

EEIA: Energy Efficient Indexed Aggregation in Smart Wireless Sensor Networks

Mohamed Watfa, William Daher and Hisham Al Azar

Abstract—The main idea behind in network aggregation is that, rather than sending individual data items from sensors to sinks, multiple data items are aggregated as they are forwarded by the sensor network. Existing sensor network data aggregation techniques assume that the nodes are preprogrammed and send data to a central sink for offline querying and analysis. This approach faces two major drawbacks. First, the system behavior is preprogrammed and cannot be modified on the fly. Second, the increased energy wastage due to the communication overhead will result in decreasing the overall system lifetime. Thus, energy conservation is of prime consideration in sensor network protocols in order to maximize the network's operational lifetime. In this paper, we give an energy efficient approach to query processing by implementing new optimization techniques applied to in-network aggregation. We first discuss earlier approaches in sensors data management and highlight their disadvantages. We then present our approach "Energy Efficient Indexed Aggregation" (EEIA) and evaluate it through several simulations to prove its efficiency, competence and effectiveness.

Keywords—Sensor Networks, Data Base, Data Fusion, Aggregation, Indexing, Energy Efficiency

I. INTRODUCTION

SENSOR networks have become an important source of data acquisition with numerous applications being developed in monitoring various real-life phenomena. One example of monitoring applications includes organizing vehicle traffic in a large city [1]. Unfortunately, sensor data is subject to several sources of errors resulting from power limitations, wireless communication, latency, throughput, and various environmental effects. Such errors may seriously impact the answer to any query posed in the network. Current production of motes are roughly $2\text{cm} \times 4\text{cm} \times 1\text{cm}$ and are equipped with a radio, a processor, memory, a small package of AA batteries, and a collection of sensors. Additionally, new nodes are not passive devices [2], they are capable of sharing, computing and combining sensor readings; therefore, they are becoming tiny computers with different functionalities. Smart-sensor devices have been developed to an extent that it is now feasible to deploy large, distributed networks of such nodes [3,

4, 5, and 6] and extracting the data from the network is an essential step for the applications to work.

In-network aggregation is a well known technique to achieve energy efficiency when propagating data from information sources (sensor nodes) to multiple sinks. The main idea behind in-network aggregation is that, rather than sending individual data items from sensors to sinks, multiple data items are aggregated as they are forwarded by the sensor network. Data aggregation is application dependent, i.e., depending on the target application, the appropriate data aggregation operator (or *aggregator*) will be employed. From the information sink's point of view, the benefits of in-network aggregation are that in general it yields more manageable data streams avoiding overwhelming sources with massive amounts of information, and performs some filtering and preprocessing on the data, making the task of further processing the data less time and resource consuming. Because of its well-known power efficiency properties, in-network aggregation has been the focus of several recent research efforts on sensor networks. As a result, a number of data aggregation algorithms and data base systems targeting different sensor network scenarios have been proposed [2, 8, 9, 10, 11, 12, 13, and 14].

In contrast, our approach is a general one unrelated to any application but includes in-network aggregation because of the central usefulness of aggregations in sensor networks. Unless all the problems introduced by the resource constraints of nodes are handled, the future of handling data in sensor networks will not be rewarding and these restrictions are:

1. The system is highly volatile
2. Relational data tables are not static since new data is continuously being sensed
3. Energy costs in communication is high so in-network querying is desirable
4. We need to continuously monitor the rates and availability of data leading to long delays.
5. We need to maintain high-level of statistical summaries since we have limited storage on single nodes

All the current and future researches and innovations are taking into consideration these drawbacks because they are very crucial limitations on the sensor network's overall operational lifetime.

Although all researches and approaches [2, 7, 8, 9, 10, 11, 15, and 16] that were done earlier in this field provide some

M. K. Watfa is with the Department of Computer Science at the American University of Beirut, Beirut Lebanon. Phone#: 961-70-159757; fax: 961-1-744461; e-mail: mw11@aub.edu.lb.

W. Daher is with the Department of Computer Science at the American University of Beirut, Beirut Lebanon. Phone#: 961-70-159757; fax: 961-1-744461; e-mail: wsd03@aub.edu.lb.

H. Al Azar is with the Department of Computer Science at the American University of Beirut, Beirut Lebanon. Phone#: 961-70-159757; fax: 961-1-744461; e-mail: hga08@aub.edu.lb.

advantages over traditional centralized approaches, they still face some pitfalls and disadvantages. Studying the disadvantages of previous work, we concluded that we need to work on several major areas such as: power consumption, reliability, less aggregation overhead, less contention, fault-tolerance, and concurrency. In addition to that, the queries should also consider data integrity, security and efficiency. Our new distributed algorithms should take into consideration that nodes might have unequal battery levels opposed to unrealistic assumptions made in previous work. Taking into consideration these facts, our algorithm should evaluate the remaining power on each node and the varying power consumption that might differ from node to node to be able to maintain longer network lifetime to transmit useful data. We propose a new combinational improvement of all the available solutions taking into consideration the above constraints to develop the sensor node's ability to handle data locally in a very efficient way.

The rest of the paper is organized as follows: Section II will present a glance on previous approaches focusing on their disadvantages. Section III will present the EEIA approach. In Section IV, we present our simulation model. We will evaluate and simulate and analyze our proposed algorithms using our own simulator in Section V. We conclude this paper in Section VI with possible improvements.

II. RELATED WORK

Many researchers such as Yao et al., Bonnet et al., Gray et al., and Madden et al. [2, 7, 8, 9, 10, 11, 15, 16, and 17] tackled the data management topic in wireless sensor networks including query processing and data handling but none generic useful results and findings were originated for implementation. There has been a lot of work and approaches on query processing in distributed database systems [8 and 9], but most related work on distributed aggregation did not consider the physical limitations of sensor networks [18 and 19]. In addition, the TinyDB Project at Berkeley [15] conducted by Madden et al. also investigates query processing techniques for sensor networks including an implementation of the system on the Berkeley motes and aggregation queries. The basic approach used in both TinyDB [15] and TAG [16] is to compute a partial state record (*partial aggregation value*) at each intermediate node in the routing topology. During the epoch after query propagation, each mote listens for messages from its children during the interval it specified when forwarding the query. It then computes a partial state record consisting of the combination of any child values it heard with its own local sensor readings. Finally, during the transmission interval requested by its parent, the mote transmits this partial state record up the network. Figure 1 illustrates the process.

