in 2015
10. T. Zou, Y. Tian, **F. Özcan**, and H. Pirahesh, "Joining Data across a Shared-Nothing Parallel Database and a Hadoop Distributed File System", filed in 2015
11. A. Floratou, N. Megiddo, **F. Özcan**, U. Kale, and N. Potti, "Caching Policies for Selection and Replacement of Objects", filed in 2015
12. U. F. Minhas, **F. Özcan**, N. Megiddo, U. Jain, and, B. van der Linden, "System and method for assigning splits of tables to processors", filed in 2016
13. A. Floratou, and **F. Özcan**, "Transforming Ontology Queries to SQL", filed in 2016
14. A. Floratou, N. Potti, **F. Özcan**, and H. Pirahesh, "Accessing Tables with Heterogeneous Partitions", filed in 2017
15. A. Quamar, **F. Özcan**, and K. Xirogiannopoulos, "Discovery and Creation of Rich Entities for Knowledge Bases", filed in 2017
16. C. Lei, R. Geis, **F. Özcan**, J. Kreulen, and D. Miller, "Query Relaxation using External Domain Knowledge for Query Answering", filed in 2019
17. A. Quamar, R. Alotaibi, and **F. Özcan**, "Ontology-based Query Routing for Domain Specific Knowledge Bases", filed in 2019
18. A. Quamar, C. Lei, R. Alotaibi, and **F. Özcan**, "Ontology-based Data Storage for Distributed Knowledge Bases", filed in 2019
19. C. Lei, A. Quamar, R. Alotaibi, and **F. Özcan**, "Property Schema Generation from Ontology", filed in 2020
20. A. Quamar, **F. Özcan**, and D. Miller, "Natural Language Processing and Machine Learning Based Ontology Artifact Generation", to be filed

21. A. Quamar, R.J. Moore, D. Miller, **F. Özcan**, and J. Kreulen, “Ontology-driven Conversational System for Business Intelligence”, under evaluation

PUBLICATIONS

Books and Book Chapters

- V.S. Subrahmanian, Piero Bonatti, Juergen Dix, Thomas Eiter, Sarit Kraus, **Fatma Özcan** and Robert Ross, "Heterogeneous Agent Systems", MIT Press, Cambridge, USA, June 2000, ISBN 0-262-19436-8
- Torsten Grust, H.V. Jagadish, **Fatma Özcan**, and Cong Yu, "XQuery Processors", Encyclopedia of Database Systems, Springer, Ling Liu and Tamer Özsu (editors-in-chief)
- **Fatma Özcan** and Ippokratis Pandis, “SQL Analytics over Big Data”, to appear in the second edition of Encyclopedia of Database Systems, Springer, Ling Liu and Tamer Özsu (editors-in-chief)

