

Rajeev Rastogi

rastogi@research.bell-labs.com

EDUCATION

- | | |
|-----------|---|
| 1988-1993 | PhD in Computer Science at The University of Texas at Austin, May 1993. |
| 1988-1990 | MS in Computer Science at The University of Texas at Austin, May 1990. |
| 1984-1988 | B.Tech in Computer Science at The Indian Institute of Technology, Bombay, May 1988. |

THESIS

- | | |
|------------------------|--|
| Title | Concurrency Control in Multidatabase Systems. |
| Supervising Professors | Dr. Henry F. Korth and Dr. Abraham Silberschatz. |

EMPLOYMENT

- | | |
|---------------------|--|
| Oct 2004-present | Executive Director, Bell Labs Research India. |
| Dec 1999-Sep 2004 | Director, Internet Management Research Department, Bell Labs. |
| Dec 1999-Dec 2001 | Director, Distributed Systems Research Department, Bell Labs. |
| June 1998-Dec 1999 | Distinguished Member of Technical Staff, Bell Labs. |
| June 1993-June 1998 | Member of Technical Staff, Bell Labs. |
| 1993 Spring | Intern, Matsushita Information Technology Laboratory, Princeton. |
| 1992 Summer | Intern, Hewlett-Packard Laboratories, Palo Alto. |

ACADEMIC AWARDS AND HONORS

- | | |
|------|---|
| 1988 | MCD fellow, The University of Texas at Austin. |
| 1998 | Distinguished Member of Technical Staff, Bell Labs. |
| 2003 | Bell Labs Fellow, Bell Labs. |

RESEARCH OVERVIEW

Rajeev Rastogi has served as the *principal investigator* for research programs in the areas of high-performance storage systems, data mining, XML publishing, and network management (program descriptions provided below). In his role as principal investigator, he was responsible for (1) providing the vision and technical direction for the programs, (2) generating intellectual property (e.g., filing patents and publishing research papers), and (3) transferring technology to Lucent's product units. His most recent project is the Bell Labs Research Center in Bangalore, India, which he currently heads as Executive Director. He played a key role in launching the center, recruiting the research staff for it, and defining the center's research agenda which includes network monitoring of VoIP and IPTV services, peer-to-peer naming services, and low-cost networking technology to provide Internet connectivity to rural villages.

- **DATABLITZ MAIN-MEMORY DATABASE SYSTEM.** DataBlitz is a full-fledged database system with full support for ACID transactions, storage management, and indices. DataBlitz has a novel architecture in which application processes map the entire database into their virtual address space for speed. By eliminating interprocess communication and exposing lower level APIs, DataBlitz enables applications to perform lookups almost 80 times faster than most disk-based database engines available today. It is for this reason that DataBlitz forms the core platform for many of Lucent's wireline and wireless switching products requiring microsecond response times for database accesses.

Datablitz provides application developers with a rich set of features – these include data replication across geographically dispersed sites, data versioning protocols that enable non-blocking reads, code-words to detect data corruption, and hot standby support to reduce downtimes for mission-critical applications.

- **CINEBLITZ CONTINUOUS MEDIA STORAGE SERVER.** Time-sensitive applications like video require a guaranteed transfer rate of data, which conventional storage servers are not designed to provide. The CineBlitz storage server provides rate guarantees to video requests by keeping track of the available disk and buffer resources, and admitting a request only if there are sufficient resources to service the request. CineBlitz provides high overall throughput by reducing disk seek latency time, and eliminating rotational latency incurred when accessing data on disk. It also employs disk striping algorithms to uniformly distribute the workload across disks, and clever parity block placement strategies so that disk failures do not disrupt video transmission.
- **DATA MINING.** We have developed novel algorithms and systems to discover interesting patterns in call record data, Web server logs, etc. These include:

- *CURE and ROCK scalable clustering algorithms.* CURE can detect arbitrarily shaped clusters (prior clustering algorithms focused on detecting primarily spherically shaped clusters) and the ROCK system uses information on common neighbors to produce high quality clusters for categorical attributes.
- Algorithms for constructing decision trees with integrated building and pruning, and mining frequent sequences in the presence of regular expression constraints.
- The XTRACT system which uses information-theoretic techniques to infer a meaningful DTD for a database of XML documents.
- *Data stream analysis algorithms.* In a number of application domains (e.g., Telecom, network monitoring, sensor networks, etc.), massive amounts of data arrive continuously in the form of a stream. In such environments, memory and processing constraints dictate that query answering engines can examine each element in the stream only once. We have devised novel single-pass algorithms for (1) approximately answering general SQL queries on data streams using randomizing techniques that compute small “sketches” of each stream, (2) estimating the cardinality of general set expressions (using sketch-based techniques) over continuous update streams in small space and time per update, (3) detecting Distributed Denial-of-Service (DDoS) attacks in real time for large IP networks by continuously tracking the top-k destination IP addresses in terms of network traffic, and (4) continuously tracking set-expression and quantile queries in a distributed environment while minimizing the communication overhead between sites.

- **XML FILTERING AND PUBLISHING.** XML provides a self-describing, flexible mechanism for applications to exchange data. The X-Hub software supports a number of XML document processing applications like content-based routing of XML data, XML filtering in publish/subscribe systems, and XML data classification. X-Hub allows applications to specify parts of an XML document of interest in terms of patterns expressed using the XML Path Language XPath. It then processes each incoming XML document based on its matching patterns. In order to retrieve the set of XPath patterns that match an input XML document, X-Hub employs a specially designed index structure for XPath expressions. Furthermore, for scalability, X-Hub aggregates similar XPath patterns so that the total number of XPath expressions is reduced. In addition to X-Hub, we have also developed systems for publishing relational data in XML format conforming to a fixed DTD.

- **NETWORK MANAGEMENT.** Network management systems are responsible for ensuring that networks operate smoothly and stay available, and that service quality meets customer expectations.

- *NetInventory auto-discovery engine.* Knowledge of the “network map” is critical for a large number of network management tasks like fault and performance management, provisioning, traffic engineering, etc. The NetInventory system automatically discovers network elements (including unmanaged elements like hubs), and the physical and logical interconnections among them. NetInventory supports a wide range of Layer-2 technologies like Ethernet, ATM, and Frame Relay, and a multitude of access protocols to extract information from network elements like SNMP, TL1, CLI over modem/telnet, etc. NetInventory employs patented algorithms to infer Ethernet topology more accurately than existing commercial products, and as a consequence, it is at the core of a number of Lucent network management products and services.
- *IP VPN provisioning.* IP VPNs provide customers with predictable and secure network connections over a shared network. In the hose model for VPNs, each request specifies only aggregate ingress and egress bandwidth requirements for each VPN endpoint (instead of point-to-point bandwidth requirements). This has the advantages of flexibility, ease of specification, better bandwidth utilization, and scalability. Specifically for the hose model, we have developed algorithms for computing VPN paths for connecting VPN endpoints such that the total bandwidth reserved on the paths to guarantee QoS is minimized. We have also proposed solutions for computing near-optimal restoration topologies for VPNs, and optimal assignments of labels to VPN paths when MPLS is used as the signaling protocol to connect VPN endpoints.
- *OSPF and BGP protocol configuration.* IP networks use the OSPF and BGP protocols, respectively, for intra-domain and inter-domain routing. To simplify OSPF configuration, we have devised a dynamic programming scheme for computing the optimal IP aggregates to be advertised by each OSPF area so that the error in route computation (due to route aggregation) is minimized. We have also developed algorithms that balance network load by controlling ingress and egress network traffic through appropriate configurations of BGP policies at border routers.
- *Network monitoring.* In order to detect service quality problems in real time, we have developed a passive probe device that can snoop on packets traversing a network link. Our passive probes can be used to measure DNS and HTTP response times, VoIP and IPTV service quality, track subscribers internet usage, and detect DDoS attacks. Our passive probes employ high-performance packet capture cards, as well as advanced sampling and aggregation techniques, to scale to high-speed GigE and 10 GigE links.

