Infrastructure Tradeoffs for Sensor Networks*

Sameer Tilak†, Nael B. Abu-Ghazaleh† and Wendi Heinzelman‡

†Computer System Research Laboratory
Dept. of CS, Binghamton University
Binghamton, NY 13902–6000
{sameer,nael}@cs.binghamton.edu

‡Electrical and Computer Engineering
University of Rochester
Rochester, NY 14627–0126
wheinzel@ece.rochester.edu

ABSTRACT

In a sensor network, the infrastructure (in terms of the sensor capabilities, number of sensors, and deployment strategy) plays a significant role in determining the performance of the network. In this paper, we study the effect of infrastructure decisions on the performance of a sensor network. We study the effect of the infrastructure for two types of network delivery models (phenomenon driven and continuous) and different network protocols (DSR, DSDV and AODV). We show the performance both in terms of network efficiency as well as meeting the application accuracy and latency demands. By exploring the criteria for effective infrastructure configurations, we open the door for network optimizations that control the effective topology to better achieve the application requirements.

1. CATEGORIES & SUBJECT DESCRIPTORS

C.2.1 Network Architecture and Design, Wireless communication.

2. GENERAL TERMS

Performance, Design, Experimentation.

3. INTRODUCTION

Sensor networks represent a new paradigm for reliable environment monitoring and information collection. They hold the promise of revolutionizing sensing in a wide range of application domains because of their reliability, accuracy, flexibility, cost-effectiveness, and ease of deployment. Furthermore, in future smart environments, it is likely that sensor networks will play a key role in sensing, collecting, and disseminating information about the environment.

A sensor network is a tool for distributed sensing of one or more phenomena, and reporting the sensed data to one or more observers. As such, the performance of the network is best measured in terms of meeting the accuracy and delay requirements of the observer. Additional performance metrics include the life time of the network, cost of the sensors and their deployment, fault tolerance and scalability [34].

Conceptually, a sensor network is organized as a three layer system: (1) infrastructure: refers to the physical sensors (their physical characteristics and capabilities), the number of sensors and their deployment strategy (how/where they are deployed); (2) networking protocol: responsible for dissemination of the sensed data by creating and maintaining paths between the sensors and the observer(s); and (3) the application: responsible for translating the observer interests into specific network-level operations. Cross-cutting optimizations across the three levels are possible to improve the performance of the network.

Although there is a large body of work in building and networking sensors (a good bibliography of sensor network research can be found on this website [27]), these studies focus on optimizing the application and networking protocol to improve performance. In contrast, this paper considers the tradeoffs in the infrastructure design and their implications on performance and the design of the networking protocol. We also study the effect of biasing the deployment to reflect the phenomenon motion pattern on the performance of the network.

Intuitively, it appears that a denser infrastructure leads to a more effective sensor network because higher accuracy is likely and a larger aggregate amount of energy is available in the network. However, if not properly managed, a denser network will lead to a larger number of collisions and potentially to congestion in the network; this will increase latency and reduce energy efficiency. Moreover, the large number of samples reported by the sensors may exceed the accuracy requirements of the observer. Thus, simply increasing the reporting rate or the number of sensors may actually harm the performance of the network. We study this tradeoff using different application scenarios (phenomenon driven vs. continuous update data reporting) and for different infrastructure configurations.

One of the lessons learned from this study is that a form of congestion control is necessary to make sure that the reported samples do not exceed the capacity of the network. In addition, this control is necessary to optimize the lifetime of the

*This research was partially supported by NSF grant EIA-9911099

WSNA'02, September 28, 2002, Atlanta, Georgia, USA.

network while meeting the minimum accuracy requirements of the application. Thus, the congestion control must not only be based on the capacity of the network, but also on the accuracy level required at the observer. The traffic in a sensor network is different from conventional networks; it is a collective communication operation with redundancy. Thus, the network protocol has the flexibility of meeting the performance demands by controlling the reporting rate of the sensors, controlling the virtual topology of the network (by turning off some sensors for example), or optimizing the collective reduction communication operation (by fusing data along the way for example). We note that this application driven congestion control is different, and at a lower level, from proposals to incorporate application dependent processing and/or data aggregation within the network.

The remainder of this paper is organized as follows. In Section 4 we overview the role of the infrastructure and discuss the available deployment strategies. Section 5 overviews the modeling approach and the evaluation environment. In Section 6 we present the experimental study. Section 7 overviews some related work. Finally, Section 8 presents some concluding remarks.

4. INFRASTRUCTURE FEATURES

The infrastructure of a sensor network refers to the characteristics of the individual sensors, the number of sensors deployed, and the deployment strategy. We will discuss each of these in turn.