Previous studies [12-15] have shown that aggregation dramatically reduces the amount of data routed through the network, increasing throughput and extending the lifetime of battery powered sensor networks as less load is placed on power-hungry radios. Previous simulation studies have shown

that aggregation can reduce energy consumption by a factor of 5 in a large network (150-250 nodes) with five active sources and five sinks. In one experiment, we found that aggregation reduces traffic by up to 42% and nested queries reduce loss rates by 15-30% as shown in Figure 2.




Fig. 1 Partial State Records up the tree during an epoch




Fig. 2 Impact of duplicate suppression on energy savings

Previous networking research [12, 13, 14, 21] approached aggregation as an application specific technique that can be used to reduce the amount of data that must be sent over a network. In a previously proposed data dissemination scheme (directed diffusion with opportunistic aggregation), data is opportunistically aggregated at intermediate nodes on a low latency tree. In [13], the authors explore and evaluate greedy aggregation, an approach that adjusts aggregation points to increase the amount of path sharing. The greedy aggregation approach was implemented. Greedy aggregation differs from opportunistic aggregation in path establishment and maintenance. To construct a greedy incremental tree, a shortest path is established for only the first source to the sink whereas each of the other sources is incrementally connected

at the closest point on the existing tree. Figure 3 indicates that 30% on average reduction of energy savings resulted from the greedy aggregation.

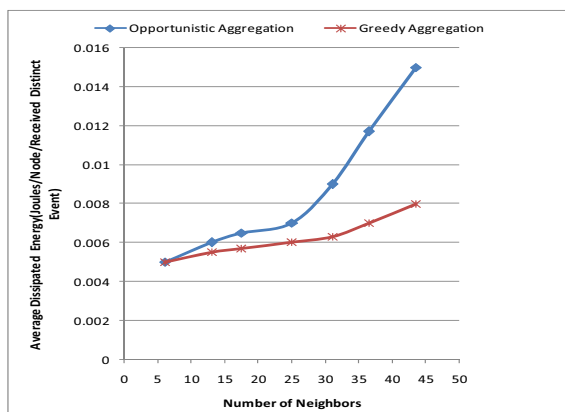


Fig. 3 Energy savings in the Greedy Aggregation vs. Energy savings in the Opportunistic Aggregation

In [2], they didn't explore all their techniques relative to mobility, and multiple queries. Thus, we can't be sure if their techniques are more efficient and reliable than the old techniques. In addition, they mentioned that in some cases that in-network aggregation performs worst than even the simplest approach "the naïve approach" [2, 8].

In [16], the authors presented TAG as a generic aggregation service for ad hoc networks of TinyOS motes. Thus, it provides a simple, declarative interface for data collection and aggregation. In addition, it intelligently distributes and executes aggregation queries in the sensor network in a time and power-efficient way, and is sensitive to the resource constraints and lossy communication properties of wireless sensor networks. TAG processes aggregates in the network by computing partial aggregation values over the flow of data from the nodes, discarding irrelevant data and combining relevant readings into more compact records when possible. Thus, the TAG paper contributes to the data management field in ad-hoc sensor networks in four aspects: simplicity, reduction in communication overhead, data messages overhead and finally reducing loss in the sensor network. In Figure 4, we present the TAG vs Centralized "naïve" approach to emphasize on its efficiency in the number of bytes transmitted.

Most of the conclusions that the above researchers are credited for can be described as: "We described a vision of processing queries over sensor networks" [11]. Some presented a prototype or some techniques they used without any actual implementation and simulation results. For example, the Cornell COUGAR system prototype [7, 8, 9, 10, and 11] is a first effort towards sensor database system. Thus, a lot of improvements are still needed in this field to achieve better generic approaches for implementation in wireless sensor networks plus taking into consideration all the drawbacks and pitfalls of earlier techniques. Unlike other

networks, wireless sensor network still need an international standard ISO to be build upon and all these future researches are directed towards this goal. One part is related to finding a general applicable approach for data management in sensor networks which will become a self-aware, self-configuring and reliable system with respect to all nodes' resource constraints.

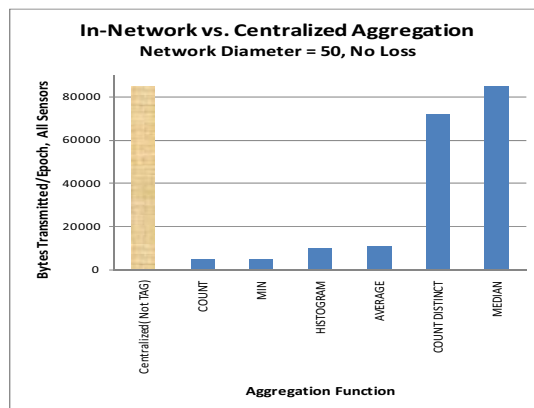


Fig. 4 In-network vs. Centralized Aggregates

As various groups around the country have begun to deploy large networks of sensors, a need has arisen for tools to collect and query data from these networks. Of particular interest are aggregates – operations which summarize current sensor values in some or all of a sensor network. For example, given a dense network of thousands of sensors querying temperature, users want to know temperature patterns in relatively large regions encompassing tens of sensors – individual sensor readings are of little value. Sensor networks are limited in external bandwidth, i.e. how much data they can deliver to an outside system. In many cases the externally available bandwidth is a small fraction of the aggregate internal bandwidth. Thus computing aggregates in-network is also attractive from a network performance and longevity standpoint: extracting all data over all time from all sensors will consume large amounts of time and power as each individual sensor's data is independently routed through the network. As noted before, aggregation dramatically reduces the amount of data routed through the network, increasing throughput and extending the life of battery powered sensor networks as less load is placed on power-hungry radios. Also, The fact that every message is effectively broadcast to all other sensors within range enables a number of optimizations that can significantly reduce the number of messages transmitted and increase the accuracy of aggregates in the face of transmission failures.

To conclude this section we present a table listing the drawbacks from all the approaches that we will try to overcome (Table I). In the next section, we provide an outline of our approach and all the necessary steps to implement our distributed in-network aggregation approach.

