

Multi-Criteria Routing in Pervasive Environment with Sensors *

Ganesh Santhanakrishnan Qinglan Li Jonathan Beaver
Panos K. Chrysanthis Ahmed Amer Alexandros Labrinidis
Department of Computer Science
University of Pittsburgh
Pittsburgh, PA 15260 U.S.A.

{ganesh, qinglan, beaver, panos, amer, labrinid}@cs.pitt.edu

Abstract

Interconnected computing nodes in pervasive systems demand efficient management to ensure longevity and effectiveness. This is particularly true when we consider wireless sensor networks, for which we propose a new scheme for adaptive route management. There have been numerous studies that have looked at the routing of data in sensor networks with the sole intention of reducing communication power. However there has been comparatively less prior art in the area of semantic and multi-criteria based routing. We look at routing in sensor networks from these perspectives and propose an adaptive multi-criteria routing protocol in the context of wireless sensor networks. Our experimental results show that our approach consistently outperforms the leading multi-criteria algorithm in its class that considers query semantics, in terms of Network Lifetime, Network Coverage and the Survivability of Critical Nodes.

1 Introduction

The computing environment today is changing quickly with the emerging of small sensor devices and sensor networks. Such sensor networks will be an integral part of a pervasive computing environment since they allow interaction with the physical environment. Consequently, sensor pervasive services would not be an exception with respect to quality of data, coverage and lifetime of the services.

A major challenge in these new environments is power conservation. More precisely, in sensor networks communication costs in power and energy subsume other costs

such as network processing. Hence many approaches toward *in-network processing* have been proposed. The main idea behind in-network processing is to perform computation in the network itself reducing the size of the data to be sent higher up to other nodes. This helps in reducing power consumption since computation is cheaper in terms of energy and power than communication. The chief among the approaches for *in-network processing* are the TAG approach [31] from the University of California at Berkley and the Cougar approach [16] from Cornell University. A more recent approach adopted successfully by researchers at the University of Pittsburgh, is TiNA: A Scheme for Temporal Coherency-Aware in-Network Aggregation [2, 26]. As we see more and more approaches adopting in-network processing of data, it is imperative that the creation of the routing tree itself be based on the semantics of the query. Also several factors that help in the processing of the sensed data need to be considered. Hence, there is a need to develop an adaptive routing protocol that considers the semantics of the query as well as several other criteria such as the energy remaining at nodes and also the power consumption model of the nodes. There have been numerous studies that have looked at the routing of data in sensor networks with the sole intention of reducing communication power or energy consumed. However there has been comparatively less prior art in the area of semantic routing and multi criteria-based routing algorithms that consider other performance goals.

The inter-communication of computing nodes in pervasive systems is an essential part of the system, and in the case of wireless sensor networks it could be described as a defining feature. The efficient management of such communication is crucial to the longevity of a system, directly affecting the power and energy requirements of the system. Additionally, the management algorithm for the

*Supported in part by the National Science Foundation award ANI-0325353 and ANI-0123705.

routing of data has a direct impact on the quality of the service provided by such system. In this paper we consider the example of a wireless sensor network, and focus on the management of routing trees to illustrate this point. We also focus on incorporating multiple criteria in sensor network routing. The problem we are looking at is as follows. In the process of creating a routing tree in a sensor network, always using the lowest energy path may not be optimal from the point of view of network lifetime and long-term connectivity. Other criteria also need to be considered such as the semantics of in-network processing, energy remaining at nodes and also query semantics.

The contributions of this paper include the introduction of a semantic and multi-criteria based routing protocol, which has shown significant performance improvement over the state of the art. Also, this scheme is inherently self-optimizing. We demonstrate performance improvements specifically in terms of *Network lifetime*, *Network Coverage* and the *Survivability of Critical Nodes*. Network lifetime is defined as the amount of time during which no less than a certain percentage of nodes remains alive. On the other hand network coverage is defined in a similar manner, but for a percentage of nodes that are distributed among equally divided groups (whereas for network lifetime it simply an overall percentage of nodes). Survivability of Critical Nodes is defined as the percentage of *critical* nodes alive, where critical nodes are the segment of nodes that need to be preserved the most in terms of energy.

The rest of the paper is structured as follows. We discuss background and related work in the next section, and go on to discuss the multi-criteria routing protocol in section 3. Then we discuss experimental results in section 4 and conclude with a brief discussion of future directions in section 5

2 Background and Related Work

While the simplest forms of routing depend solely on the passing of messages to neighbors that hear them, the most efficient routing schemes can be classified as either hierarchical or data-centric. Before we introduce our approach to routing in sensor networks, we will briefly discuss prior art, while touching upon the techniques from ad-hoc and mobile routing as well as quality assurance efforts. The simplest way to route data is to completely avoid the effort of constructing a route, and to pass the data along through flooding or gossiping [21]. This relies on a maximum number of message hops to guarantee receipt by all nodes. While this is adequate for distributing the data it's not efficient and so techniques for establishing routes

were developed that either used knowledge of data (data-centric routing) or the locations and identities of nodes (hierarchical routing).