4.1 Sensors' Capabilities

A sensor typically consists of five components: sensing hardware, memory, battery, embedded processor, and transreceiver. These components affect the performance of the sensor and ultimately that of the network. For example, the accuracy of the sensing hardware or *transducer* will affect the accuracy of the sensing at the observer. Similarly, the size of the memory affects the buffering space at the sensors and the ability of the network to handle transient bursts in traffic. The battery size determines the amount of energy available at the sensor and affects the lifetime of the network. The capabilities of the embedded processor determine the level of optimization that is possible at the sensors without introducing excessive loss of power or intolerable levels of delay. Finally, the characteristics of the transreceiver determine the transmission range of the network and the capacity of the transmission channel. Improving the characteristics of any of these subsystems increases the cost, form factor or both for the sensor. Thus, within the available budget for the sensor network, the designer must decide whether to invest in a large number of inexpensive sensors, or a smaller number of expensive, higher quality ones.

4.2 Number of Sensors

Intuitively, for a given type of sensor, increasing the number of sensors deployed in the field should result in a better performing network with respect to the metrics identified earlier; otherwise, why pay the extra cost. Consider: (1) the accuracy of the sensing should improve since there are more sensors in a position to report on the phenomena; (2) the available energy within the network increases; and (3) the additional sensor density offers the potential for a better connected network with more efficient paths between the sensors and the observers. However, increasing the number of sensors in turn results in a higher number of sensors reporting their results per unit time. If this increased load exceeds the capacity of the network in terms of access to the shared wireless medium as well as congestion in intermediate nodes, increasing the number of active sensors may end up adversely affecting the performance of the network.

With respect to capacity, the problem can be viewed in terms of collision and congestion. To avoid collisions, sensors that are in the transmission range of each other should not transmit simultaneously. Consider sensors $1 \ldots M$, each with transmission range r, that are arranged in a chain. For any given sensor S_i, sensors located in the range $loc(i)-r$ and $loc(i)+r$ should not transmit at the same time. Research by Woo et al. [35] has addressed some of the issues with the collision problem, trying to improve upon existing MAC layers. To the best of our knowledge, congestion has not been addressed by past studies.

We consider a phenomenon driven reporting model where a sensor reports if it is in range of the phenomenon. Assume that we have N sensors out of which M sensors are in range of the phenomenon at a given time T. Assume that the M sensors are in interference range with each other (e.g., the transmission range is greater than or equal to the sensing range). Of the M reporting sensors, each sensor S_i will transmit data toward the observer with bit rate $b(S_i)$. The total data in transit from time T to $T+\delta$ where δ is the average latency can be expressed as

$$Data = \sum_{i=1}^{M} b(S_i) \quad (1)$$

If this value reaches a certain fraction of the channel capacity, congestion will occur [16]. If C_{total} is the total channel capacity then

$$\sum_{i=1}^{M} b(S_i) <= \alpha C_{total} \quad (2)$$

where α is a fraction of the capacity dictated by the self-interference that arises in multi-hop connections (α is typically around 0.25 [17]). Thus, the upper bound on the reporting rate is dictated by the channel capacity. On the other hand, application specific criteria such as the required accuracy places a lower bound on the reporting rate; the reporting rate should be high enough to satisfy the desired accuracy. At any point in time the number of active sensors should be such that the application specified accuracy requirements are met. If, in order to meet the accuracy requirements, $C_{application}$ is the required channel capacity then we have:

$$C_{application} <= \sum_{i=1}^{M} b(S_i) <= \alpha C_{total} \quad (3)$$

$$C_{application} <= \alpha C_{total} \quad (4)$$

to support the application requirements.

Note that not all sensors are equal in terms of accuracy: depending on the location, a specific sensor may have a higher quality data sample, or a combination of sensors may together provide a higher accuracy than another combination. However, we can qualitatively comment on the factors on which the number of active sensors depends. From a networking perspective, it depends on factors such as the geographic locations of the reporting sensors, buffer lengths, and packet processing times. From an application perspective, the value of information sensed by the sensor needs to be considered as well. If a sensor is providing some unique information about some feature of the phenomenon, then the application might require that sensor to report irrespective of the location of that sensor. Thus, application level information must be used in determining what sensors to report and when to meet the application performance metrics. We intend to pursue such protocols in the future.

4.3 Deployment Strategies

Finally, it is important to consider the deployment strategy for the sensors (e.g., their distribution within the phenomena field). We consider three deployment strategies: (1) random deployment – the sensors are "sprayed" with a uniform distribution within the field; (2) regular deployment – the sensors are placed with some regular geometric topology in the sensor field (for example, a grid); and (3) planned deployment – sensor deployment is planned (for example, biased to provide higher sensor density in areas where the phenomenon is concentrated). It is unclear whether regular deployment will offer advantages over uniformly distributed random deployment; if it does not, random deployment is preferable because of its low cost.

In the remainder of this paper, we will evaluate these infrastructure tradeoffs for two types of monitoring disciplines (phenomenon driven and continuous reporting), and different routing protocols. The evaluation environment and modeling approach are presented in the next section.