TABLE I
DRAW BACKS IN PREVIOUS APPROACHES

Paper	Disadvantages
“The Design of an Acquisitional Query Processor For Sensor Networks” (TinyDB) [10]	<ul style="list-style-type: none"> • Costs for disseminating, executing and forwarding false query results in terms of incorrect initiation can be high. • SRT are limited to constant attributes • SRT maintenance and construction costs are high. • Building SRT could be costly • No handling of multi-queries (one query runs at a time) • Radio communication has high link-level losses typically about 20% @ 5m • Central node failures cause lost of data from a sub-tree
“The Cougar Approach to In-Network Query Processing in Sensor Networks” [11]	<ul style="list-style-type: none"> • Unreliable data due to several facts such as noise • Node failure especially Leader node failure • Multi-hop routing consumes energy over large networks • Time Synchronization between sensor nodes • Uncertainty decisions by the optimizer • Changes in network topology and power level at nodes affect the Cougar approach • Storage and update of metadata in each node is crucial • No handling of multi-queries • No experimental results are mentioned in the paper • Constantly changing conditions • Meta-data management • Fault tolerance
“Supporting aggregate queries over ad-hoc sensor networks” [2]	<ul style="list-style-type: none"> • Assumption of TinyOS as a sending-receiving API • Symmetry in Radio links (not valid) • Routing tree maintenance and flooding requests consume a great deal of energy • Generalized aggregation predicate as median and mode are not considered • Time delay depends on the depth of the routing tree • Long computation time • Pipelined approach: an extra number of message for 1st aggregate → extra overhead • Failure independence assumption (messages to both parents can't be lost) • Sensor values are independent and randomly distributed • User acceptance of error bounds • Buffer overhead • Effect of sensor mobility

III. THE EEIA APPROACH

Our approach consists of providing a new distributed algorithm for query processing in wireless sensor networks which is an optimized energy efficient distributed algorithm with respect to all the sensor's resource constraints. Some similarities to recent approaches are also used such as upgrading the tinyDB [10, 15] approach, an ACQP engine that is a distributed query processor which runs on each of the nodes in a sensor network, and the TAG approach [16].

Our goal is to provide significant reductions in power consumption through reducing the number of query related messages in the whole network. Low energy consumption, and limited storage and memory usage are the three main constraints which we focus on in our approach. This section will provide a detailed explanation of our approach by presenting the problem and the corresponding solution. We evaluate the approach in the next section through simulation.

Sensor networks have very limited power, small memory computational power and limited bandwidth so some possible unanswered questions related to in network data aggregation schemes in sensor networks are:

- *How can we decrease power consumption in a data management algorithm?*
- *How can we decrease the number of collisions and thus reduce the overall end-end latency?*
- *How can we decrease the number of computations at each node?*
- *How could we let our algorithm be self-adaptive to the changing network conditions?*

Our distributed algorithm offers three new ideas in order to answer the previous questions. Our first aim is to decrease the packet size and the second is to decrease the number of packets sent. To decrease the packet size each sensor should have values in its buffer of all its children nodes to perform

partial aggregation before sending this value to its parent. In large sensor networks, aggregation of data having small packets and small values decreases the power consumption and the computation overhead. Our second approach is to index the network so to be able to query data with minimum number of exchanged packets. We will start by building an index tree (IST) that is similar to the SRT of TinyDB [10] but for not only fixed attributes but also variable ones. The problem with building such a tree is the maintenance overhead, but we will prove that our algorithm maintains the tree with little or no maintenance overhead.

Our algorithm is to first build an index routing tree (*Section A*). Secondly, each child in the tree sends its index to its parent. Since the parent knows the number of children it has, it compares the indexes received from each child, if they are equal, the parent indexes itself with their index and sends the index to its parent otherwise it does nothing. Upon a change in the index of one node, this node sends the new index to its parent, the parent checks again to see if the indexes are equal; if not and this parent is indexed, it removes its index and informs its parent, but if this parent is not indexed, it doesn't have to inform its parent. With our network, indexing a query could take less time and computation power to return the result. For example if we have a query that asks for the average temperature where the temperature is above 36. When this query reaches a node with index 1, the node doesn't forward the query to its children. Our third idea is to conserve energy as much as possible using indexing with the power evaluation criteria available in TinyDB at each node. We can use an index of 0 to note that a certain node is low in power try eliminate it in the execution flow of the query. In the next sections, we present our algorithm in details.

A. Building the Routing Tree

After the nodes are randomly deployed, an index routing tree is built. The routing tree is built as follows. The closest node to the base station is chosen to be the root of the tree ($level = 0$). Once chosen, the root broadcasts requests containing its level to all its one hop neighbors (within its transmission range). When receiving the request, a neighbor node assigns itself a $level = level + 1$ and chooses its parent to be the level up node from which it received the request, then re-broadcasts new requests containing its new level to all its neighbors and so on until no neighbors are found; thus the last nodes become leaf nodes. Whenever a node receives two requests from two different nodes, if it has a level, it discards the second request; and if not, it selects the first arrived request. Thus it chooses one parent and one new level ($level = level + 1$). In our tree algorithm, we intend to let every node have only a single parent. Figure 5 displays the flowchart of building the tree. After building the tree each node sends its reading value to its parent starting from the leaf and up. Every node stores its last sent value. Every parent node receiving values from multiple children calculates the average of the values received and sends it to its parent and so on until the value reaches the base station (Figure 6).

As an example, we are going to implement an application that calculates the average temperature of all the nodes which are part of the tree. The number of nodes is relative to the application under study (*we incorporated 11 nodes in our example*)




Fig. 5 Flowchart of building the tree

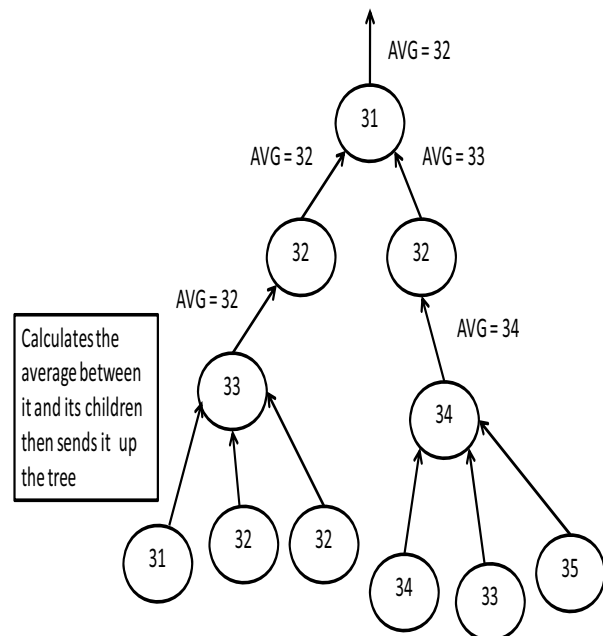


Fig. 6 Initial deployment

B. Building the Index Table in the Routing Tree

When the base station receives the values, it sends a packet containing the index table to all the nodes. The first time the index table is sent, the value ranges of each index will be large; the reason behind this approach is not to send large packets in the network. If the index table is large, it may lead to collisions. When a node receives the index table, it compares its readings to the index table and indexes itself accordingly (Figure 7.1).