Examples of data-centric routing include the early Sensor Protocol for Information via Negotiation (SPIN) [8], and the later Directed Diffusion [5, 11] and its numerous variants. SPIN used high-level meta-data to allow advertising and the on-demand retrieval of data. In Directed Diffusion, the main idea is to query the sensors in an on-demand manner, while the data has been maintained as attribute-value pairs to effectively name the data. This allows nodes to express and maintain lists of attribute-value pairs that represent interests, as well as reply links to neighbors based on the interests they expressed. These reply-links are known as gradients. As such the nodes in Directed Diffusion have the ability to do in-network data aggregation, which is modeled as a minimum Steiner tree problem [13]. Other algorithms include Constrained Anisotropic Diffusion Routing (CADR) [4], as well as the suggestion of employing multiple pre-planned paths to facilitate the choice of alternate paths without incurring the cost of searching for new routes [6], Rumor routing [3], Gradient-Based Routing [24] and a novel idea of information-directed routing proposed by Liu *et al* [15]. Prior art has also looking into using the query semantics to create a routing tree for the entire network [2].

While many data-centric routing protocols achieve energy efficiency and overhead advantages by avoiding the need to maintain topological information, hierarchical routing schemes actively use and develop this kind of information in constructing routes. The main idea behind hierarchical routing is efficient energy consumption of sensor nodes by involving them in multi-hop communication within a particular cluster and performing data aggregation and fusion in order to decrease the number of transmitted messages to the sink. Low-Energy Adaptive Clustering Hierarchy (LEACH) [9] was one of the first hierarchical routing algorithms for sensor networks. Other algorithms include Threshold-sensitive Energy Efficient sensor Network protocol (TEEN) [19], Adaptive Threshold sensitive Energy Efficient sensor Network protocol (APTEEN) [20], algorithms based on a three-tier architecture [32, 33] and techniques that use the router nodes to keep all the sensors connected by forming a dominating set [28].

Location awareness was utilized in many routing protocols originally intended for ad-hoc and mobile applications, but are amenable to sensor networks. Examples include GAF [30], Geographic and Energy Aware Routing (GEAR) [34], Minimum Energy Communication Network (MECN) protocol [23], Small Minimum Energy

Communication Network (SMECN) [14], and other protocols that actively attempted to improve the overall network lifetime of an ad-hoc network [12, 17, 25].

The Quality of Service (QoS) in terms of timing has been considered in the context of routing in Sequential Assignment Routing (SAR) [1, 27], by Maleki *et al* [18], and in the Stateless Protocol for Real-Time Communication in Sensor Networks (SPEED) [7].

COUGAR [16] is a data-centric routing protocol that views the network as a huge distributed database. The main idea is to use declarative queries in order to abstract query processing from the network layer functions and utilize in-network data aggregation to save energy. This abstraction is supported through a new query layer between the network and application layers. An architecture for the sensor database system where sensor nodes select a leader node to perform aggregation and transmit the data to the gateway (sink) was proposed. Thus, COUGAR provides a network layer-independent solution for querying the sensors and hence is one of the most popular data-centric protocols to date.

Another data-centric routing protocol is ACtive QUery forwarding In sensoR nEtworks (ACQUIRE) [22], which provides efficient querying via an adjustable range of neighborhood nodes. The routing problem in sensor networks has been addressed using various techniques of which TAG [31] and COUGAR [16] are among the very best, in the context of data-centric routing protocols. In this paper we present a data-centric routing protocol, which operates as a layer on top of existing in-network aggregation schemes, such as those provided by COUGAR and TAG, and adds the ability to minimize energy and power consumption as its primary goals.

Although it can be implemented over any data-centric routing protocol, our proposed scheme was tested over COUGAR. Both the TAG and COUGAR protocols consider communication in the form of a tree. The root of the tree is the base station. In-network aggregation is performed and the transmission needs to be synchronized to this root. Nodes are sensors in the network and they are organized in the form of a parent-children hierarchy. Both TAG and COUGAR differ in how the synchronization of message receipt and transmission is done. In COUGAR, synchronization is achieved using the idea of waiting list while in TAG it is achieved using the idea of communication slots.

Before we go on to describe our algorithm, we shall now briefly describe the Group aware Network Configuration (GaNC) algorithm [2]. The GaNC algorithm is similar to the First Heard From (FHF) algorithm that is used in TAG and COUGAR. The basic idea behind the FHF

network configuration algorithm is as follows: Starting from the root node, nodes transmit the new query. Child nodes select as their parent the first node they hear from and continue the process by further propagating the new query to all neighboring nodes. The process terminates when all nodes have been connected via the routing tree. In the GaNC algorithm, the main differing point is that a child can switch to a better parent while the tree is still being built unlike in the case of the First Heard From (FHF) algorithm. This switch is based on a set of tie-breaker conditions that go beyond the network characteristics and introduce the semantics of aggregation. The authors also suggested GaNCi [2], which is essentially the same as GaNC except that a child node could consider nodes from the same level as possible parents as well, during the process of parent selection. The authors show in terms of experimental results that the Group aware Network Configuration (GaNC) algorithm is currently the best performing algorithm in its class of multi-criteria algorithms that consider query semantics.

