Energy-efficient Search for Finite-lifetime Resources in Sensor Networks with Time-constrained Queries*

Christopher R. Mann[a] **Rusty O. Baldwin**[a] **Jeffrey P. Kharoufeh**[b] **Barry E. Mullins**[a]
christopher.mann@afit.edu *rusty.baldwin@afit.edu* *jkharoufeh@coe.neu.edu* *barry.mullins@afit.edu*
[a]Dept of Electrical and Computer Engineering, Air Force Institute of Technology
2950 Hobson Way, Wright-Patterson Air Force Base, OH 45433, USA
[b]Dept of Mechanical and Industrial Engineering
334 Snell Engineering Center, Northeastern University, Boston, MA 02115, USA

We examine the performance of a random-walk search algorithm for wireless sensor networks when resources are subject to limited lifetimes and queries are constrained by application-specific deadlines. Specifically, via the time-to-live and transmission range parameters, we estimate the appropriate number of resource copies that must be created within the network to minimize the total node arrival rate (the energy-centric approach) or to ensure the total proportion of queries failures does not exceed a specified threshold (the failure-centric approach). The effect of node transmission range on network performance is also investigated. We compare the results of our network simulations to our queueing-based analytic node model and find that there is an inverse relationship between transmission range and the time-to-live value required to minimize the total node arrival rate.

I. Introduction

Wireless sensor networks (WSN) are created through the cooperative operation of hundreds or thousands of sensing devices called *nodes*. Although present-day networks are necessarily limited in size due to technological and cost limitations, future WSNs may be composed of millions of nodes and perform tasks such as weather and environmental monitoring, or provide additional security enhancements such as border monitoring or treaty compliance applications. These nodes are linked via a wireless transmission medium to perform network tasks in a distributed manner. To keep total deployment costs low, individual node capabilities are necessarily limited. For example, a typical wireless sensor node has a finite energy reserve, a small amount of local memory, and limited computational capabilities. To conserve power and reduce contention for access to the transmission medium, it is usually advantageous to limit each node's effective transmission range to that required to ensure a connected network. This requires that nodes act as routers for neighboring nodes' packets when the intended receiver is beyond the transmission range of the origin node.

As a consequence of limited node capabilities and the distributed operation of wireless sensor networks, nodes are forced to find information and services when local access is not available. To locate these resources within the network, nodes initiate requests which are propagated though the network. These requests, however, not only force the requesting node to expend a portion of its available energy reserves, but also require neighboring nodes to expend a portion of their available energy as well. Since the effective lifetime of the network is directly related to the lifetimes of the component nodes, it is desirable to minimize the total network energy expenditure required to answer the originating node's request.

In addition to the energy expended by nodes to transmit and receive requests, two other considerations also affect the total energy expended to answer a specific request. First, nodes may be subject to application constraints which require the needed information to be returned to the originating node prior to the expiration of a specific deadline. From the perspective of energy efficiency, it is also unwise to permit requests to be forwarded indefinitely in the event the desired resource proves difficult to locate, does not exist within the network, or can no longer be used by the requesting node. To prevent such excessive energy expenditure, then, it is prudent to ensure resource requests are issued with specific deadlines. Second, when a node generates a particular resource (e.g., by sensing a reportable event), such observations are typically valid only for a finite time period. Additionally, wireless sensor networks have a dynamic phys-

*The views expressed in this paper are the views of the authors and do not necessarily reflect the views or policies of the United States Air Force.

ical topology as a result of node failures and mobility. This limits the availability of a particular resource and also means the utility of results obtained from past searches will diminish over time. To prevent the use of stale results for current operations, it is sensible to assign an expiration time to each resource at the time it is generated.

Although requests could be flooded to every node in the network in an effort to locate the desired resource, this approach is wasteful in terms of the total energy expended by the network [4]. To decrease reliance on flooding and reduce the total energy expenditure required to answer a node's request, several search paradigms are available. Two of the most common approaches can be broadly characterized as "geo-centric" and "data-centric." The former applies when the originating node has no information on the probable location of the desired resource. The request is forwarded from node to node until the resource is located. To reduce the energy expenditure and latency associated with the geo-centric approach, nodes "advertise" the availability of a particular resource to a subset of the network's nodes, and these advertisements are distributed in manner independent of resource characteristics (e.g., along a random walk or curvilinear path). This is the approach taken by [2, 3, 4, 5, 8, 16]. To control the number of resource copies in the network, a *time-to-live* (TTL) counter is sometimes associated with each resource advertisement. Additionally, most "quorum-based" search algorithms, i.e., algorithms in which advertisements and requests are propagated along cardinal directions thus guaranteeing advertisement request intersection, can be classified as geo-centric (e.g., [7, 10, 15]).