In the second round the index table changes as value for index ranges becomes smaller. After couple of rounds the index values will be more accurate (Figure 7.2). The number of rounds depends on the size of the index table decided once the network is deployed. The final index table will be derived on each node. Deciding on the index ranges of the system depends on the type of sensor node; sensor nodes with readings that vary in large ranges should have index ranges with large values (Figure 7.3).




Fig. 7.1 Round 1 for assigning indexes

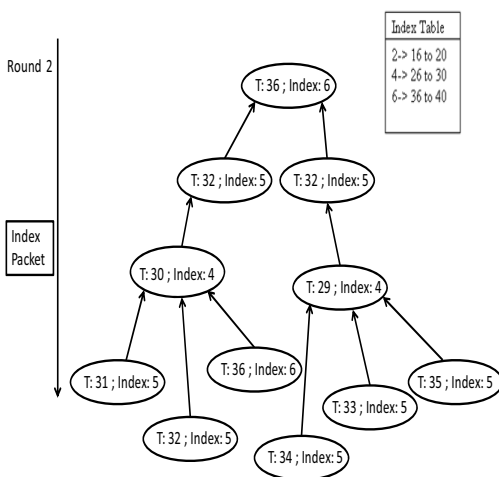


Fig. 7.2 Round 2 for assigning indexes

Index Table	
1-	10 to 15
2->	16 to 20
3->	21 to 25
4->	26 to 30
5->	31 to 35
6->	36 to 40

Fig. 7.3 Final index table

The child sends its index to its parent. When a parent receives an index from its children, the parent compares all its indexes with its own, if they are all similar, the parent indexes itself as such. However, if all indexes are not the same, the node examines the percentage of the similarity, if the similarity is larger or equal to 75% (based on the simulation results in Section V), it indexes itself with the dominant index and ignores the others. If the index similarities are lower than 75% then the parent indexes itself as between the smallest and largest index (Figure 7.4). After all the nodes are indexed in the network, the parents and children agree on a common value.

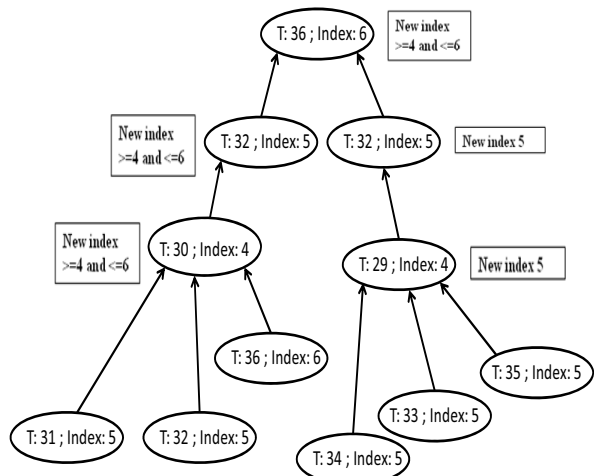


Fig. 7.4 Assigning parent indexes

C. Common Value Agreement (CV)

After a parent receives values from its children, it first calculates the average of the values; it stores the calculated average and sends it back to its children. We call this value "common value" and we denoted it by *cv* (Figure 8). Each node stores two values: the *cv* of its parent and the *cv* of its children (except leaf nodes). When a node needs to send a new reading to its parent, it subtracts the *cv* from its reading and forwards the value to its parent. The *cv* will be updated in case there is a major change in the average of the children.

When a parent notices a large change between its children and the *cv*, the parent resends the new average to its children as the new *cv*. With our *cv* approach, sent packets are smaller

and therefore leading to less collisions, more energy efficiency and less calculation overhead.




Fig. 8 Agreement on the value

D. Aggregate Functions Evaluation

The calculation approach defers between different aggregate functions. In our algorithm, we evaluate the 5 basic aggregate functions Sum, Average, Count, Min, Max.

Average Function

We start with the Average function. To elaborate more on this function let us examine the following query:

```
SELECT Avg(temp)
FROM Sensors
```

Since this query asks for the average temperature of the whole network, the query should reach all the nodes where values will be extracted. In our algorithm this is not the case; our algorithm offers the user two approaches to calculate the average. In the first approach, when the query reaches the root node, the root node doesn't forward the query to its children but returns his *cv* since the *cv* is the common average between it and its children. In the second approach the query reaches all the nodes but not all the nodes return a value. When a query reaches a node, the node examines its current reading and index. If his current reading still lies within the same index the node doesn't forward any value since his value will not have a noticeable change to the final result. If the current reading doesn't lie in the same index the node changes its index, and sends the *cv* subtracted from his reading to the parent node. After receiving the new reading the parent notices a large value from his child thus updates his index status and *cv* if needed according to the previously discussed approaches. Then the parent node calculates the new average. Assume *Avg* is the old average value, *Avg_{new}* the new average, *nv* the new value received from the child and *p* the children count involved in the query. The parent calculates the new average using the following formula:

$$Avg_{new} = \frac{(Avg \times p) + nv}{p} \quad (1)$$

In the second approach, sending the value depending on the index change decreases the overhead of sending packets where the change in reading will not cause a notable change to the overall value; thus, using this approach results in sending a small number of packets. Deciding on what approach to use depends on how accurate the data needs to be.