3 Multi-Criteria Routing Protocol

We now describe our algorithm for multi-criteria routing. In the sensor network, data is routed from various locations to a central sink of data, which is the root or the base station. Routing of data is assisted by assuming the sensor network to be configured in the form of a tree. The routing tree is created along with the propagation of the query. The base station propagates the query down the network. Traditionally, signal strength is the main factor considered when constructing the routing tree. A sensor node selects its parent based on the best link strength. Figures 1 and 2 present the overall view and the illustrate separation of the user specification and per-node algorithm.

3.1 Credit-Based Dynamic Route Update

We now propose our network configuration method that considers the semantics of the query and the properties of the sensors themselves: sensor properties refers mainly to the *energy remaining* at each sensor node and the *power consumption model* of each sensor node. The construction of the routing tree starts with a tree *build request* initiated by the root node, and propagated to the set of neighboring nodes. This message contains an identifier for the sender, the query specification, and a value representing the current level in the tree being constructed, $L(sender)$. For each node receiving the *build request*, the following steps are taken:

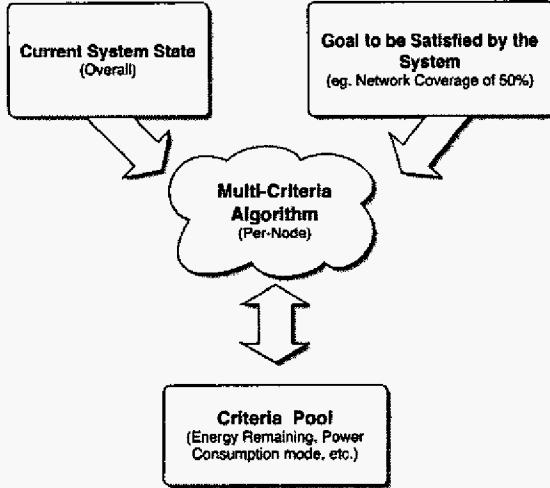


Figure 1: Overview of the System

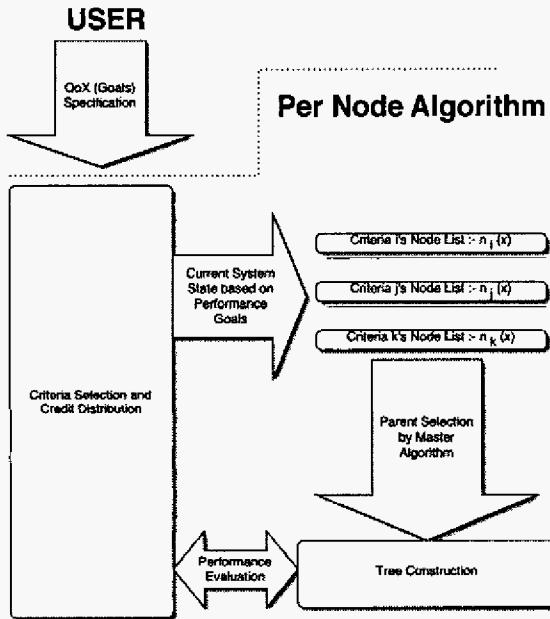


Figure 2: Separation of the user specification and per node algorithm

1. Upon initial receipt of the *build request*, a sensor node i sets its level value $L(i)$ to $L(sender) + 1$. It also records the parent value (Id) of the sender node, and its group ID. It then sends the tree build request to all its neighbors, but modifies it to reflect the new level, and itself as the sender.
2. A node will likely receive multiple *build requests* (from each of its neighbors), and upon subsequent receipts a node can switch to a “better” parent. The

definition of a better parent is determined by a summation of credit values, and is explained in detail below.

3. Steps 1 and 2 are repeated until all nodes have propagated the *build request* message.

For selecting a node’s parent, we consider the following criteria: *power consumption model* per node (in Watts), *energy remaining* at nodes (in Joules) and the *group membership* information. While the first two criteria are intuitive (nodes that have more energy or lower power consumption are preferred for increasing overall network lifetime), the significance of group membership requires clarification. Specifically, a favorable outcome is for nodes that will perform in-network aggregation of their data to fall along a common path in the routing tree (*i.e.*, share a parent-child relationship). In-network aggregation depends on the query attributes and the aggregation function. The list of attributes in the group-by clause divides the query result into a set of groups. The number of groups is equal to the number of combinations of distinct values of the attribute list. Hence readings from two different sensors are aggregated only if they are part of the same group. Since aggregation essentially combines all readings of a particular group into one, a tree in which all members of a group are in the same path is better in terms of overall energy consumed.

