

Collaboratively Querying Sensor Networks through Handheld Devices

Tsz Wai Chiu Qiong Luo

*Department of Computer Science and Engineering
The Hong Kong University of Science and Technology
{cswai, luo}@cse.ust.hk*

Abstract

We envision that in some wireless sensor network applications, such as environmental monitoring, assisted living, and industrial control, handheld devices will be used from time to time to query the sensor networks. However, there is no full-fledged query processor for this purpose. Therefore, we propose WinyDB, a relational query processing system on Windows-CE based PDAs (Personal Digital Assistants) for sensor networks. One of the main features of WinyDB is that multiple PDAs running WinyDB can answer queries collaboratively. This feature is useful in that it improves both the energy efficiency and the data quality. Our WinyDB prototype package is available online at <http://www.cse.ust.hk/winydb> and our simulation experiments have shown promising results on collaborative query processing.

1. Introduction

Wireless sensor networks (WSNs) have been widely deployed in applications such as habitat monitoring, medical care, intelligent building, and surveillance [3][5][6]. A common way to acquire sensory data from these WSNs is to issue a query from a PC-grade base station to the WSNs and have the sensors return their query results to the base station [17][18]. There has also been some work [3][11][14] on using handheld devices to collect WSN data, but the query processing capabilities of these data collection tools are limited. In this paper, we propose WinyDB, an embedded query processing system on Windows-CE based PDAs (Personal Digital Assistants) to query WSNs.

We envision that WinyDB will be useful for many WSN applications. For example, in a country park with WSNs deployed in the natural environment, Wi-Fi covered in the air, and WinyDB-installed PDAs provided to the park rangers and the visitors, people can issue queries about the environment through their PDAs as they move around. Furthermore, if multiple PDAs are roaming around in the park, they can also help one another answering queries about their local WSNs. These queries can be answered conveniently and promptly without contacting a central server. Most importantly, this collaborative query answering will save the energy consumption of WSNs as well as improve the data quality, because the data acquisition is done locally and is shared among multiple queries.

Now that mobile devices become base stations for their local WSNs, we need to consider the resource limitations of the devices as well as the WSNs in collaborative query processing. Specifically, we address the following questions:

1. What query processing capabilities do we support on PDAs? Given multiple WSNs deployed in a wide area, how do we answer a query involving multiple streams of sensory data?
2. Given queries of different kinds running on multiple PDAs for multiple WSNs, what collaborative query processing schemes do we design to answer these queries?
3. Given the resource constraints of WSNs and PDAs, how we consider these constraints in our cost model to direct WinyDB to process queries efficiently?

The remainder of this paper is organized as follows. Section 2 gives an overview of WinyDB and Section 3 presents our collaborative query processing framework. We discuss experimental results in Section 4, review the related work in Section 5, and conclude in Section 6.

2. System Overview

We consider multiple WSNs deployed in a wide area, e.g., in a park. Each WSN has a unique network ID and a sink node. A PDA may roam into any of the WSNs and can communicate with the sink node of the WSN if the sink node is within the PDA's communication range. Both PDA-WSN and PDA-PDA communications are via IEEE 802.11b.

WinyDB consists of three software components - *query processing (QP)*, *peer collaboration (PeerCo)*, and *resource management (ResM)*. The QP module is in charge of query parsing, query optimization, and query result processing. Because PDAs move around, and queries are issued and removed dynamically, the PeerCo module on each PDA keeps track of the current status of its peer PDAs and coordinates collaboration with these devices. Finally, the ResM module monitors the consumption of the resources, including WSN energy and PDA memory.

When a PDA issues a query, the query is parsed and optimized by the QP module at this *query issuer*. During the query optimization, QP consults PeerCo to check whether any other PDAs (we call them *data providers*) can collaborate to answer part of or the whole query. Then it generates candidate query plans, works with ResM to estimate the goodness of each plan, and selects the best plan. This best plan may be to send

the query to the WSNs directly or to share the query results from some other PDAs. Finally, QP works with PeerCo to execute the query and to return query results to the user. Since most of the queries are continuous and the environment is dynamic, QP may change a query plan to adapt to the current setting.