Count, Sum, Min and Max Functions

The Count function is evaluated in a normal approach where the node, if meeting the criteria, sends 1 to his parent where the parent adds the count of his children and forwards them to his parent and so on. The Sum function can also be evaluated using two approaches. The first approach is the usual one where values are sent to the parent node that in his turn sums them and sends them to his parent and so on. The second approach of evaluating sum is to break the Sum query into two queries, an average query and a count query. In this approach the advantages of average evaluation discussed previously can be used. After a node receives a sum function it sends its reading as if it is calculating the average and then sends the count. The base station calculates the Sum as Average \times Count.

Deciding on the approach to use depends on the query and the exactness of the result. Our engine on the base station decides what approach to use. The Min and Max function are evaluated in similar approach to the average where the node sends the *cv* subtracted from his value. The parent node in its turn chooses the largest or smallest (Min or Max) value received, adds to it the children *cv* then subtracts from it its parent *cv* and sends it to its parent.

E. Queries with conditions

For other types of queries that have a condition, our approach should increase the throughput of the query since indexing will help in the injection of the query. Let us elaborate more by examining the following query:

```
SELECT Avg(temp)
FROM Sensors
WHERE temp > 35
```

Our engine on the base station will parse this query and translate the condition into index. The condition "Where *temp* > 35" will be translated into "Where *tempIndex* > 5" assuming index 5 and its preceding indexes are between 0 and 30. After this translation the query is injected into the network. From the root and on, every parent node checks if it has an index smaller or equal to 5, if yes, it will not forward the query to its children. Thus the query is filtered through the injection state. As shown in Figure 9 below, the root broadcasts the query to its children. Once arrived to each child, they check if they have an index smaller or equal to 5. Thus, for node having the index 5, it ignores the query request, but in the case of the other node, it re-broadcasts the request to

all of its children (index > 5) which in their turn, each of them forwards the request if its index is greater than 5.

This approach removes the overhead of sending the query to unneeded nodes. This approach increases the energy efficiency of the network where nodes that do not satisfy the condition will not need to spend energy since filtering is happening in the injection of the query rather than the base station.

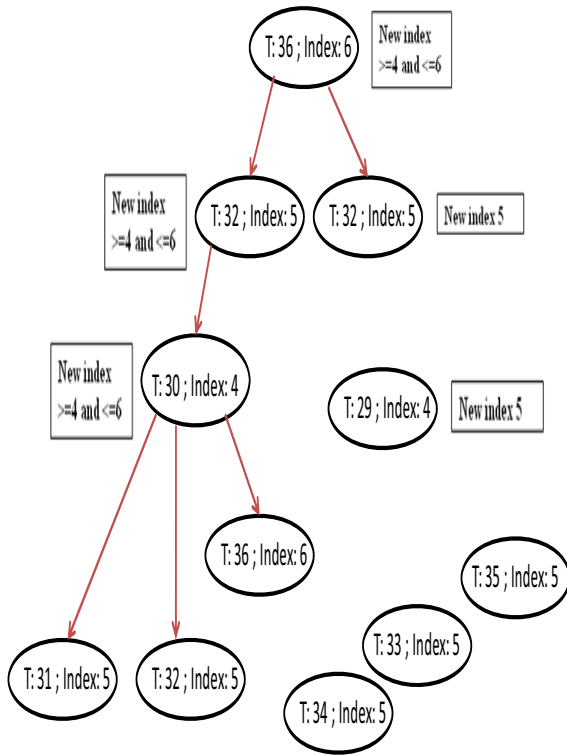


Fig. 9 Handling the “Where” condition

F. Power Management

Power consumption and network life time are major issues in the wireless sensor network design. In our algorithm, we try to increase the lifetime of the network through two different approaches.

In the first approach, the node keeps track of the number of messages sent and number of messages received. From these numbers the node can approximate the energy consumed and therefore the amount of energy left. We can also incorporate a battery model in our algorithm. When the node reaches a state where its energy is close to a predefined threshold, it informs its parent. The parent will therefore decrease the number of packets sent to this child thus will send packets to this child every two rounds rather than every round hence decreasing the transfer data rate to this child. If the child reaches a very low state of energy, it informs its parent where the parent stops sending any packets to this child.

The second approach is achieved by changing the root parent every two rounds. If a child can have multiple parents, the child, after couple of rounds, changes to another parent if the other parent has more energy than his current parent. When a parent is low on energy it informs its children, the children in their turn will ask another parent node if they could join it. If another parent with more energy is found, the child switches to this new parent. A child could know if the new parent has more energy than his old one by knowing the amount of energy consumed by both parents derived from the formula discussed before. We will just incorporate the first approach in our simulator leaving the second approach for future work. The first approach should reduce the energy consumption of every single node in the network thus increasing the network lifetime.

G. Query Optimization

In this section, we will discuss our approach in optimizing a query. The base station keeps record of the last queries with a time stamp. Upon issuing a new query, the query optimizer checks for similar queries issued before and their results. If the results are close, it concludes that the network readings are not changing so instead of sending the query to the whole network, the query is sent to different parts of the network. To send the query to different parts of the network, the query optimizer sends the query to the level 1 parent nodes which in turn will choose some of their children and forward them the query. The number of children nodes chosen depends on the query optimizer decision. With this approach a smaller amount of nodes participate in the query (Figure 10).

This approach increases energy efficiency and throughput of the network but gives an approximation of the result. This optimization technique doesn't apply on all kinds of queries.

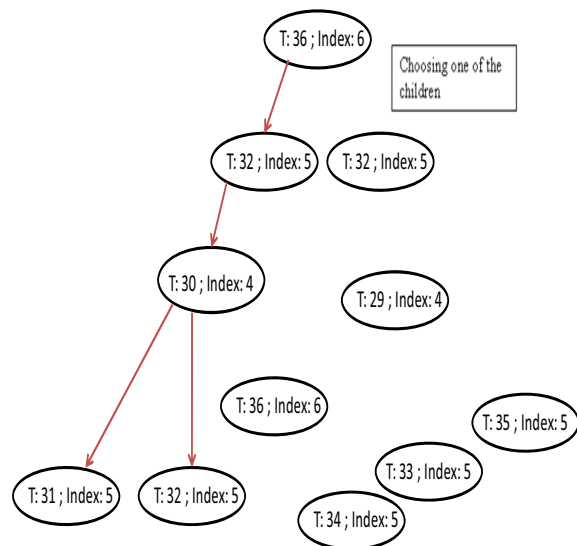


Fig. 10 Decrease in the number of nodes participating

IV. SIMULATION MODEL

To test the efficiency of our algorithm, we decided to model our own simulator to achieve our goals because of the lack of database simulators. We will present the detailed information on how our simulation was built.