Journals and Magazines

1. C. Lei, **F. Özcan**, A. Quamar, A. R. Mittal, J. Sen, D. Saha, and K. Sankaranarayanan, “Ontology-Based Natural Language Query Interfaces for Data Exploration”, IEEE Data Engineering Bulletin, 41(3), September 2018
2. S. Bharadwaj et.al., “Creation and Interaction with Large-scale Domain-Specific Knowledge Bases”, PVLDB 10(12), 2017
3. D. Saha, A. Floratou, K. Sankaranarayanan, U. F. Minhas, A. Mittal, and **F. Özcan**, “ATHENA: An Ontology-Driven System for Natural Language Querying over Relational Data Stores”, PVLDB 9(12), 2016
4. D. Abadi, S. Babu, **F. Özcan**, and I. Pandis, “Tutorial: SQL-on-Hadoop Systems”, PVLDB 8(12), 2015
5. Y. Tian, **F. Özcan**, T. Zou, R. Goncalves, and H. Pirahesh, "Building a Hybrid Warehouse: Efficient Joins between Data Stored in HDFS and Enterprise Warehouse", ACM TODS, Volume 41, Issue 4, November 2016 [Invited as Best of EDBT 2015].
6. A. Balmin, K. Beyer, V. Ercegovac, J. McPherson, **F. Özcan**, H. Pirahesh, E. Shekita, Y. Sismanis, S. Tata, and Y. Tian, "A Platform for eXtreme Analytics", *IBM Journal of Research and Development*, 57(3/4), 2013.
7. **F. Özcan**, N. Seemann, and L. Wang, “XQuery Rewrite Optimization in IBM DB2 pureXML”, (*invited paper*), IEEE Data Engineering Bulletin, Special Issue on XQuery Processing: Practice and Experience, Vol. 31, No. 4, December 2008
8. **F. Özcan**, D. Chamberlin, K. Kulkarni and J-E. Michels, “Integration of SQL and XQuery in DB2”, IBM Systems Journal, Special Issue: Celebrating 10 Years of XML, Volume 45, No 2, June 2006, pages: 245--270
9. K. S. Beyer, R. J. Cochrane, M. Hvizods, V. Josifovski, J. Klewein, G. Lapis, G. Lohman, B. Lyle, M. Nicola, **F. Özcan**, H. Pirahesh, N. Seemann, A. Singh, T. Truong, B. Van der Linden, B. Vickery, C. Zhang and G. Zhang, “DB2 Goes Hybrid: Integrating Native XML and XQuery with Relational Data and SQL”, IBM Systems Journal, Special Issue: Celebrating 10 Years of XML, Volume 45, No 2, June 2006, pages: 271-298

10. **F. Özcan**, V.S. Subrahmanian, and J. Dix, "Improving Performance of Heterogeneous Agents", *Annals of Mathematics and Artificial Intelligence*, Volume 41, Issue:2-4, August 2004, pages: 339--395
11. K. Arisha, **F. Özcan**, R. Ross, V.S. Subrahmanian, T. Eiter and S. Kraus, "IMPACT: A Platform for Collaborating Agents", *IEEE Intelligent Systems*, Volume 14, No 2, March/April 1999
12. **F. Özcan**, S. Nural, P. Koksal, C. Evrendilek and A. Dogac, "Dynamic Query Optimization in Multidatabases", (*invited paper*), *IEEE Data Engineering Bulletin*, Special Issue on Query Optimization Vol. 20, No. 3, September 1997.
13. C. Evrendilek, A. Dogac, S. Nural and **F. Özcan**, "Multidatabase Query Optimization", *Distributed and Parallel Databases Journal*, Vol:5 No:1, January 1997, pages 77-114
14. **F. Özcan**, S. Nural, P. Koksal, M. Altinel and A. Dogac, "A Region based Query Optimizer through Cascades Query Optimizer Framework", (*invited paper*), *IEEE Data Engineering Bulletin*, Special Issue on Database Query Processing / Query Optimization, Vol 18 No:3, September 1995, pages 30-40
15. A. Dogac, C. Dengi, E. Kilic, G. Ozhan, **F. Özcan**, S. Nural, C. Evrendilek, U. Halici, B. Arpinar, P. Koksal, N. Kesim and S. Mancuhan, "METU Interoperable Database Systems", *ACM SIGMOD Record*, Vol 24 No:3, September 1995, pages 56-61