By contrast, data-centric search algorithms such as those in [12, 13, 14] use a hash table and the attributes of the resource to determine the appropriate storage location within the network. Assuming all nodes have access to the correct hash table, the location of a particular resource (if it exists) is immediately known to the originating node; requests are forwarded directly to the appropriate node(s). Although data-centric approaches reduce the latency and energy expenditure associated with answering requests, they are susceptible to "hotspotting," localized congestion, and insufficient local storage capacity. Furthermore, data-centric approaches tend to aggregate related data in common locations; hence, a network partition, node failure, or destruction of a portion of the network can result in total loss of a particular data type. Because all nodes must have access to the same hash table, periodic updates may be required and are expensive in terms of total energy expenditure, especially if the affected data must also be relocated.

This is not to say that one approach or the other is superior. The appropriate approach depends on the network's intended purpose and the specific requirements of the application. However, since the geographic dispersal and data redundancy achieved by geo-centric approaches is inherently greater than that attained by pure data-centric approaches, we advocate the geo-centric approach in wireless sensor networks with large numbers of low-cost, failure-prone nodes deployed in hazardous environmental conditions.

In this paper, we use a simulation model to estimate the number of resource advertisements, or *replicates*, required either to minimize each node's total arrival rate of resource requests and advertisements (the energy-centric approach) or to ensure a specified maximum level of query failures is not exceeded (the failure-centric approach). This complements our previous work [9] where we develop a Markov chain model of individual nodes to analyze the performance of a random-walk search algorithm. To ensure the tractability of the Markovian model, however, certain simplifying assumptions were required. Most importantly, both advertisements and requests for a particular resource have lead times (i.e., the time remaining until expiration) that, upon arrival at a node, are exponentially distributed with (possibly) dissimilar means. It is more likely, however, for expiration times to be assigned to requests and advertisements by the originating node at the time of generation. When a request/advertisement arrives at a node, the lead time is a consequence of the originally assigned expiration time less any processing, queueing, and transmission delays experienced at previously-visited nodes. Therefore, the actual distribution of lead times of arriving requests and advertisements may not resemble the original distribution. Moreover, the model presumes the expiration times assigned to each agent permits the desired number of agent copies to be stored by the network. That is, the agents' TTL counters are always exhausted before their expiration times occur. Additionally, the distribution of nodes possessing a local copy of a particular agent type is assumed to be uniform throughout the network. As node transmission range is reduced, however, each node's one-hop neighborhood necessarily decreases, thus decreasing both the uniformity of agent distribution and the probability of locating an agent far from its point of origin. Finally, the Markov chain node model assumes the interarrival times of both agents and queries, whether generated locally by the node itself or received from a neighboring node, are exponentially distributed. Whether or not this assumption will hold in a network composed of thousands of nodes is unclear.

While the Markov chain model is useful for predicting the mean performance of individual nodes within the scope of the original assumptions, accurate analytic modeling of the effects of various lead time distributions, agent deployment methods, and transmission range on overall network performance is difficult; studies of such parameters are currently limited to simulation models. The purpose of this paper is to determine how effects that are difficult or impossible to capture in the analytical model affect the performance of a random walk search algorithm in a network.

The search algorithm we employ in this research is most closely related to the rumor routing algorithm [4]. When a *witness node* observes an event, it informs a portion of the network's nodes via a specialized packet called an *agent*. When an agent arrives at a node, the node adds a copy of the agent's information to its *event table*, a local memory cache. When a node has information related to a specific event stored in its event table, the node is considered to be *informed* and is capable of answering locally-generated requests as well as requests received from uninformed nodes. *Queries* are generated when a node's application requires access to a specific resource. If the resource is unavailable locally, the query is forwarded to neighboring nodes until either an informed node is found or the query's lead time expires. When a node receives a query for which it has corresponding information stored in its event table, the node generates a *response* and returns the desired information to the node that originated the query.