3.2 Neighborhoods and Criteria Lists

Our algorithm uses neighborhoods of nodes, and local per-node lists of such neighboring nodes. The concept of neighborhoods is similar to the use of a neighborhood of nodes (up to d hops away) in ACQUIRE [22], while the maintenance of a local list of nodes is similar to the approach of Woo *et al* [29]. In our algorithm, at each sensor node, we actually maintain three ordered lists of the neighboring nodes, and our neighborhood range is simply limited to the collision domain of the sensors, which we restrict to a maximum of three hops in our simulations. The lists we maintain are each ordered according to one of the evaluation criteria we use for selecting a parent node during tree construction, effectively offering three independent rankings of the neighboring nodes’ desirability based on each of the evaluation criteria.

The first list consists of the neighboring nodes ranked as per the power consumption model (in Watts) of each sensor node. The second list consists of the neighborhood range nodes ranked as per the energy remaining at the nodes. The third list consists of the neighborhood range nodes partially ranked as per the group information. This simply means that nodes in the same group

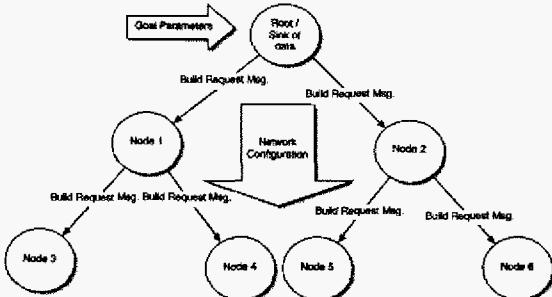


Figure 3: Initiation of Tree Construction

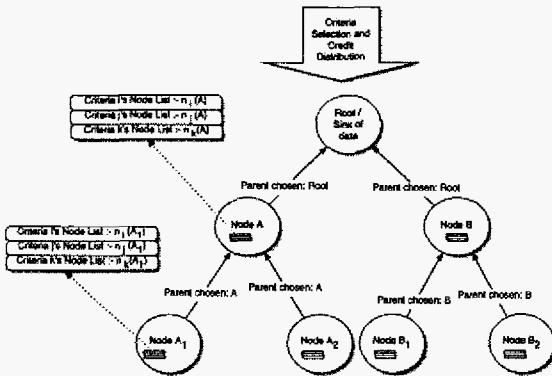


Figure 4: Actual Tree Construction

are placed at the head of the list, with no strict distinction among them. The nodes transfer this information in message headers that are transmitted back and forth between the nodes. The maintenance of tables and the overhead is limited and well justified by the performance gains that we observed. The data communication costs associated with our approach, which is essentially the overhead introduced, is in the order of magnitude of the data transmitted since we have incorporated our control information in the headers of the data transmitted. This is independent of the underlying network.

The energy remaining at each node is measured in Joules and for simulation purposes we define a maximum value equivalent to the energy of a battery cell. Over a period of time as the node transmits data and performs various computations, we reduce this value appropriately for the various operations and when the minimum value is reached, we mark the node as dead. We model the power consumption of each sensor node drawn randomly from a distribution and assume a rate of decay for all the sensor nodes. The group information is modeled as participation based on group identifier. The initiation of tree construction and the actual construction of the routing tree are illustrated in figures 3 and 4.

3.3 Updating Credits

Initially we define a set of goals that need to be satisfied. This is drawn from a pre-determined set of goals that the system might want to fulfill. For instance one possible goal is based on the number of nodes alive such as Network lifetime of 50%.

Our terminology and approach for the problem is as follows. The “criteria lists” are that of criteria that make independent recommendations; and the corresponding set of credits (for each criteria) are a representation of the suitability of the recommendations with regards to the desired goal. The ordering that a node finally uses to select its parent is based on a weighted combination of the orderings offered by the three criteria lists.

The credits of criteria lists (a number of credits for each of the lists, assigned from a central pool) represent their quality of recommendations, and this is based on desired system goals. Depending on the current desired goals, we define the distribution of credits. In other words, credits for the lists are in effect specifying the mix of criteria that best achieve the goal. This is in contrast to static schemes such as GaNC, in which the order and hence the significance of criteria (that is the tie-breaker conditions) is fixed.

Initially the credits are distributed uniformly among all criteria. This initial distribution of credits is specified in the build request message that is transmitted from the root to all nodes. Our multi-criteria routing algorithm decides the parent for each node with a weighted average of the criteria lists. Depending on the observed outcome (successful achievement of goals) the base station updates credits among criteria. Updating the credits involves a redistribution of the credits among the lists, but it should be noted that no list is allowed to drop below a minimal value for the credits. This avoids a criteria having a zero credit, which would prevent it from contributing to the algorithm and being re-evaluated. A minimal value of credits allows a criteria to have no impact, while continuing to be considered for increased value in the future. We now go on to describe the details of credit updates.

3.4 Proportional Credit Updates

The redistribution of credits is done globally. In other words, we check periodically if the goal is satisfied. If a certain goal is not satisfied, then the credits are redistributed proportionately and the network is reconfigured. That is after every redistribution, the credits are then sent out in the *build request* message.

We have assumed here that the base station has global information of alive and dead nodes for all those in the

network. Whenever a node transmits its reading, we also require it to send this information in the header of the message.