A. Our Simulator Model

Our simulator is written in VB and it incorporates our algorithm to be tested with the naive and the simple TinyDB approaches. We will evaluate our algorithm by comparing it to these approaches in terms of energy consumption, network lifetime and time delays. We randomly deploy a large number of nodes, then a routing tree is built in which the query is sent from the root to the leaf nodes to be evaluated, processed. The leaf nodes will send the results back to their parents where they are aggregated and sent over to the parent's parents until an aggregated value reaches the root which in its turn, sends the aggregated value back to the base station.

In the tree, each node is randomly colored to present its level (*number of hops away from the base station*). An edge connects two neighbor nodes if there are in the communication range of each other i.e. they can communicate by sending and receiving messages. Every node contains a cache in which it saves its level number, its index, its parent id and its children's ids. Figure 11.1 portrays the building of the routing tree. Figure 11.2 shows our simulator at work (building indexes in the tree). Every parent node is colored by the color of its children if the percentage of itself and its children's color is greater than 75%. Each node is colored by its index value.

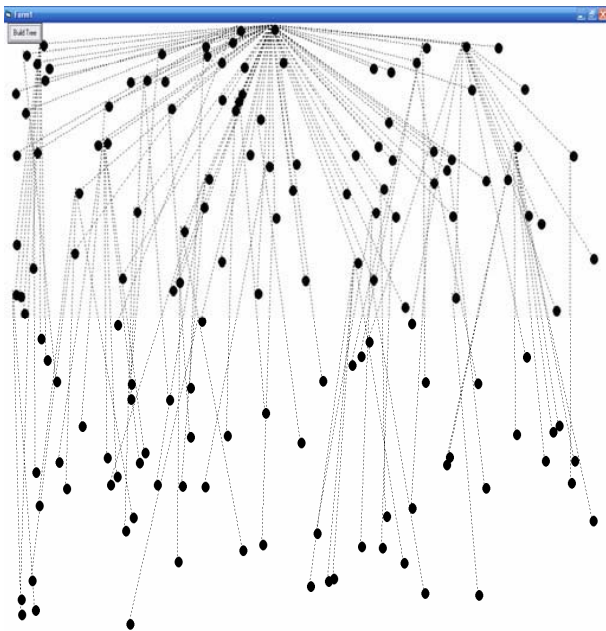


Fig. 11.1 Building the routing tree simulation

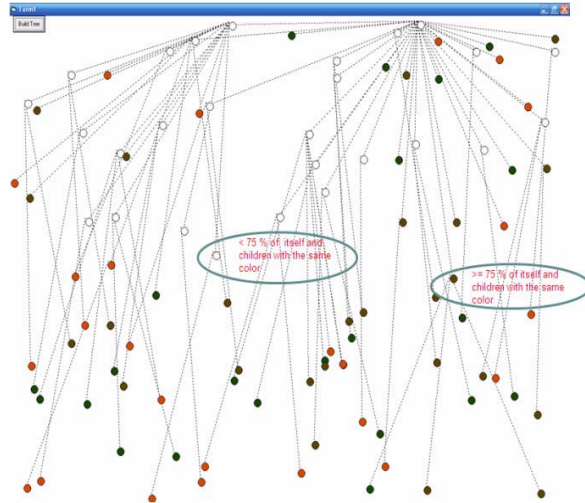


Fig. 11.2 Building indexes in the routing tree

B. Evaluation Metrics

Our Algorithm will be evaluated based on two basic metrics: power consumption and network lifetime. We are going to compare the power consumed in retrieving queries using our algorithm with other known algorithm (*the naive and the simple TinyDB approaches*). The performance of the algorithm over time will also be studied to determine the benefits of using in-network aggregation. This is done by assuming that each sensor node has a limited energy supply of and is deactivated when the available energy is used up. The performance is evaluated in terms of *network lifetime*. The *network lifetime* is the continuous operational time of the system before the coverage drops below a specified threshold (for example 0.8). In addition to that, we will incorporate the energy wastage resulting from building the IRT (index routing tree).

C. Investigated Parameter Space

In our simulation, we are going to randomly spread about 500 sensor nodes in a 10000 x 10000 region in VB to investigate the change in temperature and humidity. We are going to query these nodes to get Max, Min, and Average values. Further investigation about some sensor nodes measurements (return numbers) are left for future work.

D. Scenarios

The node we are going to use are the Berkley Mica notes [28, 3, 5] with TinyOS [10, 25]. S-MAC is used as the default MAC layer protocol. The sensor nodes will be used to measure temperature, humidity and other metrics as provided by the sensor itself (*The Mica notes already support temperature sensors, light sensors, magnetometers, accelerometers, and microphones*). These nodes will be

connected to a base station that will issue the query and the results are returned to this station by in-network aggregation in the intermediate nodes, thus, providing a reduction in message overhead and in energy consumption.

In addition, several optimization techniques should be applied to prevent other problems such as nodes mobility, node failures, network changes and others, from affecting the network.

E. Expected Results

We have to investigate the efficiency of our approach in several areas through some expected results from the simulation such as:

1. Energy efficiency in terms of low power consumption while transferring the query related data through the network (*reduce the number of messages required to compute the aggregations*).
2. Overhead calculation in the network (*decrease the overhead*).
3. Result accuracy from the data aggregated (*is the data used for aggregation satisfactory and reliable?*).
4. Other factors should be taken to consideration such as time delays, communication failures and sensor mobility.

V. SIMULATION RESULTS AND ANALYSIS

In this section, we are going to evaluate our algorithm in terms of power consumption, network lifetime and finally time delay. Then, we are going to compare our results to well known data management approaches such as the naive and the simple TinyDB approaches. We are going to prove the efficiency of our distributed approach and its adaptive manner in maintenance and re-construction of the index routing tree due to unexpected failures. We compare our EEIA approach with the:

1. *Naïve algorithm*: all the query results from each node are sent to its parent until the results reach the root where they are aggregated and sent to the base station.
2. *Simple TinyDB*: the results are aggregated at each intermediate node until reaching the root.

In our simulation, we assumed the energy wasted is $1\mu\text{J}$ for sending a single bit and $0.5\mu\text{J}$ for receiving a single bit. Initially, each node has 1 J of available energy.