5. EVALUATION ENVIRONMENT

In order to model the complex relationships described above, we have developed an evaluation environment within the NS-2 simulator [2]. Contrary to most sensor network studies, we have made the phenomenon explicit and decoupled it from the sensor network organization. This allows us to study the effect of varying the design within the sensor network using scenarios that are independent of it. We model two types of phenomena: (1) discrete phenomena (for example, animals in a habitat monitoring application [5]); and (2) continuous phenomena (for example, the temperature in a temperature tracking application). For each of these types of phenomenon, the sensors wake up periodically according to some user defined schedule, take samples of the phenomenon and report their results if required by the application.

To model the transducer states, we have assumed that a transducer operates in two modes: sleep and active to model the low duty cycle necessary for power efficient operation. The transducer periodically wakes up and enters the active state to check the status of the phenomenon. For example, in the case of a discrete phenomenon, the transducer will wake up at a time interval of every δt seconds and check whether it can sense the phenomenon, such as an animal in its range. If it can then it will report the reading to the observer using its transceiver. This model can support both continuous as well as phenomenon driven model. For a phenomenon driven model, absence of a phenomenon (e.g., animal) during periodic active state will not result in reporting about the phenomenon and the sensor will go back to the sleep state.

The environment also decouples the three levels of the sensor network: infrastructure, protocol and application. The reasoning again is to provide a vehicle to allow comparison of "apples to apples"; we can study the effect of varying the design at each of these levels on the performance of the network under uniform assumptions. For example, we can study the effect of changing the network protocol for a given application and infrastructure. In this paper, we study the effect of varying the infrastructure on the performance of the network for different applications and network protocols.

In this work, we considered an application with a discrete phenomenon that moves around in a square grid (e.g., animal tracking) as well as an application with a continuous phenomenon that can always be sensed (e.g., temperature sensing). We also considered two application level scenarios: (1) continuous update: the sensors periodically report their local measurement to the observer; and (2) phenomenon driven: sensors report their measurements to the observer periodically, but only if they have data of interest to report (in this case, the discrete phenomenon is within detection range). Other scenarios can be easily constructed; for example, scenarios with multiple phenomena or multiple observers can be directly generated.

We are interested in application-level performance; conventional network performance metrics such as throughput are of secondary interest. We consider the following performance metrics.

1. Accuracy: The accuracy of a measurement at a sensor is specific to the physical transducer and the nature of the phenomenon. In general, we assume that the measurement has a tolerance that increases with the distance between the sensor and the phenomenon. At the observer, it is likely that multiple samples will be received from the different sensors. These samples must be combined intelligently to produce a more accurate estimate of the location of the phenomenon. It is possible to bias the estimate towards sensors with higher confidence (closer to the phenomenon) and towards more recent samples.

2. Latency: Latency refers to the delays in obtaining the samples at the observer due to network congestion, the duty cycle of the sensors, or intelligent filtering of sampled data. For real time sensing applications, delays in reporting the state of the phenomenon leads to a loss in accuracy. For the purposes of this study, we report only the packet latency within the network.

3. Energy efficiency and fault tolerance: the energy efficiency of the network may be measured in different ways. For now, we report the energy expenditure within the network.

4. Goodput: Goodput is the ratio of the total number of packets received by the observer to the total number of packets sent by all the sensors over the simulation time.

5. Scalability is also of interest. While we do not investigate scalability directly, efficient data reporting and reducing network load is conducive to scalability.

Table 1: Parameters used in the simulation studies.

Simulation area	$800x800\ m^2$
Transmission range	$250\ m$
Startup-Energy	$10000\ J$
Discrete phenomenon sensing range	$200\ m$
Phenomenon speed	random between $1-2\ m/s$
MAC Protocol	802.11
Bandwidth	$2\ Mbps$
Transmit Power	$0.660\ W$
Receive Power	$0.395\ W$

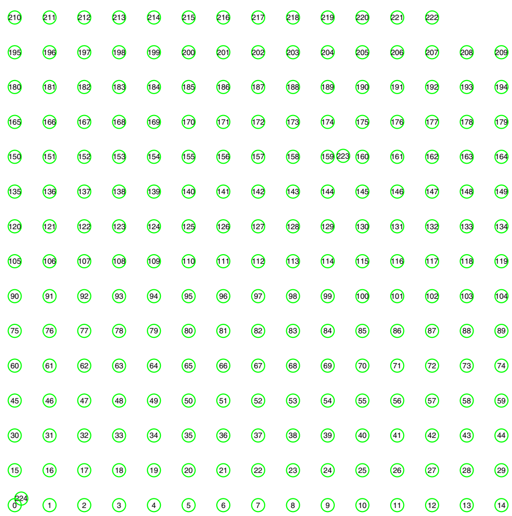

Figure 1: 15x15 grid deployment.