Passive probes are useful for troubleshooting since they help to segment the network. We have devised algorithms for determining the optimal locations at which to place passive probes so that the number of probes required is minimized (to keep probe infrastructure costs as low as possible), and (1) performance anomalies like excessive link delay or loss can be quickly detected, and (2) the problem link with the anomaly can be accurately identified.

We have also developed low-overhead software agents that run on mobile handsets, and collect a wide range of performance statistics at the RF level (e.g., signal strength), and at the application level (e.g., server response times and frequently accessed Web pages). These agents can also be used by tech support staff to initiate tests remotely, and thus more accurately diagnose connectivity problems. A key challenge here is to ensure that our agents can run on a variety of embedded operating systems (e.g., Symbian, Windows CE), and have very little impact on the activities of mobile users.

PROFESSIONAL ACTIVITIES

Journal editorships

1. **Associate Editor**, *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, January 2001-January 2005.
2. **Guest Editor**, *IEEE Transactions on Knowledge and Data Engineering (TKDE) - Special Issue on Online Analysis and Querying of Continuous Data Streams*, 15/3, 2003.
3. **Associate Editor**, *Knowledge and Information Systems (KAIS) Journal*, October 2003-present.

Conference program committees

1. *International Conference on Data Engineering (ICDE)*, 1996.
2. *International Conference on Data Engineering (ICDE)*, 1997.
3. *International Conference on Cooperative Information Systems (CoopIS)*, 1997.
4. *ACM-SIGMOD 1998 International Conference on Management of Data, 1998*.
5. *IEEE Workshop on Research Issues in Data Engineering (ICDE)*, 1998.
6. *IEEE International Workshop on Multimedia DBMS*, 1998.
7. *IEEE International Workshop on Multimedia DBMS*, 1999.
8. *Data Mining and Knowledge Discovery (DMKD) Workshop*, 1999.
9. *International Conference on Very Large Databases (VLDB)*, 2000.
10. *International Conference on Extending Database Technology (EDBT)*, 2000.
11. *IEEE Knowledge and Data Exchange Workshop*, 2000.
12. *Conference on Information and Knowledge Management (CIKM)*, 1999.
13. **Program Co-Chair**, *Data Mining and Knowledge Discovery (DMKD) Workshop*, 2000.
14. *ACM-SIGMOD 1996 International Conference on Management of Data, 1998*, Demonstrations Committee Member.
15. *ACM-SIGKDD 2000 International Conference on Knowledge Discovery and Data Mining*, 2000.
16. *Multimedia Information Systems (MIS) Workshop*, 2000.
17. *SIAM International Conference on Data Mining*, 2001.
18. *International Conference on Data Mining (ICDM)*, 2001.
19. *ACM-SIGKDD 2001 International Conference on Knowledge Discovery and Data Mining*, 2001.
20. *International Conference on Data Warehousing and Knowledge Discovery*, 2001.
21. *Workshop on Network-Related Data Management (NRDM)*, 2001.
22. *WEBKDD Workshop*, 2001.
23. *International Conference on Data Engineering (ICDE)*, 2002.
24. *SIAM International Conference on Data Mining*, 2002.
25. **Panels Chair**, *ACM-SIGKDD 2001 International Conference on Knowledge Discovery and Data Mining*, 2002.
26. *International Conference on Data Warehousing and Knowledge Discovery*, 2002.
27. *ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS)*, 2003.
28. *Data Mining and Knowledge Discovery (DMKD) Workshop*, 2002.
29. *Fifth Workshop on Next Generation Information Technologies (NGITS)*, 2002.
30. **Conference Co-Chair**, *SIAM International Conference on Data Mining*, 2003.
31. *WEBKDD Workshop*, 2002.
32. *International Conference on Data Mining (ICDM)*, 2002.
33. **Tutorials Chair**, *Seventh Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2003.

34. *First Atlantic Web Mining Conference (AWIC)*, 2003.
35. **Vice-Chair**, *International Conference on Data Mining (ICDM)*, 2003.
36. *International Conference on Data Warehousing and Knowledge Discovery*, 2003.
37. *Sixth International High Performance Data Mining Workshop*, 2003.
38. *Data Mining and Knowledge Discovery (DMKD) Workshop*, 2003.
39. *Managing and Processing of Data Streams (MPDS) Workshop*, 2003.
40. *WEBKDD Workshop*, 2003.
41. *International Conference on Data Engineering (ICDE)*, 2004.
42. *International Conference on Very Large Databases (VLDB)*, 2004.
43. *International Conference on Distributed Computing Systems (ICDCS)*, 2004.
44. *ACM Symposium on Applied Computing (SAC)*, 2004.
45. **Program Co-Chair**, *International Conference on Data Mining (ICDM)*, 2004.
46. *International Conference on Data Warehousing and Knowledge Discovery*, 2004.
47. *International Computer Engineering Conference*, 2004.
48. *Second Atlantic Web Mining Conference (AWIC)*, 2004.
49. *Eighth Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2004.
50. *ACM-SIGKDD 2004 International Conference on Knowledge Discovery and Data Mining*, 2004.
51. *EDBT 2004 Workshop on Database Technologies for Handling XML Information on the Web*.
52. *International Conference on Database Technology (ICDT)*, 2005.
53. *International Conference on Data Engineering (ICDE)*, 2005.
54. *Tenth International Conference on Database Systems for Advanced Applications (DASFAA)*, 2005.
55. *Fifth SIAM International Conference on Data Mining*, 2005.
56. *Ninth Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2005.
57. **Conference Co-Chair**, *International Conference on Data Mining (ICDM)*, 2005.
58. *2005 ACM Internet Measurement Conference (IMC)*, 2005.
59. *International Conference on Mobile Ubiquitous Systems: Networking and Services (MobiQuitous)*, 2005.
60. *ACM SIGCOMM 2005 Workshop on Mining Network Data (MinNet-05)*.
61. *Sixth International Conference on Web-Age Information Management (WAIM)*, 2005.
62. *Twelfth International Conference on Management of Data (COMAD)*, 2005.
63. *Conference on Information and Knowledge Management (CIKM)*, 2005.
64. *International Conference on Data Engineering (ICDE)*, 2006.
65. *International Conference on Very Large Databases (VLDB)*, 2006.
66. *SIAM International Conference on Data Mining*, 2006.
67. *ACM-SIGMOD 2006 International Conference on Management of Data*, 2006.
68. *International Conference on Mobile Ubiquitous Systems: Networking and Services (MobiQuitous)*, 2006.