The remainder of this paper is organized as follows. Section II discusses the literature related to our work. In Section III, we develop a model of a wireless sensor node that incorporates each node's event table, transmission queue, transceiver, sensors, and applications. Two important indicators of network performance—the total arrival rate and the total proportion of query failures—are discussed in Section IV. The results of simulations of networks with large node populations are analyzed in Section V. Section VI provides our concluding remarks.

II. Related Work

In addition to the aforementioned papers, there are several geo-centric search protocols based on rumor routing (e.g., [2, 3, 5, 16]), but these algorithms do not attempt to determine the degree of resource replication required to minimize total network energy expenditure. The Trajectory-based Selective Broadcast Query (TSBQ) protocol [8], a rumor routing-based search protocol, uses straight-line trajectories and a "selective broadcast" primitive to conduct energy efficient searches. An analytic model of TSBQ determines the resource replication level that minimizes the total network energy expenditure required to advertise a resource's availability, locate the resource via one or more queries, and return a response to the originating node(s). The optimum replication level for expanding ring search algorithms was determined analytically in [6]. A simulation model was used to ascertain the appropriate replication level in the REDMAN algorithm [3]. However, none of these protocols associate expiration times with agents or queries, and only [8] integrates the effect of resource popularity on optimal replication levels. To the best of our knowledge, only [9] provides an analytic means to evaluate optimum resource replication levels when both agents and queries have finite lifetimes. This work expands on [9] by using a simulation model to investigate the effects of agent/query lifetime and node transmission range on total energy expenditure, the total proportion of query failures, and the optimal time-to-live value.

III. Node Model

To examine the effects of various parameters on the performance of random-walk search algorithms, each node is modeled in OPNET as a wireless transceiver with a fixed maximum transmission/reception range, an event table, and a transmission queue (see Figure 1). The activity of an on-board sensor is represented by a processor which creates new agents in response to external stimuli, and the "Application" creates queries for information needed to complete node tasks. The purpose of the "Splitter" is to ensure copies of agents received from neighboring nodes are forwarded to the event table and—if the agent's TTL counter has not been exhausted—also to the transmission queue to be scheduled for forwarding to a neighboring node. The splitter has no effect on queries other than to forward the query or its corresponding response directly to the transmission queue. Since the splitter represents a simple decision function, it adds no additional processing delay to arriving agents or queries.

Each agent arriving to the event table is retained until its expiration time passes. Hence, the operation of

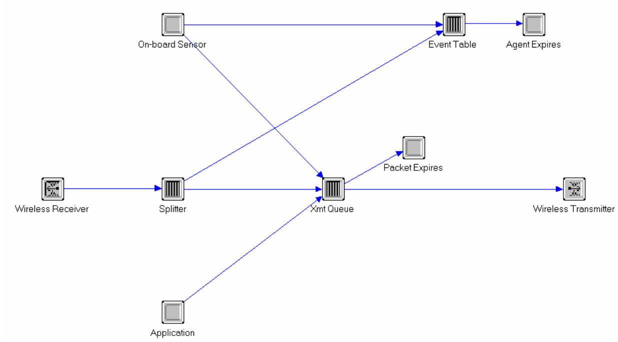

Figure 1: Wireless sensor node model in OPNET

the event table resembles that of a G/G/∞ queue. If the event table contains at least one unexpired agent of a particular type, the node is considered to be informed of that event and is capable of answering related queries. When the node's application generates a query, the node first checks its local event table for a corresponding agent. If a matching agent is found, the query is answered locally; there is no need to add the query to the transmission queue. If, however, the node is uninformed or if the query originated externally, the query (response) is sent to the transmission queue and scheduled for transmission using a FIFO service discipline. Due to contention for access to the transmission medium as well as the potential for retransmissions, we assume each agent/query requires an exponentially distributed amount of time to be successfully transmitted to the designated receiver. Prior to the beginning of each query transmission, the node checks its event table for an agent that matches the query's request. If the desired information is found, the node transmits the appropriate response in place of the query. If no corresponding agents are found, the node transmits the query to a randomly-chosen neighbor. Agents and queries expiring prior to transmission are removed from the transmission queue. The transmission queue is therefore a FIFO G/M/1 queue with customer reneging as described in [9].