6. EXPERIMENTAL STUDY

We considered a scenario where a discrete phenomenon is being tracked by sensors placed in a square grid of dimensions 800 meters by 800 meters. We assumed that each sensor data sample has a uniformly distributed tolerance of +/- 5% of the actual distance between the sensor and the phenomenon. In the phenomenon driven scenarios, only the sensors within a discrete phenomenon sensing range report their estimate of the location of the phenomenon to the observer. The packet size was fixed at 100 bytes unless specified. We used the energy model from the Directed Diffusion sensor network protocol study [13]. The buffer space available at each sensor is of size 5 packets; a larger buffer size will enable the network to withstand a higher level of transient congestion but will not help with sustained overloading of the network. From its initial position, the phenomenon walks towards a random destination with a speed randomly chosen between 1 m/s and 2 m/s.

The parameters used in our simulations are summarized in Table 1. Changing these parameters will have an effect on the capacity of the network and the offered load, but due to space limitations this effect is not pursued.

Figures 1, 2 and 3 show a sample of the topologies that were used for the experiments. Figure 1 shows a 15x15 grid topology, which includes 223 sensor nodes, 1 observer node and 1 phenomenon node. Figure 2 shows a network with 100 sensors randomly distributed. Figure 3 shows a biased network, where the designer has some idea about the phenomenon mobility. In this case if the designer knows that the phenomenon moves in general in upper right corner (hot-spot), then sensors can be deployed with non-uniform density so that more sensors are located in the hot-spot region.

We first establish the basic infrastructure configuration tradeoffs using the following parameters. We used Dynamic Source Routing (DSR) [14] as the networking protocol and explore both grid deployment and random deployment of the sensors. Finally, we explore biasing the deployment pattern to match the phenomenon motion pattern. Each simulation was run for 50 seconds, and every point represents the average of three different random seeds. Unless stated otherwise, the data delivery model was phenomenon driven.

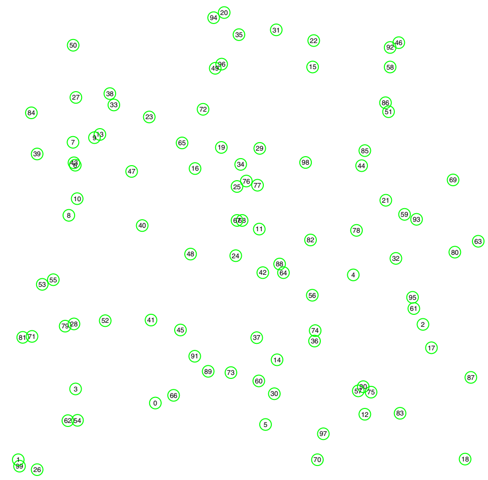

Figure 2: Random network of 100 sensors.

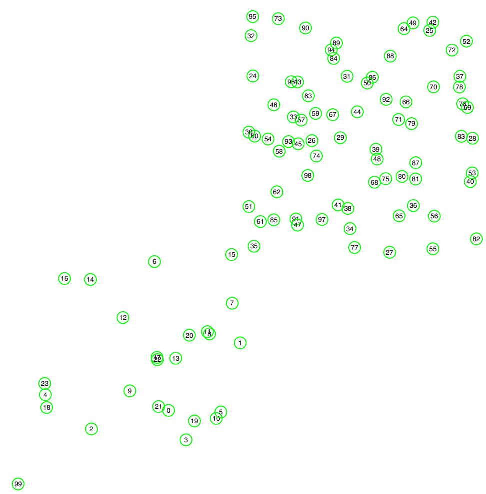

Figure 3: Biased network of 100 sensors.

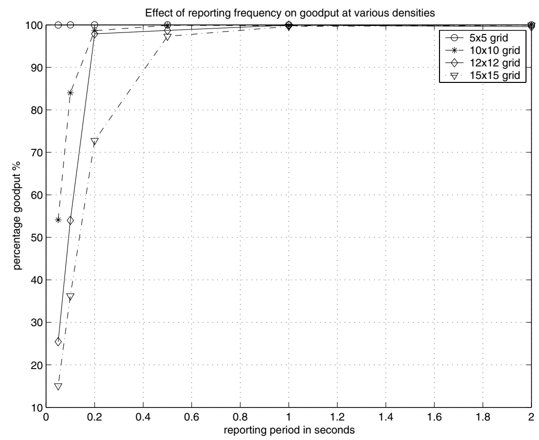

Figure 4: Goodput as a function of network density and sensor reporting period (grid deployment).

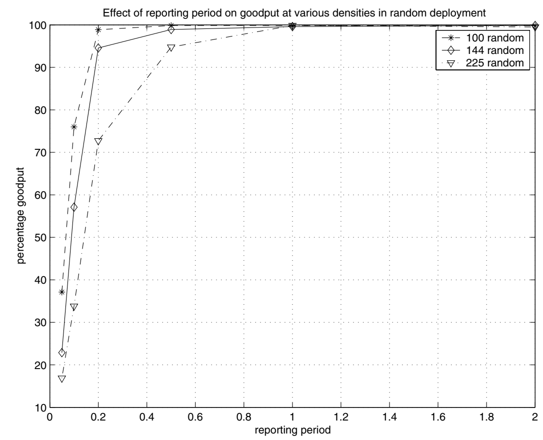

Figure 6: Goodput as a function of network density and sensor reporting period (random deployment).