69. *Third International VLDB Workshop on Data Management for Sensor Networks (DMSN)*, 2006.
70. *ACM 15th Conference on Information and Knowledge Management (CIKM)*, 2006.
71. *ACM-SIGKDD 2004 International Conference on Knowledge Discovery and Data Mining*, 2006.
72. *Seventeenth European Conference on Machine Learning (ECML)/PKDD*, 2006.
73. *International Conference on Data Mining (ICDM)*, 2006.
74. *WEBKDD Workshop*, 2006.
75. *International Conference on Data Engineering (ICDE)*, 2007.
76. *ACM-SIGMOD 2007 International Conference on Management of Data*, 2007.
77. **Industrial Co-chair**, *Thirteenth International Conference on Management of Data (COMAD)*, 2006.
78. *SIAM International Conference on Data Mining*, 2007.
79. *International World Wide Web Conference (WWW)*, 2007.
80. *First Workshop on Wireless Systems: Advanced Research and Development (WISARD)*, 2007.
81. *International Conference on Very Large Databases (VLDB)*, 2007.

Ph.D. thesis committees

1. Phil Bohannon. Rutgers University, New Jersey, 1999. Thesis: Fault resilience in main-memory databases: Handling process halting failure and data corruption.
2. Jun Rao. Columbia University, New York, 2000. Thesis: Advanced query processing.
3. Edwin Knorr. University of British Columbia, Vancouver, 2002. Thesis: Outliers and data mining: Finding exceptions in data.
4. Jingren Zhao. Columbia University, New York, 2000. Thesis: Architecture-sensitive Database Query Processing.

TUTORIALS

1. R. Rastogi and K. Shim. Data mining on large databases. In *Conference on Information and Knowledge Management (CIKM)*, 1998.
2. R. Rastogi and K. Shim. Data mining on large databases. In *International Conference on Data Engineering*, 1999.
3. R. Rastogi and K. Shim. Data mining on large databases. In *ACM-SIGKDD 1999 International Conference on Knowledge Discovery and Data Mining, San Diego, 1999*.
4. R. Rastogi and K. Shim. Scaling data mining techniques to large databases. In *Fifth Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2001.
5. M. Garofalakis and R. Rastogi. Data mining meets the internet: techniques for web information retrieval and network data management. In *First SIAM International Conference on Data Mining (SDM'2001), Chicago, Illinois, April 2001*.
6. M. Garofalakis, J. Gehrke and R. Rastogi. Querying and mining data streams: you only get one look. In *ACM-SIGMOD 2002 International Conference on Management of Data, Madison, 2002*.
7. M. Garofalakis, J. Gehrke and R. Rastogi. Querying and mining data streams: you only get one look. In *ACM-SIGKDD 2002 International Conference on Knowledge Discovery and Data Mining*, 2002.

8. M. Garofalakis, J. Gehrke and R. Rastogi. Querying and mining data streams: you only get one look. In *Proceedings of the Very Large Database Conference (VLDB), Hong Kong, 2002*.
9. M. Garofalakis, J. Gehrke and R. Rastogi. Querying and mining data streams: you only get one look. In *International Conference on Web-Age Information Management, (WAIM), 2002*.
10. M. Garofalakis, J. Gehrke and R. Rastogi. Querying and mining data streams: you only get one look. In *International Conference on Management of Data, (COMAD), 2005*.

PANELS

1. S. Cha, P. Anandan, M. Hsu, C. Mohan, R. Rastogi, V. Sikka and H. Young. Globalization: challenges to the database community. In *Proceedings of the Very Large Database Conference (VLDB), Seoul, 2006*.

BOOK CHAPTERS

1. B. Ozden, R. Rastogi and A. Silberschatz. On the storage and retrieval of continuous media data. In *Multimedia Database Systems: Issues and Research Directions*, Editors - S. Jajodia and V. S. Subrahmanian, 1996.
2. B. Ozden, R. Rastogi and A. Silberschatz. Buffer replacement algorithms for multimedia databases. In *Multimedia Information Storage and Management*, Editor - S. M. Chung, Kluwer Academic Publishers, 1996.
3. B. Ozden, R. Rastogi, A. Silberschatz and P. S. Narayan. The Fellini multimedia storage server. In *Multimedia Information Storage and Management*, Editor - S. M. Chung, Kluwer Academic Publishers, 1996.
4. S. Guha, R. Rastogi and K. Shim. Techniques for clustering massive data sets. In *Clustering and Information Retrieval*, Editors - W. Wu, H. Xiong and S. Sekhar, Kluwer Academic Publishers, December 2003.

JOURNAL PUBLICATIONS

1. B. Ozden, R. Rastogi and A. Silberschatz. Research issues in multimedia storage managers. In *ACM Computing Surveys*, 27/4, December 1995, 617-620.
2. B. Ozden, A. Biliris, R. Rastogi and A. Silberschatz. A disk-based storage architecture for movie on demand servers. In *Information Systems Journal (Special Issue on Multimedia Information Systems)*, Editor - Stavros Christodoulakis, 20/6, 1995, 465-482.
3. B. Ozden, R. Rastogi and A. Silberschatz. On the design of a low-cost video-on-demand storage system. In *Multimedia Systems Journal*, Editor - P. V. Rangan, 4/1, February 1996, 40-54.
4. P. Bohannon, D. Lieuwen, R. Rastogi, S. Seshadri, A. Silberschatz and S. Sudarshan. The architecture of the Dali main-memory storage manager. In *Journal of Multi-media Tools and Applications*, 4/2, March 1997.
5. P. Bohannon, J. Parker, R. Rastogi, S. Seshadri, A. Silberschatz and S. Sudarshan. Distributed multi-level recovery in main-memory databases. In *Distributed and Parallel Database Systems Journal*, 6/1, January 1998.
6. S. Mehrotra, R. Rastogi, H. F. Korth, and A. Silberschatz. On the correctness of non-serializable executions. In *Journal of Computer and System Sciences*, 56/1, February 1998, 68-82.
7. S. Mehrotra, R. Rastogi, H. F. Korth, and A. Silberschatz. Ensuring Consistency in Multidatabases by Preserving Two-Level Serializability. In *ACM Transactions on Database Systems*, 23/2, June 1998.
8. R. Rastogi and K. Shim. PUBLIC: A decision tree classifier that integrates building and pruning. In *Data Mining and Knowledge Discovery Journal*, 4/4, October 2000, 315-344.