Let us illustrate this proportional update of credits scheme with an example in the following manner. Let the desired goal be a network lifetime of 50%. After a regular evaluation period, which was recorded as 10 minutes in our experiments, we check if more than 5% of nodes have died since the last evaluation period. If not, we do not update the credit distribution for the three criteria lists. If more than 5% of nodes have died since the last evaluation period, then proportional to the difference between the dead nodes and the number of nodes that are in fact alive, the credits for the first two criteria lists (energy remaining and power consumption rate) are increased relative to the third criteria list (participation in the group). We note here that we have considered the two criteria lists of energy remaining and power consumption since they exclusively favor the reduction of power consumption and the increased longevity of sensor nodes. We also note here that an ideal dynamic list update scheme is the subject of continued investigation. In the next section, we describe our data collection in concrete terms.

4 Performance Evaluation

We analyze our approach using simulation. We assume a grid of sensors in which the range of transmission is restricted to a single hop. This is in keeping with the basic assumption of various other in-network aggregation schemes. This can be generalized to non-uniform grid configurations simply by reduction in the size of the grid. We evaluate performance in terms of three performance metrics: network lifetime, network coverage and survivability of critical nodes. Network lifetime and coverage deal with the time during which a percentage of nodes can remain alive, whereas survivability of critical nodes focuses on the need to maximize the run time of critical nodes. In our simulator, we assume the piggybacking of data that is used for sending the control information.

4.1 Experimental Setup and Workload

The simulator was written using C++ and csim. The credit points were shaped from a pool of size 100. We simulated various sensor network grid sizes ranging from 15×15 to 50×50 .

Network lifetime is defined as the amount of time during which no less than a certain percentage of nodes remains alive. On the other hand network coverage is defined in a similar manner, but for a percentage of

nodes that are distributed amount equally divided groups (whereas for network lifetime it simply an overall percentage of nodes). Survivability of Critical Nodes is defined as the percentage of *critical* nodes alive, where critical nodes are the segment of nodes that need to be preserved the most in terms of energy.

For our experiments, we focused on the standard SQL aggregation functions SUM, AVERAGE, and MAX. We did not include the MIN function, which is similar to MAX, nor did we include COUNT, which is similar to SUM.

We simulated Network lifetime ranging from 40% to 70%. We simulated Network coverage ranging from 40% to 70%. The range was chosen to reflect the most generic values that might be considered reasonable when deploying a wireless sensor network. We measured the Survivability of Critical Nodes for Network coverage and Network lifetime ranging from 40% to 70%. We present results for Network coverage in figures 5 and 6, Network lifetime in figures 7 and 8, and the Survivability of Critical Nodes in figure 9. For simulation purposes, we restrict to the collision domain by getting information upto a maximum of three hops.

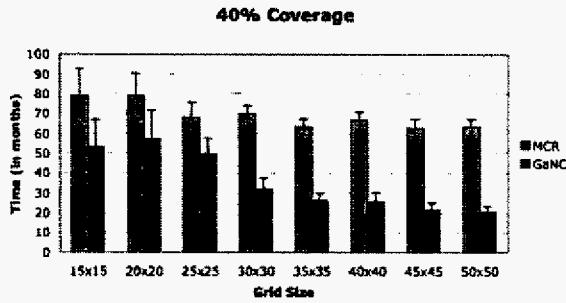
In our simulations the multi-criteria routing protocol was shown to outperform the Group aware Network Configuration (GaNC) [2] algorithm in all the measured metrics: network lifetime, network coverage and the Survivability of Critical Nodes.

4.2 Network Coverage

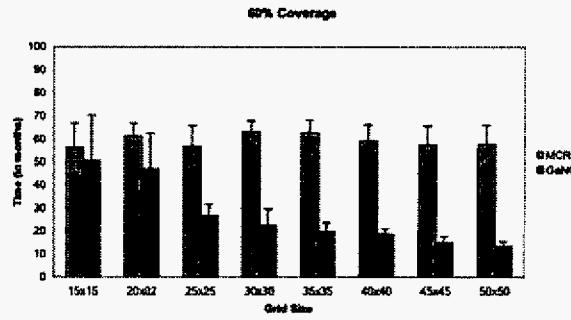
Figures 5 and 6 look at the comparison of the two protocols on the basis of Network Coverage. In these figures, the x-axis is varying grid sizes from 15×15 to 50×50 . The y-axis is the time at which the network dies. We consider Network Coverages of 40%, 50%, 60% and 70%.

As mentioned above, in all our experiments, the redistribution algorithm executes periodically every 10 minutes and checks for decrements of 5% in active nodes. All results are time values measured in terms of months. The values used for transmission consumption are simulated as a percentage of the total energy of each sensor node, in the same manner as described by Hill *et al* [10]. The energy values are maintained per node and updated (locally to each node) with every transmission that the node performs. Its important to note that these values are kept current. In this manner our simulation faithfully considers the energy overheads of the routing tree construction. All experiments were run multiple times to evaluate statistical variation.