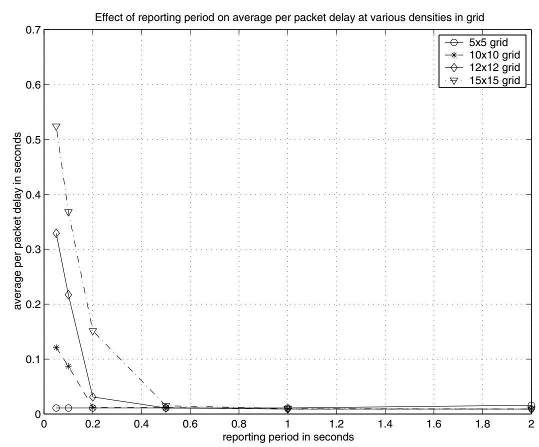

Figure 5: Delay as a function of network density and sensor reporting period (grid deployment).

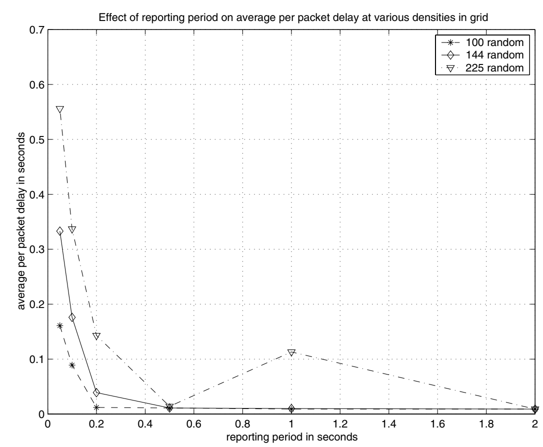

Figure 7: Delay as a function of network density and sensor reporting period (random deployment).

6.1 Basic Infrastructure Tradeoffs

6.1.1 Goodput and Delay Study

In the first set of experiments, we study the effect of increasing the sensor density on the efficiency of the network. Figure 4 shows the goodput of the network as a function of the reporting period for several levels of network density. The deployment strategy was regular; sensors were placed in square grids with the stated number of sensors per side. We first note that as the data rate increases (reporting period decreases), the goodput drops when the rate exceeds the capacity of the network and sensed packets start to be dropped. It is interesting to note that the drop in goodput is more pronounced for the denser networks. This is due to the larger number of sensors close to the phenomenon effectively increasing the offered load to the network, resulting in more collisions and a higher number of packets dropped due to congestion. This effect is corroborated by the packet latency results (Figure 5): the latency increases with the data rate as well as the density of the network.

We repeated these experiments for random deployment, keeping the same number of sensors as each of the studied grids. The results for goodput (Figure 6) and delay (Figure 7) do not show appreciable differences in comparison to grid deployment. Note that we do not consider the scenario with 25 sensors, as was done in the grid case, because the network was too sparse to maintain connectivity with random deployment.

6.1.2 Accuracy Study

In terms of application performance, we measured the accuracy of the tracking of the phenomenon position. More specifically, the observer generates an estimate of the phenomenon location based on the samples it receives from the sensors. We measured the error in these samples in the following way. We discretized time into small slots and averaged the samples received in each slot. We then compared this average to the actual location of the phenomenon at that time. The error is the square root of the sum of the square of the difference between the estimated location and the actual location averaged over

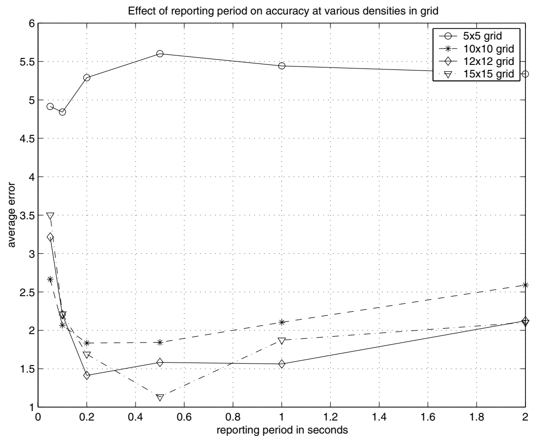

Figure 8: Error as a function of network density and sensor reporting period (grid deployment).

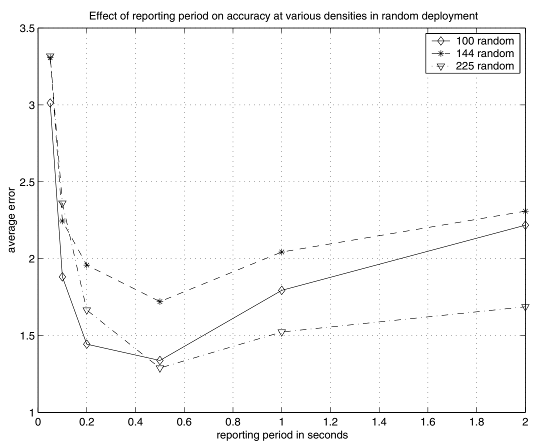

Figure 9: Error as a function of network density and sensor reporting period (random deployment).