9. R. Rastogi, S. Seshadri, P. Bohannon, D. Leinbaugh, A. Silberschatz and S. Sudarshan. Versioning algorithms for improving transaction predictability in real-time main-memory databases. In *Real Time Systems Journal*, 19/3, November 2000.
10. S. Guha, R. Rastogi and K. Shim. ROCK: A robust clustering algorithm for categorical attributes. In *Information Systems Journal*, 25/5, 2000, 345-366.
11. S. Mehrotra, R. Rastogi, Y. Breitbart, H. F. Korth, and A. Silberschatz. Overcoming heterogeneity and autonomy in multidatabase systems. In *Information and Computation*, 167/2, 2001, 137-172.
12. S. Guha, R. Rastogi and K. Shim. CURE: An efficient algorithm for clustering large databases. In *Information Systems Journal*, 26/1, 2001, 35-58.
13. R. Rastogi and K. Shim. Mining optimized support rules for numeric attributes. In *Information Systems Journal*, 26/6, 2001, 425-444.
14. Minos Garofalakis, Kaushik Chakrabarti, Rajeev Rastogi, and Kyuseok Shim. Approximate query processing using wavelets. In *VLDB Journal (Special Issue on Best of VLDB 2000)*, 10/2-3, 2001, 199-223.
15. R. Rastogi and K. Shim. Mining optimized association rules with categorical and numeric attributes. In *IEEE Transactions on Knowledge and Data Engineering*, 14/1, 2002, 29-50.
16. M. Garofalakis, R. Rastogi and K. Shim. Mining Sequential patterns with regular expression constraints. In *IEEE Transactions on Knowledge and Data Engineering*, 14/3, 2002, 530-552.
17. A. Kumar, R. Rastogi, B. Yener and A. Silberschatz. Algorithms for provisioning virtual private networks in the hose model. In *IEEE/ACM Transactions on Networking*, 10/4, 2002, 565-578.
18. C. Y. Chan, P. Felber, M. Garofalakis and R. Rastogi. Efficient filtering of XML documents with XPath expressions. In *VLDB Journal (Special Issue on XML Data Management)*, 11/4, 2002, 354-379.
19. A. Gionis, M. Garofalakis, R. Rastogi, S. Seshadri and K. Shim. XTRACT: A system for extracting document type descriptors from XML documents. In *Data Mining and Knowledge Discovery Journal*, 7/1, 2003, 23-56.
20. P. Felber, C. Y. Chan, M. Garofalakis and R. Rastogi. Scalable XML routing for web services. In *IEEE Internet Computing*, 7/1, 2003, 49-57.
21. M. Garofalakis, D. Hoon, R. Rastogi and K. Shim. Efficient algorithms for constructing decision trees with constraints. In *Data Mining and Knowledge Discovery Journal*, 7/2, 2003, 187-214.
22. S. Brin, R. Rastogi and K. Shim. Mining optimized gain rules for numeric attributes. In *IEEE Transactions on Knowledge and Data Engineering*, 15/2, 2003, 324-338.
23. R. Rastogi, Y. Breitbart, M. Garofalakis and A. Kumar. Optimal configuration of OSPF aggregates. In *IEEE/ACM Transactions on Networking*, 11/2, 2003, 181-194.
24. C. Y. Chan, M. Garofalakis and R. Rastogi. RE-Tree: An Efficient Index Structure for Regular Expressions. In *VLDB Journal (Special Issue on Best of VLDB 2002)*, 12/2, 2003, 102-119.
25. P. Bohannon, R. Rastogi, S. Seshadri, A. Silberschatz and S. Sudarshan. Using codewords to protect database data from a class of software errors. In *IEEE Transactions on Knowledge and Data Engineering*, 15/5, 2003, 1120-1136.
26. A. Natsev, R. Rastogi and K. Shim. WALRUS: A similarity retrieval algorithm for image databases. In *IEEE Transactions on Knowledge and Data Engineering*, 16/3, 2004, 301-316.
27. Y. Breitbart, B. Jai, C. Martin, M. Garofalakis, R. Rastogi, S. Seshadri and A. Silberschatz. Topology discovery in heterogeneous IP networks: The NetInventory system. In *IEEE/ACM Transactions on Networking*, 12/3, 2004, 401-414.
28. A. Gupta, A. Kumar and R. Rastogi. Traveling with a pez dispenser (or, routing issues in MPLS). In *SIAM Journal of Computing*, 34/2, 2004, 453-474.

29. M. Garofalakis, S. Ganguly and R. Rastogi. Tracking set-expression cardinalities over continuous update streams. In *VLDB Journal (Special Issue on Data Streaming)*, 13/4, 2004, 354-369.
30. Y. Bejerano, Y. Breitbart, A. Orda, R. Rastogi and A. Sprinston. Algorithms for computing QoS paths with restoration. In *IEEE/ACM Transactions on Networking*, 13/3, 2005, 648-661.
31. Y. Bejerano and R. Rastogi. Robust monitoring of link delays and faults in IP Networks. In *IEEE/ACM Transactions on Networking*, 14/5, 2006, 1092-1103.

CONFERENCE PUBLICATIONS

1. S. Mehrotra, R. Rastogi, H. F. Korth, and A. Silberschatz. Non-serializable executions in heterogeneous distributed database systems. In *Proceedings of the First International Conference on Parallel and Distributed Information Systems, Miami Beach, Florida*, 1991.
2. S. Mehrotra, R. Rastogi, Y. Breitbart, H. F. Korth, and A. Silberschatz. The concurrency control problem in multidatabases: characteristics and solutions. In *Proceedings of ACM-SIGMOD 1992 International Conference on Management of Data, San Diego*, 1992.
3. S. Mehrotra, R. Rastogi, Y. Breitbart, H. F. Korth, and A. Silberschatz. Ensuring transaction atomicity in multidatabase systems. In *Proceedings of the Eleventh ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, San Diego*, 1992.
4. S. Mehrotra, R. Rastogi, H. F. Korth, and A. Silberschatz. A transaction model for multidatabase systems. In *Proceedings of the Twelfth International Conference on Distributed Computing Systems, Yokohama, Japan*, 1992.
5. S. Mehrotra, R. Rastogi, Y. Breitbart, H. F. Korth, and A. Silberschatz. Efficient global transaction management in multidatabase systems. In *Proceedings of the Third International Symposium on Database Systems for Advanced Applications, Daejeon, Korea*, 1993.
6. R. Rastogi, S. Mehrotra, Y. Breitbart, H. F. Korth, and A. Silberschatz. On Correctness of non-serializable executions. In *Proceedings of the Twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, Washington D.C.*, 1993.
7. R. Rastogi, H. F. Korth, and A. Silberschatz. Strict histories in object-based database systems. In *Proceedings of the Twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, Washington D.C.*, 1993.
8. R. Rastogi and M. A. Neimat. Enhancing pre-existing data managers with atomicity and durability. In *Proceedings of the First International Conference on the Applications of Databases, Vadstena, Sweden*, 1994.
9. B. Ozden, A. Biliris, R. Rastogi and A. Silberschatz. A low cost storage server for movie on demand databases. In *Proceedings of the Very Large Database Conference (VLDB), Chile*, 1994.
10. H. V. Jagadish, D. Lieuwen, R. Rastogi, A. Silberschatz and S. Sudarshan. Dali - a high performance main memory storage manager. In *Proceedings of the Very Large Database Conference (VLDB), Chile*, 1994.
11. B. Ozden, R. Rastogi and A. Silberschatz. On the storage and retrieval of continuous media. In *Conference on Information and Knowledge Management (CIKM)*, 1994.
12. B. Ozden, R. Rastogi and A. Silberschatz. A framework for the storage and retrieval of continuous media data. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems, Washington D.C.*, 1995.
13. B. Ozden, R. Rastogi, A. Silberschatz and C. Martin. Demand paging for video-on-demand servers. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems, Washington D.C.*, 1995.