3. Collaborative Query Processing

Our database model is similar to those in Cougar [18] and TinyDB [17], where a sensor network is represented as a virtual table and every sensor node generates a data tuple for each sample period. The difference is that in our work, this virtual table contains multiple physical WSNs, each of which is identified by a network ID. We focus on continuous queries as opposed to snapshot queries, because they create more opportunities for as well as benefit more from collaborative query processing.

The goals of our collaborative query processing are to improve data quality and to reduce WSN energy consumption. These two goals may be conflicting; when this happens, we take a user-centric stand and put data quality before resource efficiency. We measure the data quality of a given query in the query result output rate, because it reflects both the network efficiency and the data quality. We estimate WSN energy consumption by the amount of data communication. Additionally, because a PDA is resource constrained, we take its available memory into consideration.

In the following, we present our collaborative query processing schemes for different types of queries.

3.1. Selection, projection queries

We consider selection and projection queries with numeric comparison predicates. Note that, even though a query appears as a single table query, it may involve multiple physical WSNs unless the network ID of a WSN is specified in the selection condition. Query 1 is an example of selection query.

Query 1:
 SELECT nodeid, temperature
 FROM sensors
 WHERE NetworkID=1 AND temperature>30
 SAMPLE PERIOD 60s

If the query issuer is out of the range of the sink of WSN 1, it must ask for the help of the PDAs that are in the range of WSN 1, which we call the *candidate data providers* for WSN 1; otherwise, the query issuer itself is also a candidate data provider and it will probe itself first before probing other candidate data providers.

A candidate data provider for a WSN can answer a query either by (1) directly issuing the query to the WSN or by (2) evaluating the query on the cached results of another query. We call option (1) *direct evaluation* and (2) *cache evaluation*. We always choose cache evaluation over direct evaluation whenever applicable, because it can reduce the WSN workload and therefore improve the data quality and save the WSN energy consumption.

The cache evaluation of query A over the result of query B is applicable when query B *contains* query A [8]. Given selection and projection queries A and B in their conjunctive normal forms (CNFs), we determine

that B contains A when (1) B's list of attributes is a superset of A's, (2) each of B's selection predicate conjunct is less restrictive than one of A's, and (3) if A is a continuous query, its sample period is a multiple of B's. For example, $temperature > 30$ is less restrictive than $temperature > 40$ or $temperature = 50$.

When multiple candidate data providers can answer a query by cache evaluation, we choose the one with the highest estimated result output rate. A higher output rate is preferred by the user and usually indicates the data provider has a lighter workload and/or a better capability.

In WinyDB, each candidate data provider provides its estimated result output rate for a given query. This estimation is done by examining the *input rate*, the *processing rate*, the *transmission rate*, the *buffer size*, and the *selectivity*. The input rate is the number of input tuples from a WSN or from another data provider per unit of time. The processing rate is the number of result tuples processed per unit of time. The transmission rate is the number of result tuples that can be transmitted to the query issuer per unit of time. The buffer size is the maximum number of result tuples that can be cached. The selectivity is the number of final result tuples divided by the number of input tuples.

Based on the input rate (in_t), selectivity (sel_t), processing rate (p_t), buffer size (buf_t), and the transmission rate (x_t) at time t, the output rate (out_t) during a sample period $[t_1, t_2]$ can be estimated in one of the following four cases in Table 1. Only for the four cases in Table 1 will we consider cache evaluation on the data provider. Otherwise, we will exclude this data provider from performing cache evaluation because its buffer will overflow and it will lose query results.