International Conferences and Workshops

1. A. Quamar, C. Lei, D. Miller, **F. Özcan**, J. Kreulen, R. J. Moore, and V. Efthymiou, "An Ontology-Based Conversation System for Knowledge Bases", *SIGMOD 2020 (industrial)*
2. C. Lei, V. Efthymiou, R. Geis, and **F. Özcan**, "Expanding Query Answers on Medical Knowledge Bases", *EDBT 2020*
3. J. Sen, **F. Özcan**, A. Quamar, G. Stager, A. R. Mittal, M. Jammi, C. Lei, D. Saha, and K. Sankaranarayan, "Natural Language Querying of Complex Business Intelligence Queries", *SIGMOD 2019*, pages 1997-2000
4. M. Karpathiotakis, A. Floratou, **F. Özcan**, and A. Ailamaki, "No Data Left Behind: Real-time Insights from a Complex Data Ecosystem", *SoCC 2017*, pages 108-120
5. V. Rosenfeld, R. Mueller, P. Tozun, and **F. Özcan**, "Processing Java UDFs in a C++ Environment", *SoCC 2017*, pages 419-431
6. **F. Özcan**, Y. Tian, and P. Tozun, "Hybrid Transactional/Analytical Processing: A Survey", tutorial, *ACM SIGMOD 2017*
7. C. Garcia-Arellano, R. Barber, M. Huras, R. Grosman, C. Mohan, R. Mueller, **F. Özcan**, H. Pirahesh, V. Raman, R. Sidle, A. Storm, Y. Tian, P. Tözün, D. Zilio, and G. Lohman, "Evolving Databases for New-Gen Big Data Applications", *CIDR 2017*
8. R. Barber, M. Hurras, G. Lohman, C. Mohan, R. Mueller, **F. Özcan**, H. Pirahesh, V. Raman, R. Sidle, O. Sidorkin, A. Storm, Y. Tian, and P. Tozun "Wildfire: Concurrent Blazing Data Ingest and Analytics", *SIGMOD demo*, 2016
9. J. Shi, Y. Qui, U. F. Minhas, L. Jian, C. Wang, B. Reinwald, and **F. Özcan**, "Clash of the Titans: MapReduce vs. Spark for Large Scale Data Analytics", *PVLDB 8(13)*, 2015
10. Y. Tian, T. Zou, **F. Özcan**, R. Goncalves, and H. Pirahesh, "Joins for Hybrid Warehouses: Exploiting Massive Parallelism in Hadoop and Enterprise Data Warehouses", *EDBT 2015*, pages 373-384