A network of nodes based on the analytic node model in [9] resembles a Jackson network of queues. The random arrival of agents and queries to each node is assumed to occur according to a Poisson process, the random time between successive departures of agents and queries from a node's transmission queue is exponentially distributed, and agents/queries are either forwarded to another node or depart the system with specific probabilities. However, the problem is complicated by the existence of three customer types (i.e., agents, queries, and responses), and each customer type must vie for access to the transmission medium at each node. Moreover, the rate of arrival of agents to each node as well as the expiration time assigned to each agent/query determines the probability that a query will be forwarded to a neighboring node or depart the system (i.e., fail). Even so, we show in Section 5 the analytic node model provides an accurate prediction of mean network performance.

Node parameters that can be specified by the user prior to execution of the simulation model are summarized in Table 1. All nodes within the network are assumed to be indistinguishable with respect to these parameters. The primary means for controlling the number of resource copies per agent stored in the network is through the TTL parameter. The next section discusses the TTL parameter and the significance of our chosen metrics.

IV. Metrics

There are two primary indicators of network performance we wish to estimate: the mean total arrival rate of agents and queries and the total proportion of failed queries. Using these metrics, we estimate the appropriate agent TTL to minimize the total transmission energy expended by the network while not exceeding the maximum tolerable level of query failures.

During its useful lifetime, a node may undertake several energy-consuming activities: transmission, reception, computation, sensing, and sleeping. Of these activities, transmission normally requires the greatest total energy expenditure in present-day devices; the expenditure for transmission can be nearly twice as much as the next most energy-expensive activity of reception [11]. The sleep state consumes the least energy and is the preferred state when a node is not needed for support of other network activities. Since the lifetime of a network is directly related to the lifetimes of its nodes, minimizing energy expenditure is critical to ensuring the longevity of the network [1]. Since our model assumes agents, queries, and responses are forwarded by the transmitting node to a single receiver, estimating the total rate of transmission arrivals at each node is indicative of the network's total energy expenditure and, hence, network lifetime. The goal of our energy-centric metric, then, is to minimize the total rate at which transmissions are received by each node and, as a consequence, to reduce the network's total energy expenditure. Sole reliance on an energy-centric metric, however, cannot guarantee nodes receive information at a rate that is sufficient to satisfy application requirements and also accomplish the network's objectives.

Table 1: User-adjustable simulation parameters

Module	*Parameter*	*Description*
On-board Sensor	TTL	The maximum number of times a single agent may be transmitted
	λ	The mean arrival rate of reportable (i.e., agent-generating) events
	$1/\delta$	The mean lead time assigned to an agent upon its generation
Application	γ	The mean arrival rate of queries generated by the node's application
	$1/\beta$	The mean lead time assigned to a query upon its generation
Transmission Queue	$1/\mu$	The mean time required to process and successfully transmit an agent/query to the intended recipient

If a sufficient proportion of each node's queries remain unanswered, the probability of general network application failure increases. Therefore, we must ensure the total proportion of failed queries observed by each node is less than the application-specific threshold. We define a query failure in the following manner:

> **Definition**: A *query failure* occurs when a query (or, if the node is informed, the query's corresponding response) expires in the node's transmission queue before it can be transmitted [9].

Based on this definition, the proportion of query failures in the network, ε, is obtained by dividing the total number of expired queries/responses observed in the network by the total number of unique queries generated. The goal, then, is to ensure ε does not exceed a specified maximum.

V. Simulation Results

A necessary first step is to validate the simulation model by configuring it to adhere as closely as possible to the assumptions made in the analytical queueing model. Most importantly, the analytic model assumes agents are uniformly spatially distributed throughout the network. As noted previously, however, short node transmission ranges affect the uniformity of agent dispersal. Therefore, to ensure the simulation achieves a uniform distribution of informed nodes, we artificially extend the transmission range of the nodes such that each node is a one-hop neighbor of every other node in the network (i.e., $>$5000m); the effects of contention are momentarily ignored. The nodes are configured according to the parameters in Table 2.

