

Annual Report for Period:09/2009 - 08/2010

Submitted on: 09/24/2010

Principal Investigator: Yang, Jun .

Award ID: 0713498

Organization: Duke University

Submitted By:

Yang, Jun - Principal Investigator

Title:

III-COR: Scalable Publish/Subscribe: Unifying Data Processing and Dissemination

Project Participants

Senior Personnel

Name: Yang, Jun

Worked for more than 160 Hours: Yes

Contribution to Project:

Jun Yang has been serving as the faculty lead on this project with Pankaj K. Agarwal.

Name: Agarwal, Pankaj

Worked for more than 160 Hours: Yes

Contribution to Project:

Pankaj K. Agarwal has been serving as the faculty lead on this project with Jun Yang.

Post-doc

Graduate Student

Name: Chandramouli, Badrish

Worked for more than 160 Hours: Yes

Contribution to Project:

Badrish Chandramouli was with the project when it was first conceived, and had been the main student contributor until he graduated in 2008. He had been involved in most of the research and development efforts, including developing the ProSem system (server and networking components) and demonstrating it at SIGMOD 2008, the original vision paper proposing the new database/network interfaces for publish/subscribe (first author on SIGMOD 2006 paper), scalable processing and dissemination of value-based notification conditions (first author on VLDB 2007 paper), and scalable processing and dissemination of select-join subscriptions (first author on VLDB 2008 paper). He defended his PhD dissertation on this topic in July 2008. He is now a researcher at Microsoft Research.

Name: Yu, Albert

Worked for more than 160 Hours: Yes

Contribution to Project:

Albert Yu is a PhD student who started working on the project in Fall 2007. He is now the main student contributor to the project. He is primarily working on the problem of dissemination network design for wide-area publish/subscribe, considering not only similarities in subscription interests but also their network locations and event distribution. He built the system visualization component for the ProSem system, which was demonstrated at SIGMOD 2008. He also developed a generator of wide-area content-based publish/subscribe workloads by extrapolating the limited amount of publicly available data; this work is currently under submission.

Name: Zheng, Ying

Worked for more than 160 Hours: Yes

Contribution to Project:

Ying Zheng, a PhD student, was with the project in Spring 2008. She built the Java-based subscription client for the ProSem system, which was demonstrated at SIGMOD 2008.

Name: Sharathkumar, R

Worked for more than 160 Hours: Yes

Contribution to Project:

R. Sharathkumar is a PhD student working on maintaining summaries of multidimensional data acquired through sensors. In the context of this project, he is developing indexing structures for reporting the summaries of streaming data.

Name: Thonangi, Risi

Worked for more than 160 Hours: Yes

Contribution to Project:

Risi Thonangi, a PhD student, worked on the problem of efficiently computing weighted proximity best-joins over multiple lists, with applications in information retrieval and extraction. This work was published in ICDE 2009, and is part of the effort to expand the project beyond traditional relational data and queries. Risi has been funded by the department thus far, although his travel to ICDE 2009 was supported by this grant.

Name: Gujgunte, Shashidhar

Worked for more than 160 Hours: Yes

Contribution to Project:

Shashidhar Gujgunte is working on aggregating data distributed over a sensor network. Currently he is developing algorithms to compute the topology of 'broker' nodes over a sensor network for facilitating aggregation of data to a desired node and for scheduling data transmission.

Name: Phillip, Jeff

Worked for more than 160 Hours: No

Contribution to Project:

Undergraduate Student

Name: Zhang, Weiping

Worked for more than 160 Hours: No

Contribution to Project:

In Summer 2009, Weiping Zhang, a rising junior, is working with the PI on expanding the project scope to consider richer query languages with statistical analysis. In particular, he plans to develop an efficient storage manager for array data types over the summer.