In our simulation, we also incorporated the energy consumption of building an index tree. As you can see in Figure 12, as the number of nodes increases, the energy consumption increases linearly since all nodes participate in building the query with same amount of energy. We concluded that the energy consumption of building the index tree is equivalent to initiating one query in the network.

As for the maintenance of the index tree, we see from Figure 13 that the average energy consumed in the maintenance depends on the readings of the nodes in the network. If the network readings change significantly in a small amount of time, the energy consumed in maintaining the tree increases. On the other hand, if the network readings change slowly then the energy consumption of the index maintenance decreases. The energy consumption in the first 15 seconds is high since the index tree was being built.

We issued a number of different queries on 500 nodes and compared the energy consumption, delay and number of instructions using our approach compared to the normal approach of broadcasting the query to the network. To simulate the same queries, we implemented in our simulator two approaches, the first approach queries the network by broadcasting the query to all the nodes and aggregating the results back to the base station. As for the second approach, we added our index querying approach. To make our simulation more realistic, we maintained the same condition on the network while using the different querying approaches.

In Figure 14, we compared the energy consumed for the same queries using the 2 approaches. We issued 12 different queries and calculated the energy consumed by these queries to return the result. As can be seen from the graph, all queries using the indexed approach consume less energy than using the other approach. You can also note from the graph that some queries have energy consumption that is close to the normal approach while others have larger energy consumption. This difference depends on the conditions of the issued query since more selective queries tend to have a larger advantage using our approach (*indexing with more selectiveness decreases the number of messages sent and received in the network*).

In Figure 15, we compare the number of packets sent for the same query using the 2 approaches. As the figure shows, the number of packets sent decreases with the indexing approach since some nodes will not forward the packets to their children if their children don't satisfy the query conditions.

As can be seen in Figure 16, the delay is decreased using our approach since as discussed before the number of packets sent decrease hence collisions decrease. In the delay simulation, we assumed a packets needs 0.01 sec to be resent.

In the last simulation, we added the *cv* (*common value*) approach to our simulator and compared the lifetime of the network using indexing and *cv* querying to the normal approach. As can be seen from Figure 17, the first node dies after 83 queries in the normal approach while using our approach; the first node dies after 130 queries. This increase in the lifetime of the network is due to two factors. The first factor is that nodes are sending fewer packets using the indexed approach thus less energy consumption per node. The second factor is the decrease in packet size with *cv* approach where less energy is consumed in sending the packet.

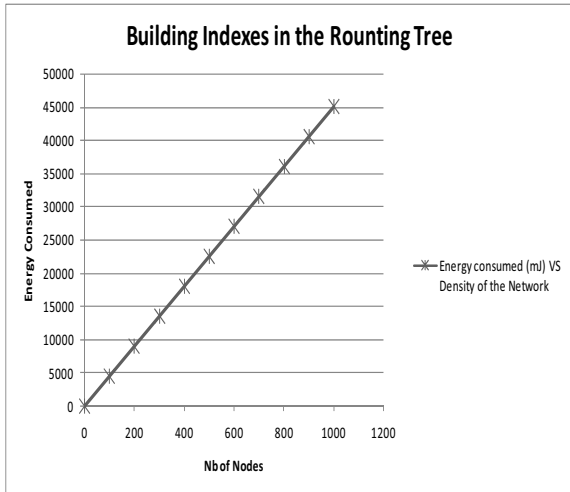


Fig. 12 Building the routing tree

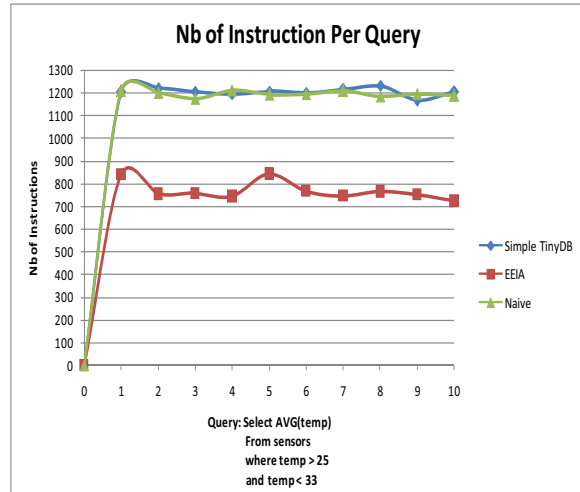


Fig. 15 Nb of instructions per query




Fig. 13 Maintaining the routing tree

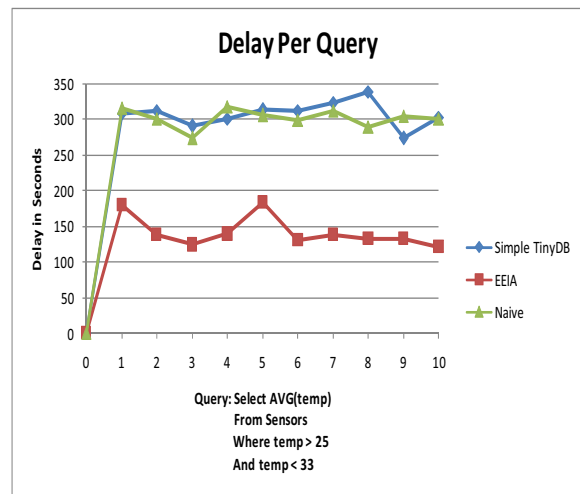


Fig. 16 Time delay per query




Fig. 14 Energy consumption per query




Fig. 17 Network lifetime

VI. CONCLUSIONS

In summary, we have showed how aggregate queries are efficiently executed over wireless sensor networks in a distributed manner. We have proved that our in-network distributed approach performed better in terms of energy reduction and network lifetime than the naïve and simple TinyDB approaches. Furthermore as for future work, our approach should confront with the difficulties of topology construction, data routing, loss tolerance by including several optimization techniques that further decrease message costs and improve tolerance to failure and loss. In addition to implementing these techniques, we need to rethink some of these techniques to present more efficiency to network changes and external factors which could affect our approach such as node mobility, obstacles and other issues. In addition as future work, we could also extend our simulator to incorporate a 3D tree construction technique plus other methodologies mentioned above.