The placement of nodes within the confines of the deployment area is determined randomly using the random topology generating feature of OPNET prior to the beginning of the simulation. This topology, once created, is held constant throughout each set of simulations and experimental replicates to ensure any effects due to node placement are identical across each test set. Experimental testing indicated the simulation model reliably achieved steady state operation within 60 seconds of elapsed simulation time. Therefore, for each set of parameters, the network is permitted to operate for a period of 60 seconds prior to the collection of performance data. After initialization is complete, performance data is collected at every node in the network for a simulated time period of 900 seconds. The total arrivals per node per second and the total proportion of failed queries in the network are shown in Figures 2 and 3. Where depicted, 95% confidence intervals are used.

The results of the simulation using a large node transmission range indicate the analytic node model closely predicts the performance of the network. However, for TTL values less than 36, the arrival rate per node in the simulations is slightly higher than predicted; the maximum differential is a very modest 2.7 additional packets per node per 100 seconds of simulation time. This additional traffic is attributed to the fact that agents generated in the simulation may expire prior to exhausting their TTL counters whereas the analytic model assumes each agent is replicated exactly TTL times prior to expiration. The result is that the actual proportion of the network informed of an event at any given instant is smaller than that assumed by

Table 2: Parameters for simulation validation

Parameter	*Distribution*	*Mean*
Agent interarrival time	Exponential	200.000 sec
Agent lead time	Exponential	10.000 sec
Query interarrival time	Exponential	20.000 sec
Query lead time	Exponential	40.000 sec
Packet transmission time	Exponential	0.200 sec
Number of nodes (N)	Constant	1000 nodes
Deployment area	Constant	3335m x 3335m
Node transmission range	Constant	>5000m (Isotropic)

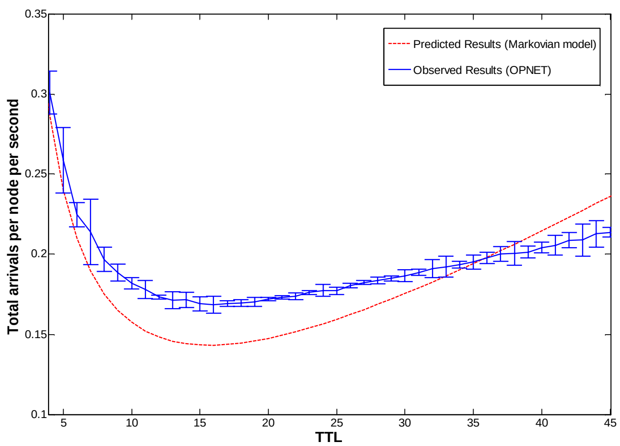

Figure 2: Total arrival rates, infinite transmission range

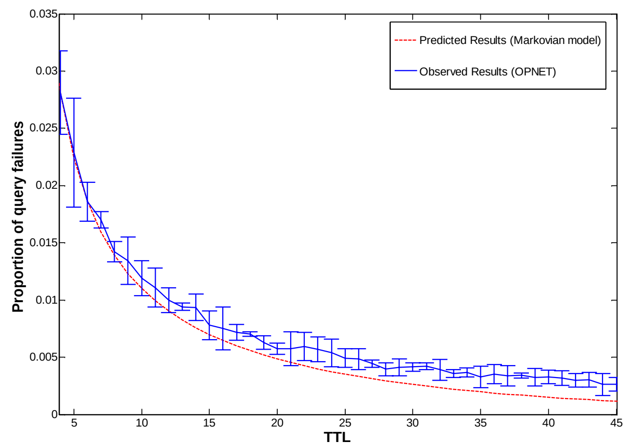

Figure 3: Total proportion of query failures, infinite transmission range

the analytic model. Lower replication levels require the network to support additional query transmissions to locate an informed node. As shown in Figure 3, the need for additional query transmissions causes a slightly higher query failure rate than predicted due to increased latency.

As TTL values increase beyond 36, the total arrival rate predicted by the analytic model is greater than that observed in the simulations. This occurs because few agents generated in the simulation can be replicated more than approximately 40 times as a consequence of the mean agent expiration time and the time required for each agent transmission, i.e., $\delta/\mu = 40$. Based on the network parameters, TTL values in excess of 40 create few additional replicates due to agent expiration; hence, total arrivals per node and the proportion of query failures remain relatively constant despite an increase in TTL. Although the analytic model predicts higher arrival rates and lower failure rates than observed, this is anticipated by means of the $\alpha_{\max}$ parameter (see [9]). The $\alpha_{\max}$ parameter recognizes that there is an upper limit to the proportion of the network that can be informed by agents as a consequence of network congestion and/or limited agent lifetimes. Momentarily ignoring the effects of congestion, the value of $\alpha_{\max}$ is approximately 40 for this network.