11. N. Katsipoulakis, Y. Tian, **F. Özcan**, H. Pirahesh, and B. Reinwald, “A Generic Solution to Integrate SQL and Analytics for Big Data”, EDBT 2015, vision paper, pages 671-676
12. A. Floratou, **F. Özcan**, and B. Schiefer, “Benchmarking SQL-on-Hadoop: TCP or not TCP?”, 5th Workshop on Big Data Benchmarking, 2014
13. K. Karanasos, A. Balmin, M. Kutsch, **F. Özcan**, V. Ercegovac, C. Xia, and J. Jackson, “Dynamically Optimizing Queries for Large Scale Data Platforms”, SIGMOD 2014 (research paper)
14. **F. Özcan**, N. Tatbul, D. J. Abadi, M. Kornacker, C. Mohan, K. Ramasamy, and J. L. Wiener, “Are we experiencing a big data bubble?”, SIGMOD Conference 2014, panel, pages 1407-1408
15. A. Floratou, U. Minhas, and **F. Özcan**, “SQL-on-Hadoop: Full Circle Back to Shared Nothing Database Architectures”, PVLDB 2014
16. O. Hassanzadeh, A. Kementsietsidis, B. Kimelfeld, R. Krishnamurthy, **F. Özcan**, and I. Pandis, "Next Generation Data Analytics at IBM Research", PVLDB, 2013
17. E. Kandogan, M. Roth, C. Kieliszewski, **F. Özcan**, B. Schloss, and M.-T. Schmidt, "Data For All: A Systems Approach to Accelerate the Path from Data to Insight", vision paper, IEEE Big Data Congress.
18. M. Eltabakh, **F. Özcan**, Y. Sismanis, P. Haas, H. Pirahesh and J. Vondrak, “Eagle-Eyed Elephant: Split-Oriented Indexing in Hadoop”, EDBT 2013, pages 89-100
19. V. Bicer, T. Tran, **F. Özcan**, O. Etzion, “CDMW 2012 - City Data Management Workshop: Workshop Summary”, CIKM 2012, pages 2756-2757
20. M. Eltabakh, Y. Tian, **F. Özcan**, R. Gemulla, A. Krettek, and J. McPherson, “CoHadoop: Flexible Data Placement and Its Exploitation in Hadoop”, PVLDB 4(9), 2011, pages 575-585
21. K. Beyer, V. Ercegovac, R. Gemulla, A. Balmin, M. Eltabakh, C-C. Kanne, **F. Özcan** and E. Shekita, “Jaql: A Scripting Language for Large Scale Semistructured Data Analysis”, VLDB 2011
22. **F. Özcan**, D. Hoa, K. Beyer, A. Balmin, C. J. Liu, and Y. Li, "Emerging Trends in Enterprise Data Analytics: Connecting Hadoop and DB2 Warehouse", SIGMOD 2011, pages 1161-1164
23. A. Balmin, L. Colby, E. Curtmola, Q. Li and **F. Özcan**, "Search Driven Analysis of Heterogenous XML Data", CIDR 2009
24. A. Balmin, L. Colby, E. Curtmola, Q. Li, **F. Özcan**, S. Srinivas and Z. Vagena, "SEDA: A System for Search, Exploration, Discovery, and Analysis of XML Data", VLDB 2008 demo, in PVLDB 1(2), pages 1408-1411
25. A. Balmin, **F. Özcan**, A. Singh and E. Ting, "Grouping and Optimization of XPath Expressions in DB2 pureXML", SIGMOD 2008 industrial paper
26. A. Balmin, **F. Özcan**, A. Singh and E. Ting, "Grouping and Optimization of XPath Expressions in System-RX", ICDE 2008, research poster
27. Z. Vagena, L. S. Colby, F. Özcan, A. Balmin and Q. Li, "On the Effectiveness of Flexible Querying Heuristics for XML Data", XSym 2007, pages 77-91
28. A. Balmin, K. S. Beyer, **F. Özcan** and M. Nicola, "On the Path to Efficient XML Queries", VLDB 2006, pages 1117--1128
29. M. Nicola and **F. Özcan**, "pureXML in DB2 9: Which way to query your XML data?", IBM developerWorks, June 2006