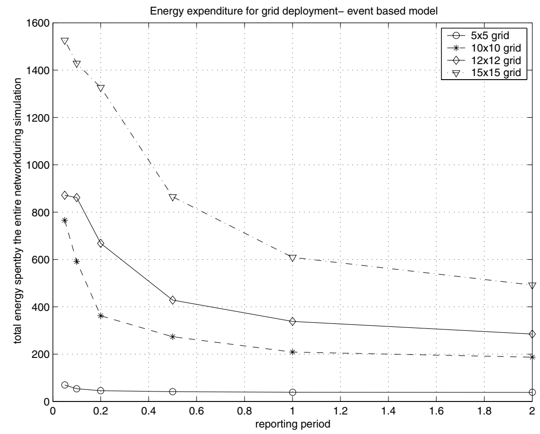

Figure 10: Energy depletion as a function of network density and sensor reporting period (grid deployment).

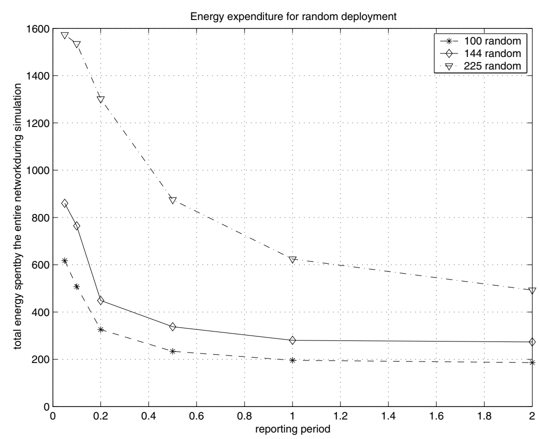

Figure 11: Energy depletion as a function of network density and sensor reporting period (random deployment).

the number of slots in the simulation. More specifically,

$$E = \frac{\sqrt{\sum_{i=1}^{n}(S(i) - A(i))^2}}{n} \tag{5}$$

where $S(i)$ is the sensed value in time slot i, $A(i)$ is the actual value at time slot i, and n is the number of slots in the duration of the simulation. This is a proof of concept approach to calculating error; any statistical measure for correlating the measured value against the actual value will suffice.

Figure 8 shows the average error for the grid deployment strategy under different densities and for different reporting periods. At high reporting rates, network capacity is exceeded, as was observed in the previous graphs. Because of the latency in the receipt of the samples and the loss of many samples, the error value is high. On the other hand, if the reporting frequency is low, not enough samples are obtained and the average error rises. With sparse networks (example 5x5 grid) the error is higher when the network is not saturated because the number of sensors in a position to measure the phenomenon and the average distance between a sensor and the phenomenon increases. For such scenarios, the error is minimized with a higher reporting frequency; the additional samples reduce the error and the network is slower to saturate because there are fewer sensors competing for the shared air space. With random deployment (Figure 9) the same pattern can be observed.

6.1.3 Energy-Efficiency Study

The energy depletion in the network is shown in Figure 10 and Figure 11 for the grid sensor deployment and random deployment respectively. The energy depletion is a function of the reporting rate as well as the density of the network. Recall that the density of the network in the phenomenon driven scenario correlates with the number of nodes that report their data. However, as suggested by the goodput results, a large portion of this energy is wasted when the capacity of the network is exceeded. Moreover, the additional cost incurred to buy more sensors will not be rewarded by a higher lifetime for the network because the depletion rate also increases. In fact, when we consider the normalized energy expenditure per sensor (as Figure 12 shows for grid deployment) the average sensor gets depleted more quickly with higher density. Thus, the lifetime

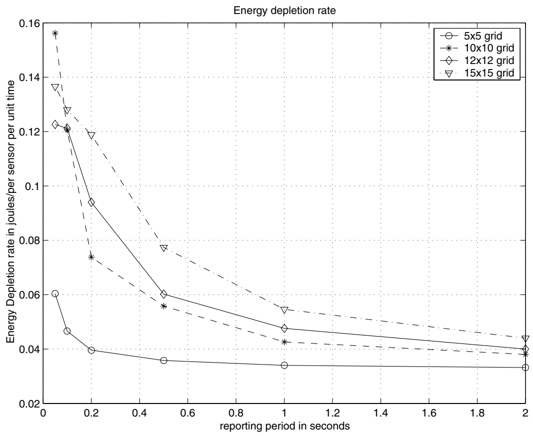

Figure 12: Energy depletion per sensor as a function of network density and sensor reporting period (grid deployment).