Despite the minor differences noted between the analytic and the simulation models, the analytic model requires a TTL value of 16 to minimize the total arrival rate of traffic to each node and, thus, to minimize the mean total node arrival rate of the network. Additionally, the predicted proportion of query failures is within 0.001 of the observed value when the TTL is 16 and does not exceed a differential greater than 0.0015 for TTL ≤ 45. Based on these results, we conclude the simulation model provides an accurate representation of the performance of a random walk search algorithm when both agents and queries are assigned expiration times. In the following subsections, we examine the effects of node transmission range and decreasing mean agent/query expiration lead times on performance.

V.A. Varying node transmission range

When a node's transmission range is limited such that its one-hop neighborhood consists of only a small subset of the total network nodes, the distribution of informed nodes is less likely to conform to the uniform distribution assumed by the analytic model. Therefore, it is expected that shorter node transmission ranges will require higher TTL values to achieve the minimum rate of arrivals, and the minimum rate of arrivals will be higher than that predicted by the analytic model. Additionally, the proportion of failed queries will increase due to the greater number of hops each query is expected to make prior to locating an informed node. Experiments using maximum effective node transmission ranges of 300m, 400m, 600m, and >5000m were conducting using the same parameters shown in Table 2. The results of these experiments are shown in Figures 4 and 5.

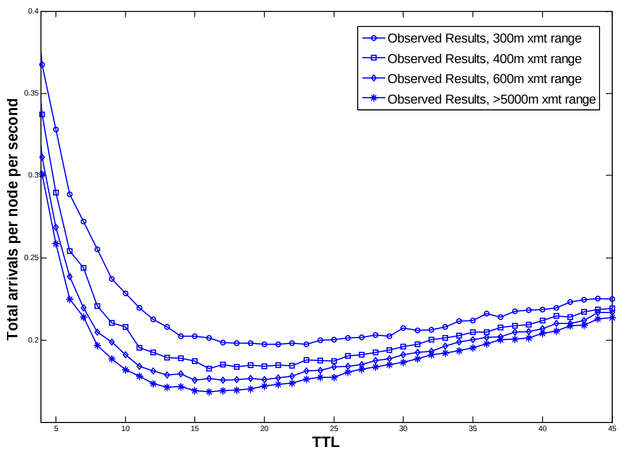

Figure 4: Mean total arrival rates, varying node transmission range

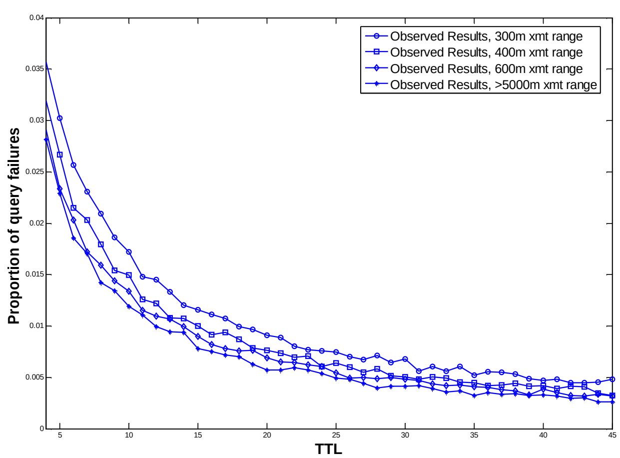

Figure 5: Proportion of query failures, varying node transmission range

As expected, the simulations confirm higher TTL values are required to achieve the minimum mean total arrival rate as the maximum effective node transmission range is decreased. These results are summarized in Table 3 (note that, for the 600m transmission range case, the results observed for TTL values of 15 and 16 are statistically identical). This implies a tradeoff between the energy required for transmission and the total number of transmissions required. While nodes with short transmission ranges expend less energy per transmission and generally experience reduced contention for medium access, the number of transmissions required per node per second is higher than that needed by nodes with longer transmission ranges. Additionally, nodes with longer transmission ranges have a smaller proportion of query failures for a given TTL value. Determination of the most energy efficient transmission range for the network (and, hence, the optimum TTL value) necessarily depends on the characteristics of the node hardware and the medium access control protocol. In general, however, the increase in total arrival rate observed when using reduced node transmission ranges is likely outweighed by the reduction in the energy required for transmission. Consequently, when considering energy efficiency, shorter node transmission ranges tend to result in less total energy expenditure despite an increase in the minimum observed total arrival rate.