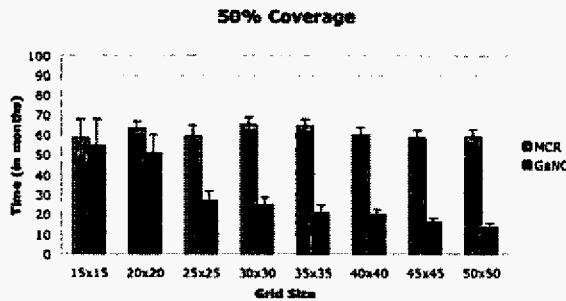
It is clear from these figures that the multi-criteria rout-



(a) Network Coverage of 40%



(a) Network Coverage of 60%



(b) Network Coverage of 50%

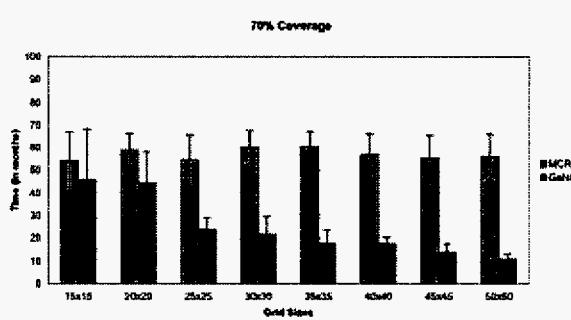
Figure 5: Network Coverage Comparison

ing policy fares consistently better. This can be attributed to the fact that the nodes' parents are redistributed, thus preventing any single node from over-expending its energy. Another factor that we have taken into account is the power consumption model of the sensors themselves, which is also being considered for all the sensors in the network as a whole.

Hence we see that the multi-criteria routing policy outperforms the GaNC scheme, offering continually prolonged network activity. That this does not degrade for larger network grid sizes is particularly noteworthy.

4.3 Network Lifetime

Figures 7 and 8 look at the comparison of the two protocols on the basis of Network Lifetime. As with figures 5 and 6, the x-axis is varying grid sizes from 15×15 to 50×50 . The y-axis is the time (in months) at which the network dies. We again look at the Network Lifetimes of 40%, 50%, 60% and 70%.



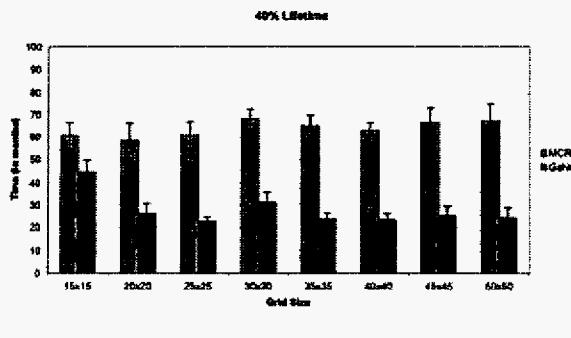
(b) Network Coverage of 70%

Figure 6: Network Coverage Comparison

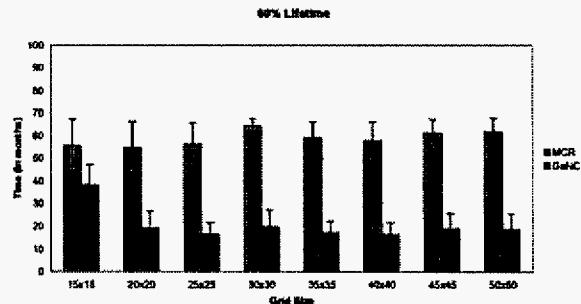
All experiments were run multiple times to evaluate statistical variation. For Network Lifetime, our multi-criteria routing policy again performed consistently better. The most impressive results appear again at larger grid sizes, and are largely a result of our multi-criteria policy continuing to perform well. This is in contrast to the degradation in performance for GaNC. It is also worthy of note that the variation in results are less at these larger grid sizes, where the difference in performance is greatest. We therefore see that for Network Lifetime, as for Network Coverage, our policy outperforms GaNC without degradation for larger network sizes.

4.4 Survivability of Critical Nodes

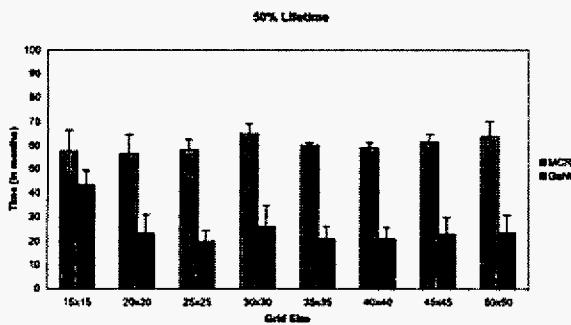
In figure 9, we compare the Survivability of Critical Nodes for varied Network lifetimes. Here we define critical nodes and measure the lifetime of those nodes as the Survivability of Critical Nodes in the network. For instance, we define the first row of nodes as critical nodes



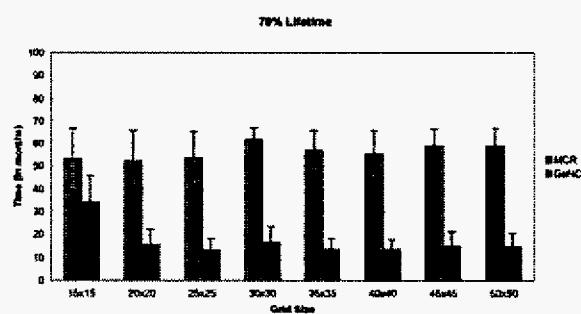
(a) Network Lifetime of 40%



(a) Network Lifetime of 60%



(b) Network Lifetime of 50%



(b) Network Lifetime of 70%

Figure 7: Network Lifetime Comparison

in the network for simulation purposes. In these figures, the x-axis is varying Network lifetimes ranging from 40% to 70%. The y-axis is the Survivability of Critical Nodes measured as a percentage of critical nodes alive.