Name: Lee, Kyung Min

Worked for more than 160 Hours: No

Contribution to Project:

Name: Butani, Aneesh

Worked for more than 160 Hours: No

Contribution to Project:**Name:** Beutel, Alex**Worked for more than 160 Hours:** No**Contribution to Project:****Name:** Prapdiyu, Aditya**Worked for more than 160 Hours:** No**Contribution to Project:****Technician, Programmer****Other Participant****Research Experience for Undergraduates****Organizational Partners****Other Collaborators or Contacts**

Mark De Berg (TU, Eindhoven), AnHai Doan (Wisconsin), Alon Efrat (Arizona), Bryon J. Gao (Texas State San Marcos), Leo Guibas (Stanford), Sariel Har-Peled (UIUC), Hao He (Google, Inc.), Raghu Ramakrishnan (Yahoo! Research), Amin Vahdat (UC San Diego), Haixun Wang (IBM Watson Research, now Microsoft Research Asia), Junyi Xie (Oracle Corp.), Ke Yi (Hong Kong UST), Hai Yu (Google, Inc.)

Activities and Findings**Research and Education Activities:**

The Digital Age has brought about unprecedented growth in the amount of data being generated, the number of data consumers, and the diversity of their interests and locations. Traditionally, users poll sources for information, but for many applications, polling is hardly scalable and may miss important events. The alternative offered by publish/subscribe systems is to push notifications to users with matching interests. This project aims at building a next-generation publish/subscribe system to face the new challenges. We propose an end-to-end solution consisting of techniques from subscription processing and indexing to dissemination network design, which work together to support efficient and powerful subscription functionalities, allowing users to control precisely what they want and when they want it.

One main feature distinguishing the proposed approach from previous work is joint consideration of subscription processing and notification dissemination. We propose a promising approach that allows complex, stateful subscriptions to be handled by simple, stateless dissemination mechanisms, with a clean system design that is easy to implement and scale. A cost-based optimizer, inspired by database query optimization, chooses the best processing and

dissemination strategies jointly and dynamically.

Besides system building, this project tackles many new algorithmic challenges, including, e.g., scalably processing a large number of complex subscriptions; exploiting event and subscription characteristics to combat worst-case complexity; balancing semantic similarity and network proximity in dissemination network design; and efficiently maintaining statistics for high-dimensional events and subscriptions.

In Year 3 of the project, we investigated the following specific research problems.

1. Publish/Subscribe Workload Generator. Making content-based publish/subscribe systems scalable in a wide-area network requires considering multiple factors, e.g., distribution of events, similarity of subscriber interests in the event space, and proximity of subscriber locations in the network. A major obstacle for this research is the lack of publicly available, realistic workloads, because of concerns of privacy and commercial interests in releasing user information. We have developed a workload generator for wide-area content-based publish/subscribe systems, which extrapolates the limited amount of public data extracted from Google Groups and PlanetLab in order to provide a reasonable guess for a large, realistic workload. Our hope is that this tool will help publish/subscribe researchers evaluate their research. This work has been published in NetDB 2009 and this tool has been made available to the research community.

2. Dissemination Network Design for Wide-Area Publish/Subscribe. We consider the problem of assigning subscribers to brokers in a wide-area content-based publish/subscribe system. A good assignment should consider both subscriber interests in the event space and subscriber locations in the network space, and balance multiple performance criteria including bandwidth, delay, and load balance. The resulting optimization problem is NP-complete, so systems have turned to heuristics and/or simpler algorithms that ignore some performance criteria. Evaluating these approaches has been frustrating because optimal solutions remain elusive for realistic problem sizes. To enable proper evaluation, we have developed a Monte Carlo approximation algorithm with good theoretical properties and robustness to workload variations. To make it computationally feasible, we combine the ideas of linear programming, randomized rounding, coresets, and iterative reweighted sampling. We show how to use this algorithm as a yardstick to evaluate other algorithms, and why it is better than other choices of yardsticks. With its help, we have demonstrated that a simple greedy algorithm works well for a number of workloads, including one generated from publicly available statistics on Google Groups. We believe that our algorithms are not only useful in their own right, but our principled approach toward evaluation will also be useful in future evaluation of solutions to similar problems in content-based publish/subscribe. This work is currently under submission.

