SENSOR: Data-Driven Sensor Networks

Final Report (September 2003 to August 2006): Activities and Findings

Johannes Gehrke Rajmohan Rajaraman

In this project, we adopted a database approach to unite the seemingly conflicting requirements of scalability and flexibility in monitoring the physical world. The objective of this research was to build a new distributed data management layer that scales with the growth of sensor interconnectivity and computational power on the sensors over the next decades. Our system is designed to create the abstraction of a single processing node without centralizing data or computation. Since sensor networks are heavily resource-constrained, our research focused on *cross-layer optimizations* that are needed to exploit opportunities for the preservation of resources. Due to the regularity of query processing patterns we designed query-layer specific routing algorithms that are optimized — not for general point-to-point communication — but for the more regular types of communication patterns that are generated by a query layer.

The main results of this project include a new MAC layer protocol that explores energy-delay tradeoffs, optimization techniques for processing multiple queries, a new approach to data aggregation based on the paradigm of universal optimization, new algorithms for efficient event notification in sensor networks, and a new algorithm for efficient data archival in sensor networks. The project also had an REU component in which we implemented a method for locating cars using a sensor ad hoc network. We now elaborate on the specific activities and findings.

1 Activities and Findings

- **Group-oblivious aggregation and multicast:** Consider a sensor network in which a gateway node t is collecting aggregate information from sources chosen from a set N of nodes. A common structure for performing an aggregation is a Steiner tree that connects the sink t to the subset of sources. However, in many scenarios, such as event monitoring, the particular group of sources with data to aggregate may change dynamically and may only be known at the time the queries are issued. In this scenario, it may be prohibitively expensive to set up a new Steiner tree on-line for each aggregate query that is issued. A natural question to ask is the following: is there a single spanning tree T such that for *any subset* S of N, the Steiner tree obtained by the subtree of T induced by S is a good approximation to the optimal Steiner tree for S.

  In joint work with colleagues Guevara Noubir and Ravi Sundaram and doctoral students Lujun Jia and Guolong Lin, we have studied the design of *universal Steiner Trees*, structures that approximate optimal Steiner trees for every possible subset of terminals [2]. We show that given a metric space over n vertices and a root vertex, we can find a universal spanning tree such that for any subset of vertices containing the root, the sub-tree induced by the subset is within $O(\log^4 n/\log\log n)$ of an optimal Steiner tree for the subset. Our algorithms rely on a new notion of sparse partitions, that may be of independent interest. For the special case of doubling metrics, which includes both constant-dimensional Euclidean and growth-restricted metrics, our algorithm achieves an $O(\log n)$ upper bound. We complement our results for the universal Steiner tree problem with a lower bound of $\Omega(\log n / \log \log n)$ that holds even for n vertices on the plane. We also show that a slight generalization of the universal Steiner Tree problem is coNP-hard. Our paper also contains other results for universal versions of the classical optimization problems TSP and Set Cover.

Our universal Steiner tree design provides an attractive approach for group-oblivious aggregation and multicast in sensor networks. In joint work with Lujun Jia, Guevara Noubir, and Ravi Sundaram, we have developed a practical and distributed protocol GIST (Group-Independent Spanning Tree) for realizing group-oblivious aggregation that requires only local knowledge. Through extensive simulations, we have shown that not only does GIST outperform traditional aggregation methods in the worst case, it outperforms them by between 30 and 60 per cent in realistic random scenarios [4].

- **Multi-query optimization:** In joint work with Al Demers, Niki Trigoni, and Yong Yao, we have developed a set of distributed algorithms for multi-query optimization in sensor network databases [7]. Our techniques are aimed at addressing a wide class of long-running as well as snapshot queries in an energy-efficient manner so as to maximize network liftetime. In particular, for queries that can be represented as linear combinations of sensor values, which includes *sum*, *avg* and *count*, we connect this problem to the problem of finding a basis for a set of vectors and give algorithms yielding optimal result compression. We show that the proposed multi-query optimization techniques are robust to node or link failures. By adding a small communication cost, our techniques continue to work even if the dissemination tree changes due to local repairs (between rounds) or global tree reconstruction (between epochs).

- **Event notification:** A canonical application of sensor networks is event monitoring, in which sensor nodes identify various events defined by physical phenomena (environmental parameters, animal localization, etc.), and convey the necessary information to designated information sinks. We have studied event notification schemes, by which a collection of information sinks can obtain information about all the events occurring within a sensor network, given a diversity of query frequencies (at the sinks) and event frequencies (at the source sensor nodes identifying the events) [1]. In recent work, we have designed algorithms that optimize communication cost by combining multicast "pulls" initiated by the sinks with multicast "pushes" initiated by the sources. For the special case of tree networks, we have developed an optimal algorithm, that is simple and amenable to a distributed implementation. We also have obtained logarithmic approximation ratios for the general case. In addition, we have studied event notification in communication models that do not allow multicast. We consider standard point-to-point unicast and a controlled broadcast model that is particularly suitable for ad hoc sensor networks. We have near-optimal approximation algorithms for certain special cases [1], and are studying the general case in ongoing research.