30. K. S. Beyer, R. J. Cochrane, V. Josifovski, J. Kleewein, G. Lapis, G. Lohman, B. Lyle, **F. Özcan**, H. Pirahesh, N. Seemann, T. Truong, B. Van der Linden, B. Vickery and C. Zhang, "System RX: One Part Relational, One Part XML", SIGMOD 2005, Baltimore, Maryland, June 2005
31. K. S. Beyer, D. Chamberlin, L. Colby, **F. Özcan**, H. Pirahesh and Y. Xu, "Extending XQuery for Analytics", SIGMOD 2005, Baltimore, Maryland, June 2005
32. K. S. Beyer, **F. Özcan**, S. Saiprasad and B. Van der Linden, "DB2/XML: Designing for Evolution", Tutorial, SIGMOD 2005, Baltimore, Maryland, June 2005
33. A. Balmin, **F. Özcan**, K.S. Beyer, R. J. Cochrane and H. Pirahesh, "A Framework for Using Materialized XPath Views in XML Query Processing", in Proc. of 30th Int. Conf. on Very Large Databases, VLDB 2004, Toronto, Canada, August 2004
34. K. S. Beyer, L. Colby, R. J. Cochrane, **F. Özcan** and H. Pirahesh, "XQuery for Analytics: Challenges and Requirements", in Proc. Of the 1st Int. Workshop on XQuery Implementation, Experience and Perspectives, XIME-P, in cooperation with ACM SIGMOD, June 2004, Paris, France
35. **F. Özcan**, V. S. Subrahmanian and L. Golubchik, "Optimal Agent Selection", in Proc. of Advances in AI, Joint German/Austrian Conf. on AI, KI-2001, Austria, September 2001, LNCS 2174, pages 2-17
36. **F. Özcan** and V. S. Subrahmanian, "Partitioning Activities for Agents", in Proc. of Int. Joint Conf. on AI, IJCAI-2001, Seattle, Washington, August 2001, pages 1218-1223
37. M. Tork Roth, **F. Özcan** and L. M. Haas, "Cost Models DO Matter: Providing Cost Information for Diverse Data Sources in a Federated System", VLDB, Edinburgh - Scotland - UK, 7th - 10th September 1999
38. K. Arisha, S. Kraus, **F. Özcan**, R. Ross and V.S. Subrahmanian, "IMPACT: The Interactive Maryland Platform for Agents Collaborating Together", ICMAS'98, Int. Conf on Multi Agent Systems, Paris, France, July 4-7, 1998, (Poster)
39. **F. Özcan**, S. Nural, P. Koksall, C. Evrendilek and A. Dogac, "Dynamic Query Optimization on a Distributed Object Management Platform", in Proc. of Intl. Conf. on Knowledge and Information Management, Maryland USA, November 1996, pages 117-124
40. A. Dogac, G. Ozhan, E. Kilic, **F. Özcan**, S. Nural, S. Mancuhan, C. Dengi, P. Koksall, U. Halici, B. Arpinar, C. Evrendilek and V. Sadjadi, "An Interoperability Structure For Developing Multidatabase Systems", ISCIS XI, November 6-8, 1996, Antalya. pages 103-112
41. S. Nural, P. Koksall, **F. Özcan** and A. Dogac, "Query Decomposition and Processing in Multidatabase Systems", in Proc. of the Intl. Conf. on Engineering Systems Design & Analysis, Montpellier, July 1996, pages 41-52
42. A. Dogac, U. Halici, E. Kilic, G. Ozhan, **F. Özcan**, S. Nural, C. Dengi, S. Mancuhan, B. Arpinar, P. Koksall and C. Evrendilek, "METU Interoperable Database System", in Proc. of ACM Sigmod Intl. Conf. on Management of Data, Montreal, June 1996, pages 552
43. G. Ozhan, A. Dogac, E. Kilic, **F. Özcan**, S. Nural, C. Dengi, U. Halici, B. Arpinar, P. Koksall, S. Mancuhan and C. Evrendilek, "Making ORACLE7, SYBASE, ADABAS D, interoperable through CORBA: MIND Project", In Proceedings of European Oracle User Group EOUG'96), Amsterdam, April 1996, pages 1047-1058
44. A. Dogac, C. Dengi, E. Kilic, G. Ozhan, **F. Özcan**, S. Nural, C. Evrendilek, U. Halici, B. Arpinar, P. Koksall, N. Kesimi and S. Mancuhan, "A Multidatabase System

- Implementation on CORBA", 6th Intl. Workshop on Research Issues in Data Engineering (RIDE-NDS' 96) New Orleans, February 1996, pages 2-11
45. Evrendilek, A. Dogac, S. Nural and **F. Özcan**, "Query Optimization in Multidatabase Systems", the second Intl. Workshop in Next Generation Inf. Tech. and Systems, Naharia, Israel, June 1995, pages 49-58
 46. A. Dogac, C. Dengi, E. Kilic, G. Ozhan, **F. Özcan**, S. Nural, C. Evrendilek, U. Halici, B. Arpinar, P. Koksal, N. Kesim and S. Mancuhan, "A Multidatabase System Implementation on CORBA", Conference on Object Oriented Programming Systems, Languages and Applications, OOPSLA '95, Austin Texas, October 1995, pages 148
 47. **F. Özcan** and A. Yazici, "Adapting K-d Trees to Fuzzy Indexing", in Proc. of International Conf. on Application of Fuzzy Systems, October 1994, pages 144-148