3. Supporting Range Top-k Subscriptions. Within the same framework that we proposed in SIGMOD 2006 for supporting stateful subscriptions using a stateless dissemination network, we have been studying how to

deal with a large number of range top-k subscriptions. Generalizing the case of range-min (considered in SIGMOD 2006) is challenging, because the effect of each event on range top-k subscriptions can be substantially more complicated than the range-min case. Nonetheless, it is possible to express this effect succinctly using a set of notification messages to be sent over the stateless dissemination network, a task that reduces to tiling the affected region of a two-dimensional subscription space with the minimum number of rectangles. We have also developed techniques that further reduce the number of messages by exploiting the likely property that, in practice, the subscriptions are often clustered and may not appear in all regions of the subscription space. We are also exploring the possibility of maintaining the top-k' results, where $k' \geq k$, to further reduce notification costs. This work is currently in progress.

4. Supporting Top-k Preference Subscription. We are also considering the problem of supporting a large number of preference subscriptions over objects with multiple numeric attributes whose values are updated by events. Each such subscription specifies a preference vector whose components correspond to how this subscriber weighs different attributes. Objects are scored by their dot products with the preference vector, and the top k are maintained for the subscription. We are considering the problem in the dual space of line arrangements and developing efficient approximation techniques. This work is currently in progress.

The progress we have made in Year 3, as summarized above, is in line with our original proposal and the directions laid out in our last annual report. Several of the problems we study (2, 3, and 4 above) need additional time to finish. Because of their complexity, the need to develop a more realistic workload generator (1 above), and the training required for students, we have spent considerably more time than anticipated in the past year. Therefore, we have requested a grantee-approved no-cost extension of the project, with the new expiration date of August 31, 2011. In the final year of the project, we plan to complete our investigation of the problems described above, and continue to explore other types of subscriptions, particularly those with a geometric flavor and novel applications such as computational journalism.

In terms of educational activities, we continue to incorporate current research topics into both undergraduate and graduate curricula at Duke University. In the undergraduate database course offered by Yang in Fall 2007, two students, Kyung Min (Jason) Lee and Aneesh Ramesh Butani, worked on components of the ProSem system as their course projects. An undergraduate student, Weiping Zhang, is working on research related to the project in the summer of 2009. Yang's graduate seminar course in Spring 2008 devoted four weeks to the latest research on publish/subscribe and related areas (stream, event, and continuous query processing). Yang's graduate seminar course in Spring 2009 and 2010, devoted to data mining and databases/programming languages (resp.), introduced students to research relevant for broadening the scope of the project. Agarwal offered a course on geometric algorithms in Fall 2008, in which he covered some of topics related to this project. Agarwal also offered a course on design and analysis of algorithms in Fall 2009 in which he covered algorithmic

challenges dealing with massive data sets. A number of PhD students have worked on the project, and two (Badrish Chandramouli, Jeff Phillips) have received the PhD degree.

Findings:

We have made significant progress in approaching the problem of wide-area publish/subscribe by unifying data processing and dissemination. Published results so far include: (1) a novel way of interfacing database processing with network dissemination based on reformulation, using range aggregate subscriptions as a motivating example (SIGMOD 2006); (2) scalable processing techniques for select-join, band-join, and window-join subscriptions, using an input-sensitive approach (ISAAC 2005, VLDB 2006; extended version in TODS); (3) scalable processing and dissemination of value-based notification conditions (VLDB 2007); (4) a system featuring per-event, cost-based joint optimization of processing and dissemination (demonstrated at SIGMOD 2008); (5) scalable processing and dissemination of select-join subscriptions (VLDB 2008); (6) techniques for computing geometric summaries on data streams (SoCG 2007); (7) indexing uncertain data (PODS 2009); (8) efficient processing techniques for information extraction and retrieval (ICDE 2009, SIGMOD 2009); (9) a generator of realistic content-based publish/subscribe workloads, based on public data about Google Groups (NetDB 2009).

Training and Development:

PhD students: Badrish Chandramouli (graduated in July 2008), Shashidhar Ganjugunte, Jeff Phillips (graduate in 2009), R. Sharathkumar, Albert Yu, Ying Zheng (Spring 2008 only).