Table 1. Output Rate Estimation

Cases	out_t
$in_t \leq p_t$ and $sel_t * in_t < x_t$	$\min(sel_t * in_t, x_t)$
$in_t \leq p_t$ and $sel_t * in_t \geq x_t$ $\int_{t_1}^{t_2} (sel_t * in_t - x_t) dt \leq buf_t$	$\min(sel_t * in_t, x_t)$
$in_t > p_t$ and $sel_t * p_t \leq x_t$ $\int_{t_1}^{t_2} (in_t - p_t) dt \leq buf_t$	$sel_t * p_t$
$in_t > p_t$ and $sel_t * p_t > x_t$ $\int_{t_1}^{t_2} (in_t - p_t + sel_t * p_t - x_t) dt \leq buf_t$	x_t

If a data provider has no cache evaluation alternatives for a given query, it will estimate its output rate for a direct evaluation of the query over the WSN. In this case, we estimate its output rate as follows:

$$out_t = \min(p_t, x_t)$$

Finally, we handle the case of using multiple data providers to answer a given query. In this case, each of the data providers can answer part of the query through cache evaluation or direct evaluation. Therefore, the query issuer generates alternative plans, estimates the

output rates of each alternative, and selects the final query plan.

3.2. Aggregation Queries

We support seven common aggregate functions, including MAX, MIN, AVG, SUM, COUNT, COUNT DISTINCT, and MEDIAN. We classify these aggregate functions into two categories, *divisible* and *indivisible*. Aggregate functions MAX, MIN, COUNT, and SUM are divisible, because they can be evaluated on partitions of the result set and the final aggregation value can be obtained by combining the partial aggregates of the partitions. In contrast, aggregate functions COUNT DISTINCT, and MEDIAN are indivisible, because partial aggregation is inapplicable to them. Finally, AVG itself is indivisible, but we treat it as divisible by representing it in the form of (SUM, COUNT) so that we can apply partial aggregation.

We handle an aggregation query in a way similar to selection and projection queries. Specifically, it can be done either through direct evaluation or cache evaluation. The cache evaluation of an aggregation query can use the results of existing selection and projection queries. Furthermore, a query with a divisible aggregate function can be evaluated at a candidate data provider and the final aggregation is performed at the query issuer by combining the partial aggregation values. However, for a query with an indivisible aggregate function, we can apply the aggregate function only at the query issuer, not at any other data provider.

3.3. Window Join Queries

A window join can be between different WSNs or within one WSN. The window is time-based. We consider only binary joins and the join predicate can use any of the numeric comparison operators.

For simplicity, we apply window join predicates only at PDAs, not at individual sensor nodes. We make this design decision because join is an expensive operation and therefore is more suitable to perform at PDAs than at sensor nodes. However, the downside of the PDA-side join approach is that many input tuples may be useless when the join is highly selective. Therefore, we have developed a synopsis based prediction method for two input sources of an equi-join to decide on whether to discard an input tuple before performing the join. This synopsis works for one or multiple join attributes, with one synopsis per join attribute.

3.3.1. Join Synopsis

We build a join synopsis at the two data providers of a join to predict if an input tuple may be discarded or not for the join. These two data providers can be the same PDA or different ones. The query issuer may be one of them or a third PDA. This join synopsis works in two phases:

1) Start-up Phase

Suppose the two data providers are A and B. Each of the data providers determines the range of its join attribute values from its local WSN. Then, they send the ranges to the query issuer. The query issuer thereby

determines a common range and a number of equal-width partitions, k . This common range covers both A and B's ranges, and the choice of k is dependent on the free memory space of the query issuer. A larger k incurs higher memory usage but leads to further reduction of communication and improvement of output rate.

Next, both A and B start to construct their join synopses. A join synopsis is a bit vector summarizing the distribution of the join attribute value. Each bit in a k -bit vector corresponds to one of the k partitions of the join attribute value over the common range. When the join attribute value falls into a certain partition, the corresponding bit is set to 1; otherwise, the bit remains to be 0. This bit vector applies to the tuples within the current join window.