14. R. Rastogi, H. F. Korth and A. Silberschatz. Exploiting transaction semantics in multidatabase environments. In *Proceedings of the Fifteenth International Conference on Distributed Computing Systems, Vancouver, 1995*.
15. T. Panagos, A. Biliris, H. V. Jagadish and R. Rastogi. Fine-granularity Locking and Private Logging for Distributed Architectures. In *Proceedings of International Conference on Extending Database Technology, 1996*.
16. T. Panagos, A. Biliris, H. V. Jagadish and R. Rastogi. Private Logging for High Performance Distributed Architecture. In *Proceedings of International Conference on Data Engineering, 1996*.
17. B. Ozden, R. Rastogi and A. Silberschatz. Disk striping in video server environments. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems, Hiroshima, 1996*.
18. B. Ozden, R. Rastogi and A. Silberschatz. Buffer replacement algorithms for multimedia databases. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems, Hiroshima, 1996*.
19. B. Ozden, R. Rastogi, P. Shenoy and A. Silberschatz. Fault-tolerant architectures for continuous media servers. In *Proceedings of ACM-SIGMOD 1996 International Conference on Management of Data, Montreal, 1996*.
20. P. Bohannon, J. Parker, R. Rastogi, S. Seshadri, A. Silberschatz and S. Sudarshan. Distributed multi-level Recovery in main-memory Databases. In *Proceedings of Parallel and Distributed Information Systems, Miami Beach, 1996*.
21. B. Ozden, R. Rastogi and A. Silberschatz. Periodic retrieval of videos from disk arrays. In *Proceedings of International Conference on Data Engineering, 1997*.
22. B. Ozden, R. Rastogi and A. Silberschatz. Multimedia support for databases. In *Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, Tucson, 1997*.
23. R. Rastogi, S. Seshadri, P. Bohannon, D. Leinbaugh, A. Silberschatz and S. Sudarshan. Logical and physical versioning in main-memory databases. In *Proceedings of the Very Large Database Conference (VLDB), Greece, 1997*.
24. R. Rastogi and K. Shim. Mining optimized association rules for categorical and numeric attributes. In *Proceedings of International Conference on Data Engineering, 1998*.
25. S. Guha, R. Rastogi and K. Shim. CURE: An efficient algorithm for clustering large databases. In *Proceedings of ACM-SIGMOD 1998 International Conference on Management of Data, Seattle, 1998*.
26. R. Rastogi and K. Shim. PUBLIC: A decision tree classifier that integrates building and pruning. In *Proceedings of the Very Large Database Conference (VLDB), New York, 1998*.
27. S. Guha, R. Rastogi and K. Shim. ROCK: A robust clustering algorithm for categorical attributes. In *Proceedings of International Conference on Data Engineering, 1999*.
28. R. Rastogi and K. Shim. Mining optimized support rules for numeric attributes. In *Proceedings of International Conference on Data Engineering, 1999*.
29. P. Bohannon, R. Rastogi, S. Seshadri, A. Silberschatz and S. Sudarshan. Detecting and recovering from data corruption. In *Proceedings of International Conference on Data Engineering, 1999*.
30. B. Hillyer, R. Rastogi and A. Silberschatz. Scheduling and data replication to improve tape jukebox performance. In *Proceedings of International Conference on Data Engineering, 1999*.
31. Y. Breitbart, K. Raghavan, R. Rastogi, S. Seshadri and A. Silberschatz. Update propagation protocols for replicated databases. In *Proceedings of ACM-SIGMOD 1999 International Conference on Management of Data, Philadelphia, 1999*.

32. S. Brin, R. Rastogi and K. Shim. Mining optimized gain rules for numeric attributes. In *Proceedings of ACM-SIGKDD 1999 International Conference on Knowledge Discovery and Data Mining, San Diego, 1999*.
33. A. Natsev, R. Rastogi and K. Shim. WALRUS: A similarity retrieval algorithm for image databases. In *Proceedings of ACM-SIGMOD 1999 International Conference on Management of Data, Philadelphia, 1999*.
34. J. Baulier, P. Bohannon, S. Gogate, C. Gupta, S. Joshi, A. Khivesera, H. Korth, P. McIlroy, J. Miller, P. P. S. Narayan, M. Nemeth, R. Rastogi, S. Seshadri, A. Silberschatz, S. Sudarshan and C. Wei. DataBlitz storage manager: Main-Memory database performance for critical applications. In *Proceedings of ACM-SIGMOD 1999 International Conference on Management of Data, Philadelphia, 1999* (Industrial Track).
35. M. Garofalakis, R. Rastogi and K. Shim. Sequential pattern mining with regular expression constraints. In *Proceedings of the Very Large Database Conference (VLDB), Edinburgh, 1999*.
36. S. Ramaswamy, R. Rastogi and K. Shim. Efficient algorithms for mining outliers from large data sets. In *Proceedings of ACM-SIGMOD 2000 International Conference on Management of Data, Dallas, 2000*.
37. A. Gionis, M. Garofalakis, R. Rastogi, S. Seshadri and K. Shim. XTRACT: A system for extracting document type descriptors from XML documents. In *Proceedings of ACM-SIGMOD 2000 International Conference on Management of Data, Dallas, 2000*.
38. K. Chakrabarti, M. Garofalakis, R. Rastogi and K. Shim. Approximate query answering using wavelets. In *Proceedings of the Very Large Database Conference (VLDB), Cairo, 2000*.
39. Y. Breitbart, C. Martin, M. Garofalakis, R. Rastogi, S. Seshadri and A. Silberschatz. Topology discovery in heterogeneous IP networks. In *Proceedings of IEEE INFOCOM, 2000*.
40. Mohana K. Lakhamraju, R. Rastogi, S. Seshadri and S. Sudarshan. On-line reorganization in object-oriented databases. In *Proceedings of ACM-SIGMOD 2000 International Conference on Management of Data, Dallas, 2000*.
41. P. Bohannon, P. McIlroy and R. Rastogi. Main-memory index structures with fixed size partial keys. In *Proceedings of ACM-SIGMOD 2001 International Conference on Management of Data, Santa Barbara, 2001*.
42. M. Garofalakis, D. Hoon, R. Rastogi and K. Shim. Efficient algorithms for constructing decision trees with constraints. In *Proceedings of ACM-SIGKDD 2000 International Conference on Knowledge Discovery and Data Mining, 2000* (Poster).
43. Y. Breitbart, C. Y. Chan, M. Garofalakis, R. Rastogi and A. Silberschatz. Efficiently monitoring bandwidth and latency in IP networks. In *Proceedings of IEEE INFOCOM, 2001*.
44. A. Deshpande, M. Garofalakis and R. Rastogi. Independence is good: Dependency-based histogram synopses for high-dimensional data. In *Proceedings of ACM-SIGMOD 2001 International Conference on Management of Data, Santa Barbara, 2001*.
45. S. Babu, M. Garofalakis and R. Rastogi. SPARTAN: A model-based semantic compression system for massive data tables. In *Proceedings of ACM-SIGMOD 2001 International Conference on Management of Data, Santa Barbara, 2001*.
46. A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi and B. Yener. Provisioning a virtual private network: A network design problem for multicommodity flow. In *Proceedings of the 33rd ACM Symposium on Theory of Computing (STOC), 2001*.
47. A. Kumar, R. Rastogi, A. Silberschatz and B. Yener. Algorithms for provisioning virtual private networks in the hose model. In *Proceedings of ACM-SIGCOMM, San Diego, 2001*.
48. R. Rastogi, Y. Breitbart, A. Kumar and M. Garofalakis. Optimal configuration of OSPF aggregates. In *Proceedings of IEEE INFOCOM, 2002*.