Undergraduate students: Kyung Min (Jason) Lee and Aneesh Ramesh Butani (Fall 2007), Weiping Zhang (Summer 2009), Alex Beutel and Aditya Prapaidyu (2010).

Outreach Activities:

Agarwal organized the Dagstuhl seminar on geometric algorithms in 2009; one of the topics in the seminar was on handling large data distributed over network. Agarwal is organizing a short course and a workshop on geometric computing this fall semester. He will be organizing another Dagstuhl seminar in March 2011.

Journal Publications

P. K. Agarwal, S. Har-Peled, and H. Yu, "Robust shape fitting via peeling and grating coresets", *Discrete Comput Geom*, p. 38, vol. 39, (2008). Published,

P.K. Agarwal, R. Poreddy, K. Varadarajan, and H. Yu, "Practical methods for shape fitting and kinetic data structures using coresets", *Algorithmica*, p. , vol. , (2008). Accepted,

Pankaj K. Agarwal, Junyi Xie, Jun Yang, and Hai Yu, "Input-Sensitive Scalable Continuous Join Query Processing", *ACM Transactions on Database Systems*, p. , vol. 34, (2009). Accepted,

Agarwal, PK; Bereg, S; Daescu, O; Kaplan, H; Ntafos, S; Sharir, M; Zhu, BH, "Guarding a Terrain by Two Watchtowers", *ALGORITHMICA*, p. 352, vol. 58, (2010). Published, 10.1007/s00453-008-9270-

Books or Other One-time Publications

- Pankaj K. Agarwal, Junyi Xie, Jun Yang, and Hai Yu, "Monitoring Continuous Band-Join Queries over Dynamic Data", (2005). Conference Paper, Published
Collection: Proceedings of the 16th Annual International Symposium on Algorithms and Computation (ISAAC '05)
Bibliography: Sanya, Hainan, China, December 2005
- Badrish Chandramouli, Jun Yang, and Amin Vahdat, "Distributed Network Querying with Bounded Approximate Caching", (2006). Conference Paper, Published
Collection: Proceedings of the 11th International Conference on Database Systems for Advanced Applications (DASFAA '06)
Bibliography: Singapore, April 2006
- Badrish Chandramouli, Junyi Xie, and Jun Yang, "On the Database/Network Interface in Large-Scale Publish/Subscribe Systems", (2006). Conference Paper, Published
Collection: Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data (SIGMOD '06)
Bibliography: Chicago, Illinois, USA, June 2006
- Pankaj K. Agarwal, Junyi Xie, Jun Yang, and Hai Yu, "Scalable Continuous Query Processing by Tracking Hotspots", (2006). Conference Paper, Published
Collection: Proceedings of the 32nd International Conference on Very Large Data Bases (VLDB '06)
Bibliography: Seoul, Korea, September 2006
- Badrish Chandramouli, Jeff M. Phillips, and Jun Yang, "Value-Based Notification Conditions in Large-Scale Publish/Subscribe Systems", (2007). Conference Paper, Published
Collection: Proceedings of the 33rd International Conference on Very Large Data Bases (VLDB '07)
Bibliography: Vienna, Austria, September 2007
- Badrish Chandramouli, Jun Yang, Pankaj K. Agarwal, Albert Yu, and Ying Zheng, "ProSem: Scalable Wide-Area Publish/Subscribe", (2008). Conference Paper, Published
Collection: Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data (SIGMOD '08)
Bibliography: Vancouver, Canada, June 2008
- Badrish Chandramouli and Jun Yang, "End-to-End Support for Joins in Large-Scale Publish/Subscribe Systems", (2008). Conference Paper, Published
Collection: Proceedings of the 34th International Conference on Very Large Data Bases (VLDB '08)
Bibliography: Auckland, New Zealand, August 2008
- P.K. Agarwal, S. Har-Peled and H. Yu, "Embeddings of surfaces, curves, and moving points in Euclidean space", (2007). Conference, Published
Collection: Proc. 23rd Annu. ACM Sympos. Comput. Geom.
Bibliography: Gyeongju, South Korea, June 2007
- P.K. Agarwal and H. Yu, "A space-optimal data-stream algorithm for coresets in the plane", (2007). Conference Proceedings, Published
Collection: Proc. 23rd Annu. ACM Sympos. Comput. Geom.
Bibliography: Gyeongju, South Korea, June 2007
- A. Danner, K. Yi, T. Molhave, P. K. Agarwal, L. Arge, H. Mitasova, "From elevation data to watershed hierarchies", (2007). Book, Published
Collection: Proc. ACM-GIS
Bibliography: Seattle, Washington, November 2007
- P. K. Agarwal, L. Arge, T. Molhave, and B. Sadri, "I/O efficient algorithms for computing contour lines in a terrain", (2008). Conference Proc., Published
Collection: Proc. 24th Annu. ACM Sympos. Comput. Geom.
Bibliography: College Park, Maryland, June 2008