For each window join, A and B exchange their bit vectors through the query issuer and only send the tuples that fall into the common partitions to the query issuer. Additionally, these bit vectors are used to determine what tuples to discard when the memory size is limited.

2) Maintenance Phase

As the sensor data values may change over time, the bit vectors and the common range of data values may change too. Therefore, we update the bit vector for each sample period. If it differs from the last sample period, it is sent to the query issuer. To prevent the common range and the number of partitions, k , from changing frequently, we expand the range to a certain degree so that it usually holds for the full length of a join window. In our experiments with the Intel Berkeley dataset [7], it is sufficient to expand this common range by 10%.

When the number of partitions is small or the data distribution is skewed, many data values may fall into one partition. When this happens, few tuples may be discarded according to the bit vectors. Therefore, we maintain join statistics at the query issuer. Specifically, for each partition, the query issuer records two counters - the number of tuples received from the data provider and the number of tuples that satisfy the join condition. With these two counters, we know which partitions contain many tuples but only a few of the tuples satisfy the join condition. The data provider then further divides up such a partition by constructing a child bit vector on the partition. Similarly, if the two counter values associated with a partition are similar, the data provider may decide to remove its child partitions, if any.

3.3.2. Collaborative Join Processing




Figure 1. Cache Evaluation on other PDAs

Given a new window join query Q on two WSNs, its collaborative processing is either (1) cache evaluation

on another PDA or (2) cache evaluation on two other PDAs.

Figure 1(a) illustrates the cache evaluation on another PDA. PDA X is the query issuer and PDA Y is the data provider. If there is an existing join query on Y that contains the new query, all processing can be done on PDA Y and only the final results are forwarded to PDA X. If there are only selection or projection queries on Y to be used for cache evaluation, the join processing can be done either on Y or on X. The choice of plan depends on the result output rate.

If the join query involves two other PDAs each of which connecting to a WSN, there are two possible communication schemes – either data shipping (Figure 1(b)) or query shipping (Figure 1(c)).

When a join query involves more than one data providers, the coordination of data providers is essential. The simplest way is to use data shipping as shown in Figure 1(b). All data are forwarded to and joined at the query issuer. However, this approach may not be cost-effective even with our join synopsis method applied, because data communication may be reduced if they are joined between the data providers before forwarded to the query issuer, as shown in the query shipping in Figure 1(c). In addition, the memory and processing requirement of the query issuer can be shared by the data providers in query shipping.

For simplicity, data shipping is used as the default at start-up. One problem with data shipping is that the query issuer may not have enough memory to cache the results from the two data providers. When this happens, the query issuer will check whether the data providers can satisfy the memory requirement for query shipping.

Suppose the query issuer is joining the results from data providers i and j using data shipping. The memory requirement for query shipping will be either data provider i 's buffer can cache (R_j) results from data provider j , or vice versa. If neither data providers satisfy this memory requirement, query shipping cannot be used. If exactly one of the data providers satisfy the memory requirement, query shipping will be used and the join will be moved to that data provider. If both data providers satisfy the memory requirement, the query issuer will determine the join site. If Formula (1) holds, the join will be at data provider i ; otherwise, it will be at data provider j .

$$N(R_i) > N(R_j) \quad (1)$$

R_i is data provided by data provider i . $N(R_i)$ is the size of R_i .

Even when the query issuer has enough memory for data shipping, query shipping will be considered for possible reduction of data communication over data shipping. This consideration is represented by Formula (2).

$$N(R_i) + N(R_j) > \min(N(R_i), N(R_j)) + N(R_i) \triangleright \triangleleft_c R_j \quad (2)$$

To avoid oscillating between data shipping and query shipping, we set a threshold for the confidence level of Formula (2). This confidence level is calculated as the ratio of the number of sample periods that Formula (2) holds to the total number of sample periods. We set the threshold to be 80% in our experiments. Only when Formula (2) holds with a

confidence level above the threshold will we consider switching to query shipping. Similar to the case of the query issuer having insufficient memory for data shipping, we will consider the memory requirement of the data providers and determine the join site for query shipping.