We can see that the multi-criteria routing protocol appears to outperform the Group aware Network Configuration (GaNC) protocol, since there is no mechanism in it to specify a particular subset of nodes as critical. We have observed similar results for Network coverage based on Survivability of Critical Nodes.

In summary, we can see that the multi-criteria routing protocol offers significant improvements over the Group aware Network Configuration (GaNC) protocol in terms of Network Coverage, Network Lifetime and the Survivability of Critical Nodes. Moreover the overhead of power consumption for tree construction is comparable to it. Hence it is a very good mechanism for in-network aggregation and is quite versatile as it can be deployed over TAG or COUGAR.

Figure 8: Network Lifetime Comparison

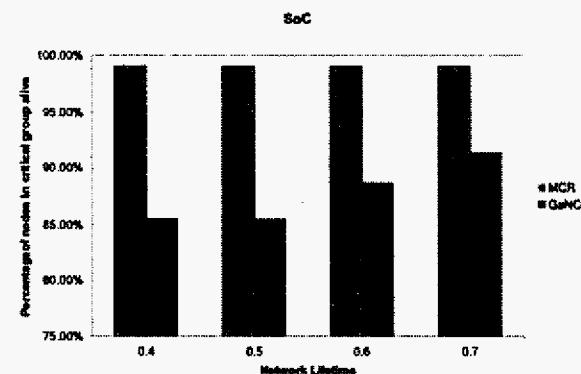


Figure 9: Comparison of Survivability of Critical Nodes (SoC) for varied Network Lifetimes.

5 Conclusions

In this paper, we have designed and implemented a multi-criteria routing scheme for sensor networks with pervasive services in mind. Our scheme exhibits significant performance improvement with minimal overhead when compared to the current state of the art routing algorithms, specifically with respect to Network Lifetime, Network Coverage, and Survivability of Critical Nodes (SoC).

While we have shown that our adaptive multi-criteria algorithm improves the longevity of system nodes, and the longevity of the connected network, we have also seen that it results in a higher quality of service by allowing the survival of more critical nodes. This example was specific to wireless sensor networks, but any pervasive computing system that depends on the interconnection of its nodes for its services, and the routing of data among them, could benefit from such adaptive multi-criteria algorithms for the management of communication routing. As part of our future work, we aim to consider varied query frequencies, and varied (e.g., non-uniform) distributions of nodes.

References

[1] I. Akyildiz, W. Su, Y. Sanakarabramaniam, and E. Cayirci. Wireless sensor networks: a survey. *Computer Networks*, 38(4):393–422, 2002.

[2] J. Beaver, M. A. Sharaf, A. Labrinidis, and P. K. Chrysanthis. Location-aware routing for data aggregation for sensor networks. In *Proc. of the Geo Sensor Networks Workshop*, 2003.

[3] D. Braginsky and D. Estrin. Rumor routing algorithm for sensor networks. In *Proc. of the 1st Workshop on Sensor Networks and Applications (WSNA)*, 2002.

[4] M. Chu, H. Haussecker, and F. Zhao. Scalable information-driven sensor querying and routing for ad hoc heterogeneous sensor networks. *The Int'l Journal of High Performance Computing Applications*, 16(3), 2002.

[5] D. Estrin, R. Govindan, J. S. Heidemann, and S. Kumar. Next century challenges: Scalable coordination in sensor networks. In *Proc. of the 5th Annual Int'l Conference on Mobile Computing and Networking (MOBICOM)*, 1999.

[6] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin. Highly-resilient, energy-efficient multipath routing in wireless sensor networks. *Mobile Computing and Communications Review*, 1(2), 2001.

[7] T. He, J. A. Stankovic, C. Lu, and T. Abdelzaher. SPEED: A Stateless Protocol for Real-Time Communication in Sensor Networks. In *Proc. of the 23rd Int'l Conference on Distributed Computing Systems (ICDCS)*, 2003.

[8] W. Heinzelman, J. Kulik, and H. Balakrishnan. Adaptive protocols for information dissemination in wireless sensor networks. In *Proc. of the 5th Annual Int'l Conference on Mobile Computing and Networking (MOBICOM)*, 1999.

[9] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-efficient communication protocol for wireless microsensor networks. In *Proc. of the Hawaii Int'l Conference on System Sciences (HICSS)*, 2000.

[10] J. L. Hill and D. E. Culler. Mica: A wireless platform for deeply embedded networks. *IEEE Micro*, 22(6):12–24, 2002.