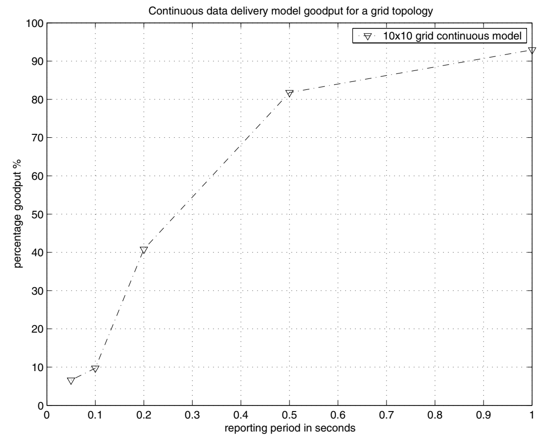

Figure 13: Goodput as a function of sensor reporting period for continuous update traffic and a 10x10 grid.

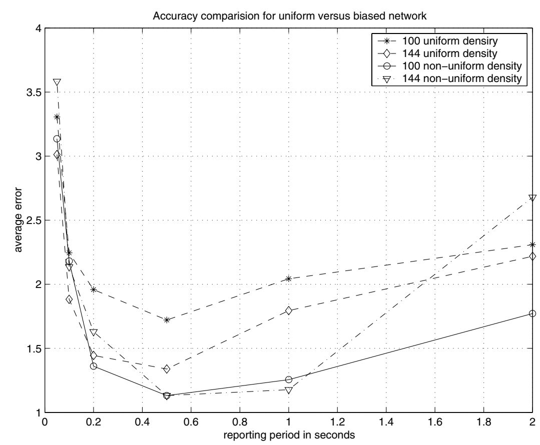

Figure 14: Average error comparison – controlled vs. grid deployment.

of the network likely drops with increased density even though we start with a much higher total available power in the network! Accordingly, there is a need for intelligent management of the infrastructure from an energy perspective as well.

To summarize, in agreement with intuition, increasing the network density can result in higher accuracy, but only if the sensing traffic is kept below the network capacity. This is an expanded form of the congestion control requirement for regular computer networks; due to the redundant collective communication nature of sensor network traffic, the network has the ability of controlling what data gets reported to meet the observer requirements. It is likely that the observer is satisfied with less than the optimal achievable accuracy. Thus, the network protocol must control the available infrastructure and the reporting discipline to meet the accuracy requirements while minimizing the energy expenditure. The sensor network must converge on a good accuracy to reporting pattern/energy solution. This may be achieved, for example, by deciding to turn off some sensors, by adapting the reporting frequency, or by fusing sampled data within the network.

6.2 Continuous Update Reporting Model

For the continuous update reporting model (all sensors report continuously), the offered load was significantly higher than the phenomenon driven model. Energy depletion results (not shown) displayed this effect. As can be seen in Figure 13, the goodput values using continuous update reporting were significantly lower than for phenomenon driven traffic. Error is not directly comparable across the two scenario types.

6.3 Controlled Deployment

In this experiment, we study the effect of biasing the deployment to the phenomenon's motion pattern. In this experiment, the phenomenon was restricted to move in the right half of the 800 meters by 800 meters field. Furthermore, the deployment of the sensors was skewed to reflect this fact: the density of the sensors in the left half was kept fixed (and low) while the density of the sensors in the right half was increased. Figure 14 shows the accuracy using biased deployment vs. grid deployment. As can be seen from the figure, the desired effect of increasing the accuracy was achieved (the average error is lower in the biased deployment case). In fact, with biased deployment, a network of 100 sensors performs better than one with 144 sensors that are deployed in a grid. However, note that with aggressive reporting the network saturates faster under biased deployment, since the average number of nodes within reporting range of the phenomenon increases. This effect was also seen in the goodput results (not shown). The increased accuracy comes at the cost of extra energy depletion as well (results not shown). These results again argue that the network protocol should carefully manage the infrastructure.

6.4 Other Supporting Experiments

In order to further investigate the congestion problem, we conducted experiments with different packet sizes, in order to study the effect on goodput and accuracy with a change in bandwidth, where an increase in packet size corresponds to a reduction in bandwidth. The results shown in Figures 15 and 16 show that congestion becomes a serious problem with low bandwidth (packets with large size), as goodput drops dramatically and the average error increases appreciably.

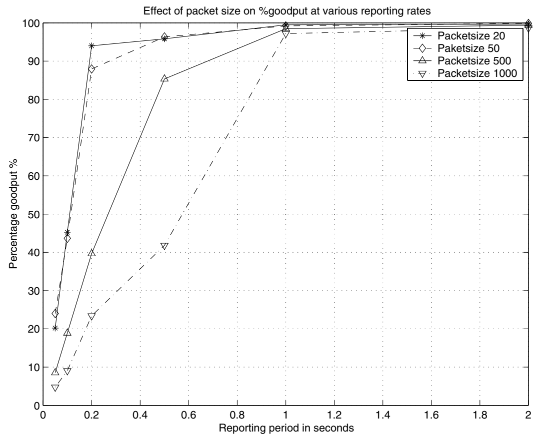

Figure 15: Goodput as a function of packet size and sensor reporting period for a 15x15 grid.