V.B. Decreased query lifetimes

If query lifetimes are reduced in response to application requirements, preventing an unacceptably high proportion of query failures will necessitate decreasing the amount of time required by a query to locate an informed node. If the mean effective node transmission rate cannot be increased, the only remaining recourse is to increase the number of informed nodes in the network. To examine the effect of shorter query lifetimes on network performance, additional experiments were conducted using exponentially-distributed query lifetimes with means of 10, 20, 30, and 40 seconds. The results of these experiments are shown in Figures 6 and 7. The maximum node transmission range for these experiments is fixed at 400m.

As shown in Figure 6, total arrival rates are only marginally reduced by decreasing the mean query lifetime (a consequence attributed to reduced traffic due to query expiration). However, the resulting increase in the proportion of query failures necessitates higher TTL values to achieve the same proportion of query failures observed when queries have longer mean lifetimes. These results verify the intuitive link between query latency (i.e., the time required by the network to answer a query) and energy expenditure.

Table 3: Observed TTL values that minimize total arrival rates

Transmission Range	*Observed TTL Value*	*Observed Arrival Rate*
300m	20	177.576
400m	16	164.297
600m	15	157.828
>5000m	16	151.611

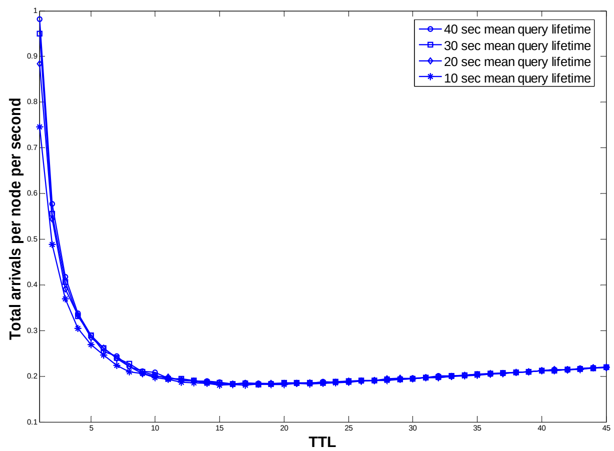

Figure 6: Total arrival rates, varying mean query lifetime

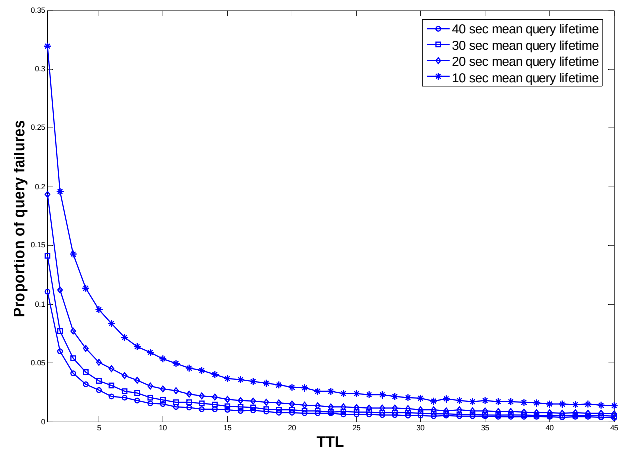

Figure 7: Proportion of query failures, varying mean query lifetime

VI. Conclusions

The choice of medium access control protocol affects the performance of the network. We assumed network traffic is very low; thus, the probability of a transmission collision is correspondingly small. This is an appropriate assumption in energy-constrained WSNs. Accordingly, the network's medium access control protocol is modeled by requiring each node to expend an exponentially-distributed amount of time to successfully transmit a query or agent to a neighboring node. Within the range of traffic intensities tested, the distribution of the random time required for a successful transmission is assumed to be constant. However, it is probable that the distribution of the time required by the medium access control protocol to facilitate a successful transmission may change as node densities and/or traffic levels increase.