- **Wave scheduling:** In joint work with Al Demers, Niki Trigoni, and Yong Yao, we have developed a novel methodology for trading energy versus latency in sensor database systems [6, 5]. We proposed a new protocol, called WaveScheduling, that carefully schedules message transmissions so as to avoid collisions at the MAC layer. Since all nodes adhere to the schedule, their radios can be off most of the time and they only wake up during well-defined time intervals. We showed how routing protocols can be optimized to interact symbiotically with the scheduling decisions, resulting in significant energy savings at the cost of higher latency.

We demonstrated the effectiveness of our approach by means of a thorough simulation study in which we compare WaveScheduling with a tree-based scheduling protocol (as used in the Tag system of Berkeley) and the IEEE 802.11 protocol (with varying duty cycles). Our experiments indicate that the three approaches provide a tradeoff between energy consumption and latency. In particular, the WaveScheduling approach is the most energy-efficient efficient protocol and is suitable for latency-tolerant applications.

- **Power control in ad hoc wireless networks:** In joint work with Guevara Noubir and doctoral students Lujun Jia and Xin Liu, we have developed a new class of power control schemes for IEEE 802.11 that explores the tradeoff between spatial reuse and collision avoidance to obtain significant improvements both for energy consumption and network capacity, with minimum modifications to the standard [3]. Through extensive simulations, we have studied the performance of this class on various measures including throughput, energy cost, and fairness to transmission distance. We show that there exists a scheme in this class that near-optimally maximizes the throughput and significantly reduces energy consumption. This scheme improves upon the throughput achieved by IEEE 802.11 by 30% to 40%, increases the packets delivered per Joule by a factor of 3, and shows better fairness toward short and long range communication in terms of achieved traffic. We have also provided a theoretical justification for our simulation results by means of an analysis under a simplified communication model.

- **Network Scheduling for Data Archiving Applications in Sensor Networks.** In joint work with Yong Yao, Nazrul Alam, and Sergio Servetto, we investigated power management in sensor networks. Since data archiving in sensor networks is a communication intensive application, careful power management of communication is of critical importance for such networks. An example is FPS, an adaptive power scheduling algorithm that combines slotted scheduling with a CSMA MAC. We propose a new global power scheduling protocol called Multi-Flow Power Scheduling (MPS) that delivers more data and consumes less energy than existing power scheduling protocols. MPS sets up a transmission schedule through standard data aggregation and dissemination operations, however since it creates a global schedule it does not scale to large networks. We also present a new power scheduling protocol called Hybrid Power Scheduling (HPS) that retains the scalability of FPS while maintaining the energy efficiency and high data delivery rate of MPS. In a thorough simulation study, we compare HPS and MPS, and our results show the efficacy of HPS [8].

The REU Component: The CAR-WHERE Project. The REU component of this project partially supported the research work of Northeastern undergraduate student Nathan Faber in 2006, who developed a protoype system for locating cars using an ad hoc sensor network. The system enables a use to issue queries from a laptop that are communicated through a network of motes. When the target car (a toy car in the prototype) hears the query, it broadcasts a signal. When a node receives a response from the car, it measures the signal strength and includes this in its response, which is propagated along the reverse path of the query flood, to the query source.

This project was jointly advised with colleague Ravi Sundaram, who took the leadership role in the project. A demonstration prototype is in our lab, a preliminary version of which also won Nathan the best undergraduate research project within Northeastern University. Our original REU proposal also included implementing a parking database and more sophisticated queries. Unfortunately, we were not able to complete it during the summer, and our (late) no-cost extension request was not granted.

2 Training and Development

2.1 Training of the Next Generation Workforce

2.1.1 Postdocs

This project partially supported a postdoc, Niki Trigoni. Niki has a PhD degree from Cambridge University in Great Britian, and she joined Cornell for nearly two years as a postdoc working on query processing in sensor networks. It is fair to say that the postdoc position helped her to launch a successful academic career, since she first obtained a faculty position at Birkbeck College in London in 2005, and she will join Oxford University in late 2007.

2.1.2 PhD Students

This project supported several doctoral students.

- Lujun Jia, who was supported on this grant, completed a dissertation on communication structures for wireless ad hoc networks, with a focus on sensornets. He defended his thesis in Fall 2005, and had several publications related to the project. He is currently a Distinguished Member of Technical Staff at Verizon Labs.

- Guolong Lin, who was partially supported by this grant, worked on approximation algorithms for network design in uncertain and dynamic environments. The work on universal Steiner trees for sensor networks is an instance of this research. Guolong defended his thesis in Spring 2006, and is now a Senior Software Engineer at Akamai Technologies.

- Adina Crainiceanu, who was supported by this grant, graduated in December 2005, and she is now an Assistant Professor at the Naval Academy in Annapolis. Adina's thesis investigates methods for building highly fault-tolerant and available distributed data management systems.

- Abhinandan Das who was supported by this grant graduated in August 2005; his thesis investigated the use of randomized projections, so-called "sketches", in particular for spatial data as generated by small-scale sensors. Abhinandan joined Google.