ACKNOWLEDGMENT

The authors would like to gratefully thank the support of the American University of Beirut university research board (URB) and the Lebanese National Council for Scientific Research (LNCSSR) through their corresponding research grants.

REFERENCES

- [1] D. Estrin, R. Govindan, J. Heidemann (Editors), "Embedding the Internet", In CACM Volume 43, Issue 5, May 2000, Pages: 38 - 41.
- [2] S. Madden, R. Szewczyk, M. Franklin, and D. Culler, "Supporting aggregate queries over ad-hoc sensor networks", In Workshop on Mobile Computing and Systems Applications (WMCSA), Callicoon, NY, June 2002, Pages: 49 - 58.
- [3] G. Pottie and W. Kaiser, "Wireless integrated network sensors", *Communications of the ACM*, Volume 43, Issue 5, May 2000, Pages: 51- 58.
- [4] Mainwaring, J. Polastre, R. Szewczyk, and D. Culler, "Wireless sensor networks for habitat monitoring", In Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications, Atlanta, Georgia, USA, 2002, Pages: 88 - 97.
- [5] J. Hill, R. Szewczyk, A.Woo, S. Hollar, and D. Pister, "System architecture directions for networked sensors", In Proceedings of the 1st international conference on Embedded networked sensor systems Los Angeles, California, USA, November 2000, Pages: 193 - 204.
- [6] Cerpa, J. Elson, D.Estrin, L. Girod, M. Hamilton, and J. Zhao, "Habitat monitoring: Application driver for wireless communications technology", In *ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean*, April 2001, Pages: 20- 41.
- [7] P. Bonnet, J. Gehrke, and P. Seshadri, "Towards sensor database systems", In *2nd International Conference on Mobile Data Management, Hong Kong*, January 2001, Pages: 3 - 14.
- [8] Y. Yao and J. Gehrke, "Query Processing in Sensor Networks", In Proceedings of the First Biennial Conference on Innovative Data Systems Research (CIDR 2003), Asilomar, California, January 2003.
- [9] Demers, J. Gehrke, R. Rajaraman, N. Trigoni, and Y. Yao, "Energy-Efficient Data Management for Sensor Networks: A Work-In-Progress Report", 2nd IEEE Upstate New York Workshop on Sensor Networks. Syracuse, NY, October 2003.
- [10] S. Madden, M. Franklin, J. Hellerstein, W. Hong, "The Design of an Acquisitional Query Processor for Sensor Networks", ACM SIGMOD Conference, San Diego, CA, June 2003, Pages: 491 - 502.
- [11] Y. Yao and J. Gehrke, "The Cougar Approach to In-Network Query Processing in Sensor Networks", *Sigmod Record*, Volume 31, Number 3, September 2002, Pages: 9 - 18.
- [12] J. Heidemann, F. Silva, C. Intanagonwivat, R. Govindan, D. Estrin, and D. Ganesan, "Building efficient wireless sensor networks with low-level naming", In Proceedings of the eighteenth ACM symposium on Operating systems principles, Banff, Alberta, Canada, October 2001, Pages: 146-159.
- [13] C. Intanagonwivat, D. Estrin, R. Govindan, and J. Heidemann, "Impact of network density on data aggregation in wireless sensor networks", In ICDCS-22, November 2001.
- [14] C. Intanagonwivat, R. Govindan, and D. Estrin, "Directed diffusion: A scalable and robust communication paradigm for sensor networks", In *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networks (MobiCOM 2000)*, Boston, MA, August 2000, Pages: 56-67.
- [15] S. Madden, M. Franklin, J. Hellerstein and W. Hong, "TinyDB: An Acquisitional Query Processing System for Sensor Networks", In *ACM Transactions on Database Systems (TODS)*, Volume 30, Issue 1, March 2005, Pages: 122 - 173
- [16] S. Madden, M. Franklin, J. Hellerstein, and W. Hong, "TAG: a Tiny Aggregation service for ad-hoc sensor networks," In *5th Annual Symposium on Operating Systems Design and Implementation (OSDI)*, December 2002, Pages: 131-146.
- [17] J. Hellerstein, P. Hass, and H.Wang, "Online aggregation", In *Proceedings of the ACM SIGMOD*, Tucson, AZ, May 1997, Pages: 171-182.
- [18] Shatdal and J. Naughton, "Adaptive parallel aggregation algorithms", In Proceedings of the 1995 ACM SIGMOD international conference on Management of data, San Jose, California, United States, Pages: 104 - 114.
- [19] W. Yan and P. Larson, "Eager aggregation and lazy aggregation", In Proceedings of the 21th International Conference on Very Large Data Bases (VLDB), September 1995, Pages: 345 - 357.
- [20] ANSI. *SQL Standard*, 1992. X3.135-1992.
- [21] P. Larson, "Data reduction by partial preaggregation", In *ICDE*, 2002.
- [22] S. Madden and M. Franklin, "Fjording the stream: An architecture for queries over streaming sensor data", In *ICDE*, 2002.
- [23] M. Garofalakis and P. Gibbons, "Approximate query processing: Taming the terabytes!" (Tutorial), In *VLDB*, 2001.
- [24] M. Stonebraker and G. Kemnitz, "The POSTGRES Next-Generation Database Management System", *Communications of the ACM*, Volume 34, Issue 10, 1991, Pages: 78-92.
- [25] UC Berkeley, "Smart buildings admit their faults", Web Page, November 2001. Lab Notes: Research from the College of Engineering, UC Berkeley. <http://coe.berkeley.edu/labnotes/1101.smartbuildings.html>.
- [26] O. Wolfson, A. Sistla, B. Xu, J. Zhou, and S. Chamberlain, "DOMINO: Databases fOr MovINg Objects tracking", In Proceedings of the 1999 ACM SIGMOD international conference on Management of data, Philadelphia, Pennsylvania, United States, June 1999, Pages: 547 - 549.
- [27] D. Kossman, "The state of the art in distributed query processing", In *ACM Computing Surveys (CSUR)*, Volume 32, Issue 4, December 2000, Pages: 422 - 469.
- [28] Crossbow, "Wireless sensor networks (mica motes)", <http://www.xbow.com/Products/WirelessSensorNetworks.htm>.
- [29] Hill. A software architecture to support network sensors. Master's thesis, UC Berkeley, 2000.
- [30] Mica motes: Crossbow technology, inc. Tech. rep. <http://www.xbow.com>.