4. Evaluation

We have implemented WinyDB in C# using Visual Studio .NET 2003 and .NET Compact Framework 1.0. We call this version the optimized WinyDB.

For comparison, we have ported TinyDB [17] of version 1.1.10 to run on PDAs and have extended it to process window joins. We call it the original WinyDB. In the original WinyDB, the query issuer can query multiple WSNs through other PDAs. These data providers will simply execute the query through direct evaluation and forward the results to the query issuer.

Since we focus on continuous queries for sensor data streams, we select the result output rate as the end-to-end performance metric. This metric reflects the network traffic in the WSNs and the packet loss rate. Also, as energy consumption is the critical constraint for WSNs, we select the average sensor node energy consumption as the other performance metric.

4.1. Experimental Setup

To evaluate WinyDB in a variety of settings, we ran WinyDB on three real PDAs to query simulated WSNs. Two of the PDAs each has a 624 MHz processor and 64 MB memory, running Windows mobile 2003. The third PDA has a 400 MHz processor and 64 MB memory, running Windows mobile 2002.

We used TOSSIM [13] to simulate WSNs of different configurations. The simulated WSNs consist of tens of MICA2 motes [2] as specified in Table 2. These nodes are randomly deployed in a grid. The sensor readings at each node are simulated using the real data from the Intel Lab data set [7].

In every experiment, each of the three PDAs connects to a WSN so that any queries about any of the three WSNs can be answered, with or without query sharing. All queries are continuous queries with sample periods varies from 8 to 64 seconds. The number of attributes returned in each query is between 1 and 4, both inclusive. Each attribute that appears in the query condition involves at most one selection predicate. The presence of selection condition, aggregation function, and join condition is all randomly generated. The window size of join queries varies from 40 to 120 seconds. The number of concurrent queries on each PDA is varied from 2 to 4. The reason for this small number is that running a large number of queries concurrently overwhelms TOSSIM.

We also conducted another set of experiments to evaluate the effect of memory size on join queries. In these join query experiments, one PDA issues a window join query with a fixed sample period of 8 seconds, whereas the other two PDAs issue queries of all types by random. The available memory is varied from 25% to 100% of the join window size.

To estimate the energy consumption of multiple queries running in the WSNs, we logged in each sensor

node the time spent in different operation modes, including sensing, processing, listening, receiving, transmitting, and sleeping. We then use the electric current level of each mode from the specification of mica2 hardware platform [2] to calculate the energy consumption.

Table 2. Simulation parameters

Parameters	Values (Default*)
# nodes per WSN	10, 30*, 50
sample period (seconds)	8*, 16, 32, 64
window size (seconds)	40*, 80, 120

4.2. Simulation Results

4.2.1. Effect of the number of queries

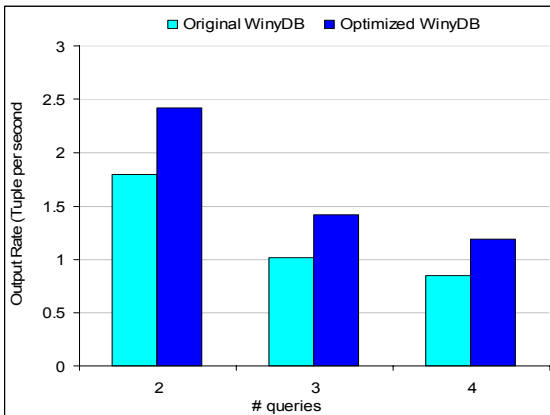


Figure 2. Output Rate

Figure 2 shows the effect of number of concurrent queries per PDA on the output rate. The main reason for the optimized WinyDB outperformed the original WinyDB is that, with query sharing, when a query can be answered by the results of some other queries, this query is not sent to the WSNs. Consequently, the workload of the WSNs is reduced. Additionally, the workload for a PDA to process the raw results from the WSNs directly is heavier than to process the intermediate results from other PDAs.