Fei Chen, Byron Gao, AnHai Doan, Jun Yang, and Raghu

Ramakrishnan, "Optimizing Complex Extraction Programs over Evolving Text Data", (2009). Conference Paper, Accepted

Collection: Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data (SIGMOD '09)

Bibliography: Providence, Rhode Island, USA, June 2009

Pankaj K. Agarwal and Ke Yi, "Indexing Uncertain Data", (2009). Conference Paper, Accepted

Collection: Proceedings of the 2009 ACM Symposium on Principles of Database Systems (PODS '09)

Bibliography: Providence, Rhode Island, USA, June 2009

Albert Yu, Pankaj K. Agarwal, and Jun Yang, "Generating Wide-Area Content-Based Publish/Subscribe Workloads", (2009). Technical Report, Submitted

Bibliography: Technical Report,

Department of Computer Science, Duke University, June 2009

Pankaj K. Agarwal, Esther Ezra, and Shashidhar Gujgunte, "Efficient Sensor Placement for Surveillance Problems", (2009). Conference Paper, Published

Collection: Proceedings of the IEEE 2009 Conference on Distributed Computing in Sensor Systems (DCOSS '09)

Bibliography: Marina Del Rey, California, USA, June 2009

Risi Thonangi, Hao He, AnHai Doan, Haixun Wang, and Jun Yang, "Weighted Proximity Best-Joins for Information Retrieval", (2009). Conference Paper, Published

Collection: Proceedings of the 25th International Conference on Data Engineering (ICDE '09)

Bibliography: Shanghai, China, March 2009

P.K. Agarwal, S. Har-Peled, M. Sharir, and Y. Wang, "Hausdorff distance under translation for points and balls", (2010). journal article, Published

Collection: ACM Transactions on Algorithms

Bibliography: 6 (2010), 1–-26

P. K. Agarwal, J. Phillips, and B. Sadri, "Lipschitz unimodel and isotonic regression on paths and trees", (2010). conference paper, Published

Collection: Ninth Annual Latin American Theoretical Informatics Symposium

Bibliography: 2010

P. K. Agarwal, S. Sankararaman, A. Efrat, and S. Ramasubramanian, "On channel-discontinuity-constraint routing in wireless networks", (2010). conference paper, Published

Collection: IEEE INFOCOM Mini-Conference

Bibliography: 2010

P. K. Agarwal, B. Aronov, M. van Kreveld, M. L'offler, and R. I. Silveira, "Computing similarity of piecewise-linear functions", (2010). conference paper, Published

Collection: Twenty-Sixth Annual Symposium on Computational Geometry

Bibliography: 2010

P. K. Agarwal, "An improved algorithm for computing the volume of the union of cubes", (2010). conference paper, Published

Collection: Twenty-Sixth Annual Symposium on Computational Geometry

Bibliography: 2010

P. K. Agarwal, R. Ben Avraham and M. Sharir, "The 2-center problem in three dimensions", (2010). conference paper, Published

Collection: Twenty-Sixth Annual Symposium on Computational Geometry

Bibliography: 2010

P. K. Agarwal, J. Gao, L. Guibas, H. Kaplan, V. Koltun, N. Rubin, and M. Sharir, "Kinetic stable Delaunay graphs", (2010). conference paper, Published