49. A. Gupta, A. Kumar and R. Rastogi. Routing Issues in MPLS (Or, How to Travel with a Pez Dispenser). In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS)*, Las Vegas, Nevada October 14-17, 2001.
50. C. Y. Chan, P. Felber, M. Garofalakis and R. Rastogi. Efficient filtering of XML documents with XPath expressions. In *Proceedings of International Conference on Data Engineering*, 2002.
51. G. Italiano, R. Rastogi and B. Yener. Restoration algorithms for virtual private networks in the hose model. In *Proceedings of IEEE INFOCOM*, 2002.
52. A. Dobra, M. Garofalakis, J. Gehrke and R. Rastogi. Processing Complex Aggregate Queries over Data Streams. In *Proceedings of ACM-SIGMOD 2002 International Conference on Management of Data, Madison, 2002*.
53. C. Y. Chan, M. Garofalakis and R. Rastogi. RE-Tree: An efficient index structure for regular expressions. In *Proceedings of the Very Large Database Conference (VLDB)*, Hong Kong, 2002.
54. C. Y. Chan, W. Fan, P. Felber, M. Garofalakis and R. Rastogi. Tree pattern aggregation for scalable XML data dissemination. In *Proceedings of the Very Large Database Conference (VLDB)*, Hong Kong, 2002.
55. M. Benedikt, C. Y. Chan, W. Fan, R. Rastogi, S. Zheng and A. Zhou. DTD-directed publishing with attribute translation grammars. In *Proceedings of the Very Large Database Conference (VLDB)*, Hong Kong, 2002.
56. A. Gupta, A. Kumar and R. Rastogi. Exploring the trade-off between label size and stack depth in MPLS routing. In *Proceedings of IEEE INFOCOM*, 2003.
57. T. Bressoud, R. Rastogi and M. Smith. Optimal configuration for BGP route selection. In *Proceedings of IEEE INFOCOM*, 2003.
58. Y. Bejerano and R. Rastogi. Robust Monitoring of Link Delays and Faults in IP Networks In *Proceedings of IEEE INFOCOM*, 2003.
59. Y. Bejerano, Y. Breitbart, A. Orda, R. Rastogi and A. Sprinston. Algorithms for computing QoS paths with restoration. In *Proceedings of IEEE INFOCOM*, 2003.
60. Y. Bejerano, Y. Breitbart, M. Garofalakis, R. Rastogi. Physical Topology Discovery for Large Multi-Subnet Networks. In *Proceedings of IEEE INFOCOM*, 2003.
61. M. Garofalakis, S. Ganguly and R. Rastogi. Processing set expressions over continuous update streams. In *Proceedings of ACM-SIGMOD 2003 International Conference on Management of Data, San Diego, 2003*.
62. M. Benedikt, C. Y. Chan, W. Fan, J. Freire and R. Rastogi. Capturing both types and constraints in data integration. In *Proceedings of ACM-SIGMOD 2003 International Conference on Management of Data, San Diego, 2003*.
63. A. Dobra, M. Garofalakis, J. Gehrke and R. Rastogi. Sketch-based multi-query processing over data streams. In *Proceedings of International Conference on Extending Database Technology*, 2004.
64. S. Ganguly, M. Garofalakis and R. Rastogi. Processing data-stream join aggregates using skimmed sketches. In *Proceedings of International Conference on Extending Database Technology*, 2004.
65. A. Das, S. Ganguly, M. Garofalakis and R. Rastogi. Distributed set expression cardinality estimation. In *Proceedings of the Very Large Database Conference (VLDB)*, Toronto, 2004.
66. M. Garofalakis, S. Ganguly, A. Kumar and R. Rastogi. Join-distinct aggregate estimation over update streams. In *Proceedings of the Twenty fourth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, Washington D.C.*, 2005.
67. P. Bohannon, M. Flaster, W. Fan and R. Rastogi. A cost-based model and effective heuristic for repairing constraints by value modification. In *Proceedings of ACM-SIGMOD 2005 International Conference on Management of Data, San Diego, 2005*.

68. G. Cormode, M. Garofalakis, S. Muthukrishnan and R. Rastogi. Holistic aggregates in a networked world: distributed tracking of approximate quantiles. In *Proceedings of ACM-SIGMOD 2005 International Conference on Management of Data, San Diego, 2005*.
69. Wenfei Fan, Jeffrey Xu Yu, Hongjun Lu, Jianhua Lu and Rajeev Rastogi. Query translation from XPath to SQL in the presence of recursive DTDs. In *Proceedings of the Very Large Database Conference (VLDB)*, Trondheim, 2005.
70. S. Agrawal, P. P. S. Narayan, J. Ramamirtham, R. Rastogi, M. Smith, K. Swanson and M. Thottan. VoIP service quality monitoring using active and passive probes. In *Proceedings of the First International Conference on Communication System Software and Middleware (COMSWARE)*, New Delhi, 2006.
71. S. Kashyap, S. Deb, K. V. M. Naidu, R. Rastogi and A. Srinivasan. Efficient gossip-based aggregate computation. In *Proceedings of the Twentyfifth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, Tucson, 2006*.
72. J. Golestani, M. Smith and R. Rastogi. Optimal scheduling for dynamic channel allocation in wireless LANs. In *Proceedings of the Second International Conference on Communication System Software and Middleware (COMSWARE)*, Bangalore, 2007.
73. Y. Breitbart, M. Garofalakis, A. Kumar, A. Gupta and R. Rastogi. On configuring BGP route reflectors. In *Proceedings of the Second International Conference on Communication System Software and Middleware (COMSWARE)*, Bangalore, 2007.
74. S. Parthasarthy, R. Rastogi and M. Thottan. Efficient design of end-to-end probes for source-routed networks. In *Proceedings of the Second International Conference on Communication System Software and Middleware (COMSWARE)*, Bangalore, 2007.
75. S. Agrawal, S. Deb, K. V. M. Naidu and R. Rastogi. Efficient detection of distributed constraint violations. In *Proceedings of International Conference on Data Engineering, 2007 (Poster)*.
76. S. Agrawal, S. Deb, K. V. M. Naidu and R. Rastogi. Efficient detection of distributed constraint violations. In *International Conference on Management of Data, (COMAD)*, 2006.
77. S. Agrawal, K. V. M. Naidu and R. Rastogi. Diagnosing link-level anomalies using passive probes. In *Proceedings of IEEE INFOCOM, 2007*.
78. P. Dutta, S. Jaiswal and R. Rastogi. Routing and channel allocation in rural wireless mesh networks. In *Proceedings of IEEE INFOCOM, 2007*.