- Yong Yao graduated in January 2007; his thesis was on query processing in sensor networks. He joined IBM Santa Theresa Lab where he will work on a data stream management system. According to Yong, the fact that he was knowledgable in both sensor networks and thus wireless communication *and* data management was a deciding factor for many employers during his job search.

- Ashwin Machanavajjhala has gained valuable experience in building distributed systems through work on this grant, but he has now moved on to work on data privacy as his thesis topic.

- Selcuk Aya worked on the implementation of a fault-tolerant data management layer. Unfortunately, for personal reasons, Selcuk decided to leave the PhD program at Cornell, and he is now working in industry.

- Biswanath Panda worked on a network abstraction for data management in sensor networks which he implemented in C++. In particular, the network abstraction has a single API that can be moved between a simulation environment and real sensors, allowing higher-level components to be programmed independently of whether the system is run as a simulation or in a real environment. Biswanath's has now moved on to apply his experience to scientific simulations (outside of this grant), where he is very successful and he plans to graduate in the summer of 2008.

- Christian Bucila investigated the possibility of applying data mining techniques to sensor network data, however the results were disappointing in that no new data mining techniques came out of this work. However, this initial work led to our later work on data compression in sensor networks as that seemed to be the one important operation that users cared about.
- Xin Liu is a senior doctoral student at Northeastern University, and is scheduled to complete his thesis work by the end of Summer 2007. Her research work has been on cross-layer design for cooperative and adversarial wireless networks. She worked on a MAC-layer power control project related to this grant.

2.2 Tutorials

In June/July 2004, co-PI Rajmohan Rajaraman presented a tutorial on 'Algorithmic Foundations for Ad Hoc Networks' at a Summer School in ETH Zurich. This tutorial had a significant component on sensor network algorithms. The tutorial slides and bibliography are available at

`http://www.ccs.neu.edu/home/rraj/Talks/AdHocTutorial/ETHZ/`

The site acknowledges NSF funding of this grant.

2.3 Curriculum Development

In Spring 2005, co-PI Rajmohan Rajaraman taught a graduate course "Wireless Networks" which had a substantial component on ad hoc and sensor network algorithms, closely related to the research conducted in this project. This course is also being taught in the Spring 2007 semester. It has had an enrolment of over 20 MS and PhD students. After covering the fundamentals of wireless communications (roughly a third of the course), we studied widely-used wireless local area network standards and TCP/IP extensions for mobile and wireless networking (another third). The final third of the course was devoted to algorithms for multihop ad hoc networks and sensor networks. Topics in this component included routing and MAC protocols, topology and power control, capacity of ad hoc networks, and query processing algorithms for sensor network databases. The course also includes several simulation assignments, a sensor network programming project using Tmotes. More information about the course, including lecture notes, may be obtained from the course web pages, both of which acknowledge NSF funding of this grant.

`http://www.ccs.neu.edu/home/rraj/Courses/G250/S05/www/`
`http://www.ccs.neu.edu/home/rraj/Courses/G250/S07/www/`

References

[1] R. Chakinala, A. Kumarasubramanian, R. Manokaran, K. Laing, C. P. Rangan, and R. Rajaraman. Playing push vs pull: Models and algorithms for disseminating dynamic data in networks. In *Proceedings of the ACM Symposium on Parallelism in Algorithms and Architectures*, pages 244–253, July 2006.

[2] L. Jia, G. Lin, G. Noubir, R. Rajaraman, and R. Sundaram. Universal approximations for TSP, Steiner Tree, and Set Cover. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 386–395, May 2005.

[3] L. Jia, X. Liu, G. Noubir, and R. Rajaraman. Transmission power control for ad hoc wireless networks: Throughput, energy and fairness. In *Proceedings of IEEE WCNC*, March 2005.

[4] L. Jia, G. Noubir, R. Rajaraman, and R. Sundaram. Gist: Group-independent spanning tree for data aggregation in dense sensor networks. In *DCOSS*, pages 282–304, 2006.

[5] N. Trigoni, Y. Yao, A. Demers, J. Gehrke, and R. Rajaraman. WaveScheduling: Energy-Efficient data dissemination for sensor networks. *ACM Transactions on Sensor Networks*. To appear.

[6] N. Trigoni, Y. Yao, A. Demers, J. Gehrke, and R. Rajaraman. WaveScheduling: Energy-Efficient data dissemination for sensor networks. In *Proceedings of the VLDB Workshop on Data Management for Sensor Networks (DMSN)*, 2004.

[7] N. Trigoni, Y. Yao, A. Demers, J. Gehrke, and R. Rajaraman. Multi-query optimization for sensor networks. In *Proceedings of the IEEE International Conference on Distributed Computing in Sensor Systems*, June 2005.

[8] Y. Yao, S. Alam, J. Gehrke, and S. D. Servetto. Network scheduling for data archiving applications in sensor networks. In *Proceedings of the International Workshop on Data Management for Sensor Networks (DMSN 2006)*, 2006.