Collection: Twenty-Sixth Annual Symposium on Computational Geometry

Bibliography: 2010

P. K. Agarwal, T. M?lhave, L. Arge, and M. Revsb?k, "Scalable algorithms for large high-resolution terrain data", (2010). conference paper, Published

Collection: First International Conference on Computing for Geospatial Research and Applications

Bibliography: 2010

P. K. Agarwal, J. Phillips and H. Yu, "Stability of epsilon-kernels", (2010). conference paper, Published

Collection: Eighteenth European Symposium on Algorithms

Bibliography: 2010

P. K. Agarwal, A. Efrat, S. K. Ganjugunte, D. Hay, S. Sankararaman, G. Zussman, "A probabilistic model for network vulnerability under physical attack", (2010). conference paper, Published

Collection: Military Communications Conference

Bibliography: 2010

Albert Yu, Pankaj K. Agarwal, and Jun Yang, "Generating Wide-Area Content-Based Publish/Subscribe Workloads", (2009). Workshop Paper, Published

Collection: Proceedings of the 5th International Workshop on Networking Meets Databases (NetDB '09)

Bibliography: Big Sky, Montana, USA, October 2009

Web/Internet Site

URL(s):

<http://www.cs.duke.edu/dbgroup/prosem/>

Description:

Other Specific Products

Contributions

Contributions within Discipline:

At the end of Year 3 of the project, we have made a series of solid contributions to the area of publish/subscribe system and its related area of continuous query processing. A number of our contributions have been published: ISAAC 2005, DASFAA 2006, SIGMOD 2006, SoCG 2006, VLDB 2006, SoCG 2007, VLDB 2007, ACM-GIS 2007, SIGMOD 2008, SoCG 2008, VLDB 2008, SODA 2009, ICDE 2009, DCOSS 2009, SoCG 2009, PODS 2009, SIGMOD 2009, NetDB 2009, TODS (2009), SODA 2010, SoCG 2010, INFOCOM 2010, MILCOM 2010, ESA 2010. For detailed descriptions of these contributions, please refer to the section of this report on research and education activities.

Contributions to Other Disciplines:

The PIs are part of another NSF project, in which they are developing techniques for managing dynamic sensor network data for ecological modeling. Some of the work done in this project is influenced by the problems they have faced in the other project, and vice versa.

Contributions to Human Resource Development:

PhD students: Badrish Chandramouli (graduated in July 2008), Shashidhar Ganjugunte, Jeff Phillip (graduated in 2009), R. Sharathkumar, Albert Yu, Ying Zheng (Spring 2008 only).

Undergraduate students: Kyung Min (Jason) Lee and Aneesh Ramesh Butani (Fall 2007), Weiping Zhang (Summer 2009), Alex Beutel and Aditya Prapdiyu (2010).

Contributions to Resources for Research and Education:

The PIs are integrating the research being carried out under this project into the classes they are teaching (geometric algorithms, databases, geometric optimization, sensor networks, and data mining).

A major obstacle for research on publish/subscribe has been the lack of publicly available, realistic workloads, because of concerns of privacy and commercial interests in releasing user information. This project has developed a workload generator for wide-area content-based publish/subscribe systems, by extrapolating the limited amount of public data. A paper describing this work has been published in NetDB 2009, and the generator is now publicly available to the research community.

Contributions Beyond Science and Engineering:

Conference Proceedings

Agarwal, PK;Sharathkumar, R, Streaming Algorithms for Extent Problems in High Dimensions, "JAN 17-19, 2010", PROCEEDINGS OF THE TWENTY-FIRST ANNUAL ACM-SIAM SYMPOSIUM ON DISCRETE ALGORITHMS, 135: 1481-1489 2010

Special Requirements

Special reporting requirements: None

Change in Objectives or Scope: None

Animal, Human Subjects, Biohazards: None

Categories for which nothing is reported:

Organizational Partners

Any Product

Contributions: To Any Beyond Science and Engineering