WORKSHOP AND OTHER PUBLICATIONS

1. S. Mehrotra, R. Rastogi, H. F. Korth, and A. Silberschatz. Relaxing serializability in multidatabase systems. In *Proceedings of the Second International Workshop on Research Issues on Data Engineering: Transaction and Query Processing, Mission Palms, Arizona, 1992*.
2. R. Rastogi, S. Mehrotra, Y. Breitbart, H. F. Korth, and A. Silberschatz. Transcending the serializability correctness in multidatabase systems. In *Data Engineering Bulletin*, June 1993.
3. B. Ozden, R. Rastogi and A. Silberschatz. Disk striping in video server environments. In *Proceedings of the Third International Workshop on Community Networking, Princeton, 1995*.
4. B. Ozden, R. Rastogi and A. Silberschatz. Disk Striping in Video Server Environments. In *Data Engineering Bulletin*, 18/4, 1995.
5. B. Ozden, R. Rastogi and A. Silberschatz. Architecture Issues In Multimedia Storage Systems. SIGMETRICS Performance Evaluation Review, 25/2, 1997.
6. J. Baulier, P. Bohannon, S. Gogate, C. Gupta, S. Joshi, A. Khivesera, H. Korth, P. Mcilroy, J. Miller, P. P. S. Narayan, M. Nemeth, R. Rastogi, A. Silberschatz and S. Sudarshan. DataBlitz: A high performance main-memory storage manager. In *Proceedings of the Very Large Database Conference (VLDB)*, New York, 1998 (Exhibits).

7. M. Garofalakis, S. Ramaswamy, R. Rastogi and K. Shim. Of crawlers, portals, mice, and men: is there more to mining the Web? In *Proceedings of ACM-SIGMOD 1999 International Conference on Management of Data, Philadelphia, 1999* (Panel).
8. M. Garofalakis, R. Rastogi, S. Seshadri and K. Shim. Data mining and the Web: Past, present and future. In *Proceedings of the Second Workshop on Web Information and Data Management (WIDM), Kansas City, Missouri, 1999*.
9. R. Rastogi. DataBlitz – providing real-time database support for telecommunication networks. In *37th Meeting of IFIP Working Group 10.4, 2000*.
10. D. Gnupolos and R. Rastogi. Workshop report: 2000 ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery. In *SIGKDD Explorations, 2/1, 2000*.
11. M. Garofalakis and R. Rastogi. Scalable data mining with model constraints. In *SIGKDD Explorations (Special Issue on Scalable Data Mining)*, December 2000.
12. P. Felber, B. Jai, R. Rastogi and M. Smith. Using semantic knowledge of distributed objects to increase reliability and availability. In *Sixth International Workshop on Object-oriented Real-time Dependable Systems (WORDS), 2001*.
13. S. Babu, M. Garofalakis and R. Rastogi. Model-based semantic compression for network data tables. In *Workshop on Network Related Data Management (NRDM), 2001*.
14. M. Garofalakis and R. Rastogi. Data mining meets network management: The NEMESIS project. In *Workshop on Research Issues in Data Mining and Knowledge Discovery (DMKD), 2001*.
15. Kenneth A. Ross, Jiawei Han, Nick Koudas, Rajeev Rastogi, Daniel J. Rosenkrantz, Peter Scheuerman, Dan Suciu. Reminiscences on Influential Papers. In *SIGMOD Record 30/3, 2001, 101-104*.
16. P. Felber, C. Y. Chan and R. Rastogi. Large-scale dissemination of XML data. In *Byte Magazine, March, 2002*.
17. M. Garofalakis and R. Rastogi. Network data mining and analysis: the NEMESIS project. In *Sixth Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD), 2002 (Industrial Track)*.
18. R. Rastogi. Single-pass algorithms for querying and mining data streams. In *Fifth International Workshop on High Performance Data Mining: Resource and Location aware Mining (HPDM), April, 2002*.
19. S. Babu, M. Garofalakis and R. Rastogi. SPARTAN: Using constrained models for guaranteed-error semantic compression. In *SIGKDD Explorations (Special Issue on Constraints in Data Mining)*, June 2002.
20. M. Garofalakis, A. Gionis, R. Rastogi, S. Seshadri and K. Shim. DTD inference from XML documents: The XTRACT approach. In *Data Engineering Bulletin (Special Issue on Structure Discovery)*, September, 2003.
21. P. Dutta, S. Jaiswal, K. V. M. Naidu, D. Panigrahi, R. Rastogi and A. Todimala. VillageNet: A low-cost, 802.11-based mesh network for rural regions. In *First Workshop on Wireless Systems: Advanced Research and Development (WISARD), 2007*.

INVITED TALKS

1. The Fellini multimedia storage server. In *Matsushita Information Technology Laboratory, Princeton, 1994*.
2. A low cost storage server for movie on demand databases. In *Ohio State University, 1995*.
3. Fault-tolerant architectures for continuous media servers. In *State University of New York (SUNY), Buffalo, 1996*.

4. WALRUS: A similarity retrieval algorithm for image databases. In *IBM Almaden Research Center*, August, 1999.
5. Data mining and the web: past, present and future. In *Second Workshop on Web Information and Data Management (WIDM), Kansas City, Missouri*, 1999.
6. DataBlitz – providing real-time database support for telecommunication networks. In *37th Meeting of IFIP Working Group 10.4*, February, 2000.
7. Data mining applications in the World Wide Web. In *Second Canadian Database Workshop, Toronto*, 2000.
8. Data mining meets the internet: techniques for web information retrieval. In *Korean Advanced Institute for Science and Technology (KAIST), Korea*, March, 2001.
9. DataBlitz – A high-performance storage manager. In *Seoul National University, Korea*, March, 2001.
10. Approximate query answering using wavelets. In *Lawrence Livermore National Laboratory (LLNL)*, 2001.
11. Data mining meets the internet: techniques for web information retrieval. In *DIMACS Summer School Tutorial on New Frontiers in Data Mining*, August, 2001.
12. Spartan: A model-based semantic compression system for massive data tables. In *DIMACS Working Group Meeting on Data Compression in Networks and Applications*, March, 2002.
13. **Keynote Address.** Single-pass algorithms for querying and mining data streams. In *Fifth International Workshop on High Performance Data Mining: Resource and Location aware Mining (HPDM)*, April, 2002.
14. **Keynote Address.** Querying and mining data streams: You only get one look. In *Sixth Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, May, 2002.
15. Data stream processing and protocol configuration in IP networks. University of Illinois, Urbana Champagne, September 2002.
16. Data stream processing and protocol configuration in IP networks. University of Chicago, Chicago, November 2002.
17. Data stream processing and protocol configuration in IP networks. Ohio State University, Columbus, November 2002.

PATENTS

Granted

1. Method Of Retrieving Continuous And Non-Continuous Media Data From A File System. U.S. Patent No. 5708796. 01/13/1998.
2. Multimedia On-Demand Server. U.S. Patent No. 5720037. 02/17/1998.
3. Method And Apparatus For Selective Buffering Of Pages To Provide Continuous Media Data To Multiple Users. U.S. Patent No. 5721956. 02/24/1998.
4. Video Data Retrieval Method For Use In Video Server Environments That Use Striped Disks. U.S. Patent No. 5724543. 03/03/1998.
5. Multimedia On-Demand Server Having Different Transfer Rates. U.S. Patent No. 5754773. 05/19/1998.
6. Method and Apparatus of Retrieving Continuous and Non-Continuous Media Data From A File System. U.S. Patent No. 5761692. 06/02/1998.