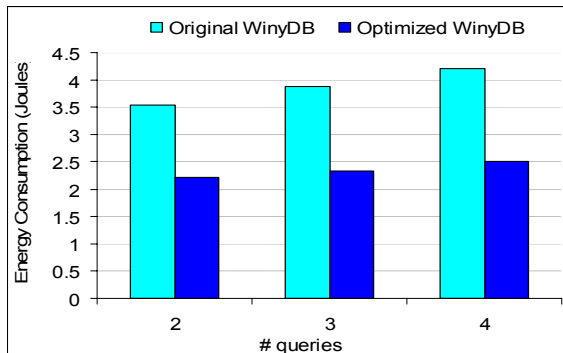


Figure 3. Energy Consumption

With the same configuration as for Figure 2, Figure 3 shows the average node energy consumption for one minute. The optimized WinyDB consumes much less energy than the original one.

The experimental results for other sizes of WSNs show that the optimized WinyDB can always improve the output rate by 30% to 40%, and the percentage of energy saving is about 40% on average.

4.2.2. Effect of the memory size




Figure 4. Output Rate

In Figure 4, we vary the memory size with respect to the sliding window size and evaluate the join query performance. The output rate of the original WinyDB drops sharply when the memory is not enough. This drop is because the original WinyDB discards some results when the memory is insufficient. In comparison, the optimized WinyDB tolerates memory insufficiency better by utilizing other PDAs to store some of the query results.

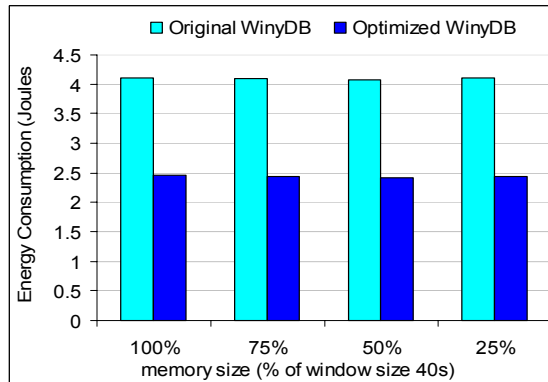


Figure 5. Energy Consumption

With the same configuration as for Figure 4, Figure 5 shows the average node energy consumption for one minute. The energy consumption is nearly the same for all memory sizes. The reason is that the number of queries running in the WSNs is independent of the memory size. Experimental results for other join window sizes are similar to those in Figures 4 and 5.

5. Related Work

There are embedded database systems for handheld devices, such as IBM DB2 Everyplace, Microsoft SQL Server 2005 Mobile Edition, Oracle Berkeley DB, and Sybase SQL Anywhere. These systems deal with disk-resident relational tables and usually have a backend database server to connect with. There have also been a few systems for sensor data collection from handheld

devices or microserver-grade sensor nodes, for example, CodeBlue [3], MSR Sense [11], and TASK [14]. However, they are either designed for special purpose (e.g., medical care applications in CodeBlue) or for digital signal processing and visualization. In comparison, WinyDB is designed as a general-purpose query processing system on PDAs for querying WSNs.

Existing query processing systems for sensor networks or data streams, such as Aurora [4], Cougar [18], NiagaraCQ [9], and TinyDB [17] assume queries are issued from a PC-grade central base station. Query parsing, optimization, and query result processing are done at this base station. In contrast, WinyDB runs on PDAs to query WSNs and multiple PDAs running WinyDB can answer queries collaboratively. As a result, the query optimization in WinyDB considers the resource constraints of both the WSNs and the PDAs and responds to the current status of PDAs dynamically.