[11] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In *Proc. of the 6th Annual Int'l Conference on Mobile Computing and Networking (MOBICOM)*, 2000.

[12] F. D. Jie Wu, M. Gao, and I. Stojmenovic. On calculating power-aware connected dominating sets for efficient routing in ad hoc wireless networks. *IEEE/KICS Journal of Communications and Networks*, 4(1):59–70, 2002.

[13] B. Krishnamachari, D. Estrin, and S. Wicker. Modelling data-centric routing in wireless sensor networks. In *Proc. of IEEE INFOCOM*, 2002.

[14] L. Li and J. Y. Halpern. Minimum energy mobile wireless networks revisited. In *Proc. of the IEEE Int'l Conference on Communications (ICC)*, 2001.

[15] J. Liu, F. Zhao, and D. Petrovica. Information-directed routing in ad hoc sensor networks. In *Proc. of the 2nd ACM Int'l conference on Wireless Sensor Networks and Applications (WSNA)*, 2003.

[16] S. R. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. The Cougar approach to in-network query processing in sensor networks. In *Proc. of the ACM Symposium on Operating System Design and Implementation (OSDI)*, 2002.

- [17] M. Maleki, K. Dantu, and M. Pedram. Power-aware source routing protocol for mobile ad hoc networks. In *Proc. of the Int'l Symposium on Low Power Electronics and Design (ISLPED)*, 2002.
- [18] M. Maleki, K. Dantu, and M. Pedram. An energy-aware QoS Routing Protocol for Wireless Sensor Networks. In *Proc. of the IEEE Workshop on Mobile and Wireless Networks (MWN)*, 2003.
- [19] A. Manjeshwar and D. P. Agrawal. TEEN : A Protocol for Enhanced Efficiency in Wireless Sensor Networks. In *Proc. of the 1st Int'l Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing*, 2001.
- [20] A. Manjeshwar and D. P. Agrawal. APTEEN: A Hybrid Protocol for Efficient Routing and Comprehensive Information Retrieval in Wireless Sensor Networks. In *Proc. of the 2nd Int'l Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile computing*, 2002.
- [21] E. L. Miller and R. H. Katz. A survey of gossiping and broadcasting in communication networks. *Networks*, 18(4):319–349, 1988.
- [22] N. Sadagopan, B. Krishnamachari, and A. Helmy. The ACQUIRE mechanism for efficient querying in sensor networks. In *Proc. of the 1st IEEE Int'l Workshop on Sensor Network Protocols and Applications (WSNA)*, 2003.
- [23] V. Rodoplu and T. Meng. Minimum energy mobile wireless networks. In *Proc. of the IEEE Int'l Conference on Communications (ICC)*, 1998.
- [24] C. Schurgers and M. Srivastava. Energy efficient routing in wireless sensor networks. In *MILCOM Proc. on Communications for Network-Centric Operations: Creating the Information Force*, 2001.
- [25] R. C. Shah and J. Rabaey. Energy aware routing for low energy ad hoc sensor networks. In *Proc. of the IEEE Wireless Communications and Networking Conference (WCNC)*, 2002.
- [26] M. A. Sharaf, J. Beaver, A. Labrinidis, and P. K. Chrysanthis. TiNA: a scheme for Temporal Coherency-Aware in-Network Aggregation. In *Proc. of the 3rd ACM Int'l Workshop on Data Engineering for Wireless and Mobile Access (MobiDE)*, 2003.
- [27] K. Sohrabi, J. Gao, V. Ailawadhi, and G. J. Pottie. Protocols for self-organization of a wireless sensor network. *IEEE Personal Communications*, 7(5):16–27, 2000.
- [28] L. Subramanian and R. H. Katz. An architecture for building self-configurable systems. In *Proc. of the 1st ACM Int'l Symposium on Mobile Ad Hoc Networking and Computing*, 2000.
- [29] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *Proc. of the 1st Int'l conference on Embedded networked sensor systems*, 2003.
- [30] Y. Xu, J. Heidemann, and D. Estrin. Geography-informed energy conservation for ad hoc routing. In *Proc. of the 7th Annual Int'l Conference on Mobile Computing and Networking (MOBICOM)*, 2001.
- [31] Y. Yao and J. E. Gehrke. TAG: a Tiny AGgregation Service for ad-hoc sensor networks. In *SIGMOD Record*, 2002.
- [32] M. F. Younis, M. Youssef, and K. A. Arisha. Energy-aware routing in cluster-based sensor networks. In *Proc. of the 10th Int'l Workshop on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, 2002.
- [33] M. A. Youssef, M. F. Younis, and K. A. Arisha. A constrained shortest-path energy-aware routing algorithm for wireless sensor networks. In *Proc. of the IEEE Wireless Communications and Networking Conference (WCNC)*, 2002.
- [34] Y. Yu, R. Govindan, and D. Estrin. Geographical and energy aware routing: A recursive data dissemination protocol for wireless sensor networks. In *UCLA Computer Science Department Technical Report UCLA/CSD-TR-01-0023*, 2001.