7. System For Video Server Using Coarse-Grained Disk Striping Method In Which Incoming Requests Are scheduled And Rescheduled Based On Availability Of Bandwidth. U.S. Patent No. 5826110. 10/20/1998.
8. System And Method For Restoring A Distributed Checkpointed Database. U.S. Patent No. 5845292. 12/01/1998.
9. System And Method For Restoring A Multiple Checkpointed Database In View Of Loss Of Volatile Memory. U.S. Patent No. 5864849. 01/26/1999.
10. Lookahead Buffer Replacement Method Using Ratio Of Clients Access Order Offsets And Buffer Data Block Offsets. U.S. Patent No. 5870551. 02/09/1999.
11. Technique For Effectively Instantiating Attributes In Association Rules. U.S. Patent No. 5946683. 08/31/1999.
12. Method And Apparatus For Providing Enhanced Pay Per View In Video Server Employing Coarse-Grained Striping Scheme. U.S. Patent No. 5964829. 10/12/1999.
13. Local Logging in a Distributed Database Management Computer System. U.S. Patent No. 5966706. 10/12/1999.
14. Method And Apparatus For Providing Enhanced Pay Per View In A Video Server. U.S. Patent No. 6012080. 01/04/2000.
15. A Method, Apparatus And Programmed Medium For Clustering Databases With Categorical Attributes. U.S. Patent No. 6049797. 04/11/2000.
16. Fault-Tolerant Architectures For Continuous Media Servers. U.S. Patent No. 6079028. 06/20/2000.
17. A Method, Apparatus And Programmed Medium For Clustering Large Databases. U.S. Patent No. 6092072. 07/18/2000.
18. System And Method For Physically Versioning Data In A Main Memory Database. U.S. Patent No. 6122645. 09/19/2000.
19. System And Method For Aging Versions Of Data In A Main Memory Database. U.S. Patent No. 6125371. 09/26/2000.
20. Cache Management System For Continuous Media Servers. U.S. Patent No. 6154813. 11/28/2000.
21. Method For Mining Association Rules In Data. U.S. Patent No. 6185549. 02/06/2001.
22. System And Method For Providing Hot Spare Redundancy And Recovery For A Very Large Database Management System. U.S. Patent No. 6205449. 03/20/2001.
23. Decision Tree Classifier With Integrated Building And Pruning Phases. U.S. Patent No. 6247016. 06/12/2001.
24. A System To Improve The Performance Of Tape Storage Systems. U.S. Patent No. 6282607. 08/28/2001.
25. On-Line Reorganization In Object-Oriented Databases. U.S. Patent No. 6343296. 01/29/2002.
26. Method And Apparatus For Detecting And Recovering From Data Corruption Of Database Via Read Prechecking And Deferred Maintenance Of Codewords. U.S. Patent No. 6374264. 04/16/2002.
27. System And Method for Serializing Lazy Updates In A Distributed Database Without Requiring Timestamps. U.S. Patent No. 6381609. 04/30/2002.
28. Method And Apparatus For Detecting And Recovering From Data Corruption Of Database Via Read Logging. U.S. Patent No. 6449623. 09/10/2002.
29. System And Method For Constraint Based Sequential Pattern Mining. U.S. Patent No. 6473757. 10/29/2002.

30. System And Method For Serializing Updates To Ancestor Sites In A Distributed Database. U.S. Patent No. 6499037. 12/24/2002.
31. Timestamp-Based System And Method For Serializing Lazy Updates In A Distributed Database. U.S. Patent No. 6502095. 12/31/2002.
32. Method For Identifying Outliers In Large Data Sets. U.S. Patent No. 6643629. 11/04/2003.
33. Determination Of Physical Topology Of A Communication Network. U.S. Patent No. 6697338. 02/24/2004.
34. Methods Of Imaging Based On Wavelet Retrieval Of Scenes. U.S. Patent No. 6751363 06/15/2004.
35. Approximate Query Processing Using Wavelets. U.S. Patent No. 6760724. 07/06/2004.

Pending

1. Method And Apparatus For Mining Optimized Gain Rules In Data. Brin 1-31-7. 04/13/1999.
2. Document Descriptor Extraction Method. Garofalakis 6-1-36-11-10. 06/16/2000.
3. System And Method For Optimizing Open Shortest Path First Aggregates And Autonomous Network Domain Incorporating The Same. Breitbart 14-8-1-39. Filed 02/01/2001.
4. Method For Provisioning Bandwidth In A Computer Network. Kumar 2-41-59-9. Filed 11/16/2001.
5. SPARTAN: A Model-Based Semantic Compression System For Massive Data Tables. Babu 1-10-42. Filed 12/28/2001.
6. System and Method for Optimally Configuring Border Gateway Selection for Transit Traffic Flows in a Computer Network. Bressoud 2-44-1. Filed 07/01/2002.
7. System and Method for Filtering XML Documents With Xpath Expressions. Chan 2-1-12-45. Filed 07/09/2002.
8. Method And Apparatus For Routing A Packet Within A Plurality Of Nodes Arranged In A Line Or A Tree Given A Maximum Stack Depth. Gupta 1-3-46. Filed 03/31/2003.
9. Restoration For Virtual Private Networks. Italiano 1-50-10. Filed 06/13/2003.
10. Techniques For Information Dissemination Using Tree Pattern Subscriptions And Aggregation Thereof. Chan 3-1-2-14-51. Filed 06/20/2003.
11. Automated Transformation Of Specifications For Devices Into Executable Module. Jai 5-4-52. Filed 06/20/2003.
12. System And Method For Provisioning QOS Paths With Restoration In A Network. Bejerano 3-16-1-49-1. Filed 09/26/2003.
13. System And Method For Determining The Physical Topology Of A Network Having Multiple Subnets. Bejerano 1-15-13-47. Filed 05/27/2003.
14. System And Method For Monitoring Link Delays And Faults In An IP Network. Bejerano 2-48. Filed 09/26/2003.
15. The Allocation Of Channels To Wireless LANS. Golestani 5-53-3. Filed 04/12/2004.
16. Method For Distinct Count Estimation Over Joins Of Continuous Update Stream. Ganguly 3-17-7-57. Filed 09/30/2004.
17. Methods And Devices For Approximating Optimal Channel Allocations. Golestani 6-59-4. Filed 09/30/2004.
18. Computing Optimal Channel Allocations Using Decomposition Methods And Related Devices. Golestani 7-60-5. Filed 09/30/2004.

19. Method For Performing Information-Preserving DTD Schema Embeddings. Benedikt 11-16-56. Filed 09/30/2004.
20. Streaming Algorithms For Robust, Real-Time Detection Of DDoS Attacks. Ganguly 2-15-55-58. Filed 09/30/2004.
21. Processing Data-Stream Join Aggregates Using Skimmed Sketches. Ganguly 6-21-63. Filed 12/29/2004.
22. Tracking Set-Expression Cardinalities Over Continuous Update Streams. Ganguly 5-19-61. Filed 12/29/2004.
23. Sketch-Based Multi-Query Processing Over Data Streams. Dobra 1-20-1-62. Filed 12/29/2004.
24. Equivalence Class-Based Method And Apparatus For Cost-Based Repair Of Database Constraint Violations. Bohannon 17-5-3-64. Filed 12/29/2004.
25. Distributed Set-Expression Cardinality Estimation. Das 1-4-18-58. Filed 12/30/2004.
26. System And Method For XML Data Integration. Benedikt 10-3-54. Filed 03/31/2005.
27. Method For Translating Recursive XPath Queries Into SQL Queries In The Presence Of Recursive XML Schemas.
28. A Method For Distributed Tracking Of Approximate Quantiles. Cormode 1-23-19-65.