PROFESSIONAL ACTIVITIES

1. SIGMOD ACM Executive Committee: Secretary/Treasurer (elected for the term July 2017-June 2021)
2. VLDB Endowment Board of Trustees, 2016-2022
3. VLDB Journal, Editorial Board
4. Editor of the data analytics chapter in the second edition of Encyclopedia of Database Systems, Springer, Ling Liu and Tamer Özsu (editors-in-chief)
5. Co-editor of the Big SQL chapter in the Encyclopedia of Big Data Technologies, Springer

Conference/Workshop Organization:

1. Program Committee Chair, VLDB 2019, Editor in Chief, PVLDB Volume 12
2. Demo Track co-chair, SIGMOD 2019
3. Associate editor, PVLDB Volume 10, 2017
4. General Chair, SIGMOD 2016
5. Associate editor, PVLDB Volume 8, 2015
6. Industrial Program Co-chair, SIGMOD 2014
7. Program Co- chair and organizer, CDMW 2012 (co-located with CIKM 2012)
8. Industrial Program Co-chair, VLDB 2012
9. Demo Track Chair, EDBT 2010
10. Program Committee Vice Chair for the "XML" track of ICDE 2009
11. Track Chair of the XML data Processing, Filtering, Routing, and Algorithms, Program Committee of CIKM 2008
12. Program Co-Chair and organizer, XIME-P 2008 (co-located with SIGMOD 2008)
13. Industrial Program Co-chair, ICDE 2007
14. Proceedings Chair, SIGMOD 2005

Program Committee Member:

PVLDB Volume 14 Review Board, VLDB 2020 (Industrial), EDBT 2020 (Industrial), CIDR 2020, CIDR 2019, SIGMOD 2018 (demo and research), VLDB 2018 (industrial and research), HPTS 2017, VLDB 2017 (Industrial and Demo), ICDE 2017, CIDR 2017, VLDB 2016, SIGMOD 2015, EDBT 2014 (Industrial), ICDE 2014 (Industrial), CIKM 2013 (Posters), VLDB 2012 (Industrial), SIGMOD 2012, EDBT 2012, VLDB 2011, SIGMOD 2011, EDBT 2011, SIGMOD 2010, IDAR 2010 ((SIGMOD PhD Workshop on Innovative Database Research), KEYS Workshop 2009, ISCIS 2009, VLDB 2009, SIGMOD 2009, ICDE 2008 (Demo), WIDM 2007, VLDB Workshop on Management of Uncertain Data (MUD), 2007, ADBIS 2007, VLDB 2007, WWW 2007, CIKM 2006, XSym 2006, EDBT 2006, XIME-P 2005, VLDB 2004, ICDE 2004, IJCAI 2003, ISCIS 2003

Others service:

1. Invited Panelist, Information and Intelligent Systems, National Science Foundation, 2019
2. Invited Panelist, Information and Intelligent Systems, National Science Foundation, 2018
3. Invited Panelist, Information and Intelligent Systems, National Science Foundation, March 2015
4. Invited Panelist, Big Data Science and Engineering, National Science Foundation, July 2014
5. Invited Panelist, Information and Intelligent Systems, National Science Foundation, September 2005
6. Invited Panelist, Information and Data Management, National Science Foundation, March 2004
7. Reviewer for SIGMOD 2001, VLDB 2001, IEEE-TC , IJCAI 2001, ATAL 2001, ICLP 2001, SIGMOD 2003, VLDB 2003, SIGMOD 2004, SIGMOD 2005, VLDB 2005, WWW2006, ICDE 2006, and VLDB 2008
8. Reviewer for ACM Transactions on Internet Technology, ACM Transactions on the Web, ACM transactions on Databases (TODS), Multimedia Tools and Applications, IEEE Proceedings, VLDB Journal, Data and Knowledge Engineering Journal, IEEE TKDE