Many techniques [1][10][12][16] are proposed to allow resource sharing between similar queries. Arasu et al. [1] proposed a suite of algorithms to exploit the resource sharing opportunities in sliding-window aggregations. Goldstein et al. [10] presented a scalable algorithm to determine whether part or all of a complex query can be computed from materialized views. Fjords [16] is a data structure for query plans such that related queries can be combined into a single Fjord. Niki Trigoni et al. [12] optimize multiple aggregation queries in the WSN to minimize the communication cost. In comparison, WinyDB focuses on sharing sensor data streams as well as query results among multiple PDAs.

6. Conclusions and Future Work

In this paper, we have presented our embedded query processing system, WinyDB, for handheld devices to query sensor networks. In particular, we have focused on the collaborative processing techniques for selection, projection, aggregation, and window join queries. Our experiments with WinyDB running on real PDAs querying simulated WSNs have shown that, with collaborative query processing, the output rate of the queries and the energy consumption of the sensor nodes can be improved significantly. Furthermore, this collaboration is especially beneficial to continuous sliding window join queries when the memory size of a handheld device is limited.

Developing a full-fledged, distributed, and efficient query processor for PDAs to query sensor networks involves numerous interesting research issues. Our work on WinyDB has only been initial steps on the collaborative query processing techniques. Future work along this line includes incorporation of more cost factors, such as processing and network communication workloads of the data providers, in selecting the collaboration schemes; extensions to WSNs with multiple sinks, count-based sliding window join queries, and other collaboration schemes.

7. Acknowledgements

This work was supported by grants HKUST6263/04E from the Hong Kong Research

Grants Council and MCCL03/04.EG01 from Microsoft Research Asia.

8. References

- [1] Arvind Arasu, Jennifer Widom, "Resource Sharing in Continuous Sliding-Window Aggregates," VLDB 2004.
- [2] Crossbow Corp. <http://www.xbow.com/>
- [3] David Malan, et al. "CodeBlue: An Ad Hoc Sensor Network Infrastructure for Emergency Medical Care," WAMES 2004.
- [4] D. Abadi, et al. "Aurora: A New Model and Architecture for Data Stream Management," VLDB Journal, 2003.
- [5] Habitat Monitoring on Great Duck Island. <http://www.greatduckisland.net/>
- [6] Ian F. Akyildiz, et al. "A Survey on Sensor Networks," IEEE Communications Magazine, vol. 40, no. 8, pp. 102-114, August 2002.
- [7] Intel Lab Data. <http://berkeley.intel-research.net/labdata/>
- [8] Jeffrey D. Ullman, "Principles of Database and Knowledge-Base Systems," Volume I. Computer Science Press 1988.
- [9] J. Chen, et al. "NiagaraCQ: A Scalable Continuous Query System for Internet Databases," SIGMOD 2000.
- [10] Jonathan Goldstein, Per-Åke Larson, "Optimizing queries using materialized views: a practical, scalable solution," SIGMOD 2001.
- [11] MSR Networked Embedded Sensing Toolkit (MSR Sense), <http://research.microsoft.com/nec/msrsense/>
- [12] Niki Trigoni, et al. "Multi-query Optimization for Sensor Networks," DCOSS 2005
- [13] Philip Levis, et al. "TOSSIM: Accurate and Scalable Simulation of Entire TinyOS Applications," SenSys 2003.
- [14] Phil Buonadonna, et al. "TASK: Sensor Network in a Box," EWSN 2005.
- [15] R. Motwani, et al. "Query Processing, Resource Management, and Approximation in a Data Stream Management System," CIDR 2003.
- [16] Samuel Madden, Michael J. Franklin, "Fjording the Stream: An Architecture for Queries over Streaming Sensor Data," ICDE 2002.
- [17] Samuel Madden, et al. "TinyDB: An Acquisitional Query Processing System for Sensor Networks," ACM TODS 2005.
- [18] Y. Yao, J. E. Gehrke, "Query Processing for sensor networks," CIDR 2003.