The simulation experiments indicate the Markovian node model [9] provides a reasonable means to predict the performance of a random-walk search algorithm in large-population sensor networks. However, it may be possible to refine the Markovian model to better predict the performance of large networks of nodes with varying transmission ranges and mean agent/query lifetime distributions. Most importantly, the proportion of nodes informed by an agent (the parameter α in [9]) could be modified to reflect the fact that some agents will not exhaust their TTL counters prior to expiration. Consequently, the proportion of informed nodes is somewhat smaller than expected.

VII. Acknowledgments

The authors acknowledge, with gratitude, the helpful comments of the anonymous referees which have improved this paper.

References

[1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam and E. Cayirci, "A survey on sensor networks," *IEEE Communications Magazine*, vol. 40, pp. 102-114, 2002.

[2] T. Banka, G. Tandon and A. Jayasumana, "Zonal rumor routing for wireless sensor networks," in Proceedings of the International Conference on Information Technology: Coding and Computing–Volume II, pp. 562-567, 2005.

[3] P. Bellavista, A. Corradi and E. Magistretti, "Comparing and evaluating lightweight solutions for replica dissemination and retrieval in dense MANETs," in *Proceedings of the 10th*

IEEE Symposium on Computers and Communications, pp. 43-50, 2005.

[4] D. Braginsky and D. Estrin, "Rumor routing algorithm for sensor networks," in Proceedings of the 1st ACM International Workshop on Wireless Sensor Networks and Applications, pp. 22-31, 2002.

[5] C. F. Chou, J. J. Su and C. Y. Chen, "Straight line routing for wireless sensor networks," in *Proceedings of the 10th IEEE Symposium on Computers and Communications*, pp. 110-115, 2005.

[6] B. Krishnamachari and J. Ahn, "Optimizing data replication for expanding ring-based queries in wireless sensor networks," *USC Computer Engineering Technical Report CENG-05-14*, October 2005.

[7] D. Liu, X. Hu and X. Jia, "Energy Efficient Protocols for Information Dissemination in Wireless Sensor Networks," *Lecture Notes in Computer Science*, vol. 3842, pp. 176-185, 2006.

[8] C. R. Mann, R. O. Baldwin, J. P. Kharoufeh and B. E. Mullins, "A Trajectory-based Selective Broadcast Query Protocol for Large-Scale, High-Density Wireless Sensor Networks," *Telecommunication Systems: Modeling, Analysis, Design and Management*, Vol. 35, No. 1-2, pp. 67-86, 2007.

[9] C. R. Mann, R. O. Baldwin, J. P. Kharoufeh and B. E. Mullins, "A Queueing Approach to Optimal Resource Replication in Wireless Sensor Networks," accepted for publication in *Performance Evaluation*.

[10] R. Melamed, I. Keidar and Y. Barel, "Octopus: A fault-tolerant and efficient ad-hoc routing protocol," in *Proceedings of the 24th IEEE Symposium on Reliable Distributed Systems*, pp. 39-49, 2005.

[11] V. Rajendran, K. Obraczka and J. J. Garcia-Luna-Aceves, "Energy-Efficient, Collision-Free Medium Access Control for Wireless Sensor Networks," *Wireless Networks*, vol. 12, pp. 63-78, 2006.

[12] S. Ratnasamy, B. Karp, S. Shenker, D. Estrin, R. Govindan, L. Yin and F. Yu, "Data-Centric Storage in Sensornets with GHT, A Geographic Hash Table," *Mobile Networks and Applications*, vol. 8, pp. 427-442, 2003.

[13] S. Ratnasamy, B. Karp, L. Yin, F. Yu, D. Estrin, R. Govindan and S. Shenker, "GHT: A geographic hash table for data-centric storage," in *Proceedings of the 1st ACM International Workshop on Wireless Sensor Networks and Applications*, pp. 78-87, 2002.

[14] S. Shenker, S. Ratnasamy, B. Karp, R. Govindan and D. Estrin, "Data-Centric Storage in Sensornets," *Computer Communication Review*, vol. 33, pp. 137-142, 2003.

[15] I. Stojmenovic, "A scalable quorum based location update scheme for routing in ad hoc wireless networks," TR 99-09, SITE, University of Ottawa, 1999.

[16] J. B. Tchakarov and N. H. Vaidya, "Efficient content location in wireless ad hoc networks," in *Proceedings of the 2004 IEEE International Conference on Mobile Data Management*, pp. 74-85, 2004.