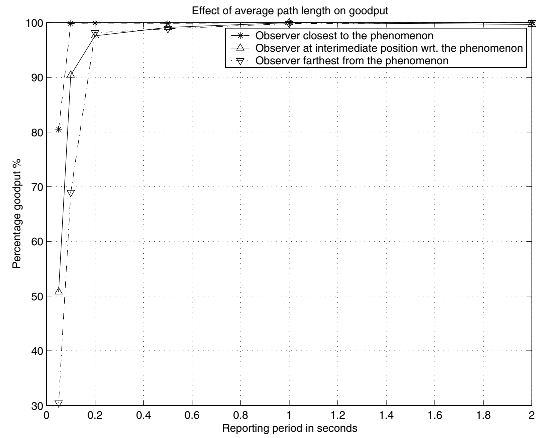

Figure 17: Goodput as a function of variation in path length and sensor reporting period for a 15x15 grid.

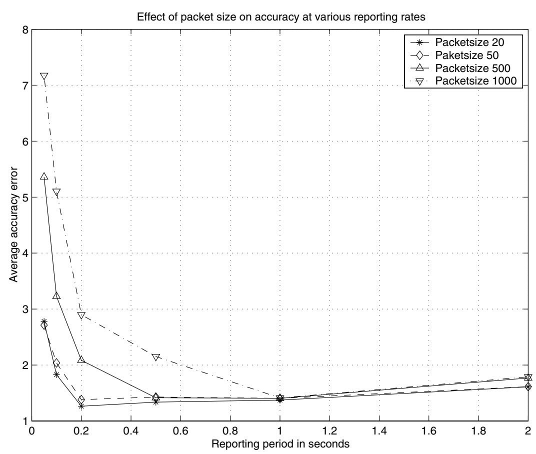

Figure 16: Error as a function of packet size and sensor reporting period for a 15x15 grid.

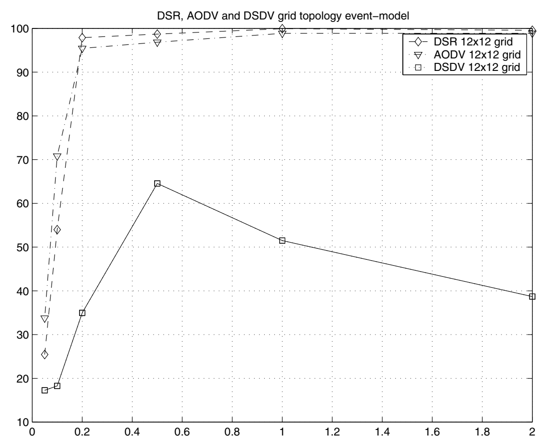

Figure 18: The effects of the routing protocol on goodput for a 12x12 grid.

We also conducted experiments with the observer at different relative distances with respect to the phenomenon, moving closer to the phenomenon in each case. Results shown in Figure 17 indicate that with an increase in path length, goodput decreases from 85% when the observer is closest to the phenomenon to 30% when the observer is furthest from the phenomenon. This agrees with the previous results in that the further the observer is from the phenomenon, the longer the average path from the sensors to the observer and the more data must be transmitted throughout the network, increasing network load and causing congestion.

Although the results we have presented and the conclusions we have drawn should be not be heavily impacted by the network/routing protocol (ignoring in-network processing), we investigated the effect of using other routing protocols. We investigated AODV [22] which, like DSR is reactive routing protocol. Also, we studied DSDV [23], which is a proactive protocol. Figure 18. indicates AODV performed almost identically to DSR, while DSDV was considerably poorer in all cases.

7. RELATED WORK

Because of the unique requirements on sensor network nodes, several groups have proposed architectures for sensor nodes [1, 6, 8, 15, 24, 25, 26, 31, 35]. On top of these architectures, several studies targeted the development of power-efficient medium access protocols (e.g., [30, 32, 35]). Networking and data dissemination issues have also received considerable interest. Due to the data-centric nature of sensor networks, researchers proposed alternative addressing schemes that take advantage of this fact [9, 12]. A number of routing/data aggregation approaches were also proposed [3, 10, 11, 13, 18]. A number of studies have explored implementing services for sensor networks, including positioning mechanisms [4, 21, 28], time synchronization [7] and energy scans [36]. Other studies considered specific sensor network applications and their implication on protocol design [5, 29, 33, 34].

Meguerdichian et al. define the problem of exposure in sensor networks [19] and propose localized algorithms to address it [20]. The exposure problem is the problem of determining whether a sensor network can keep track of a phenomenon that moves within the observation field. Depending on the sensor density/deployment, there could be blind spots in the observation field. Clearly, exposure is influenced by the deployment configuration of the sensors and is related to our work.

8. CONCLUDING REMARKS

In this paper, we investigated the effect of infrastructure tradeoffs on the performance of a sensor network. First, we systematically increased the deployed sensor density and the required reporting rate and observed the performance of the network. When the offered load from the sensors to the network exceeded the capacity of the network, the performance dropped according to both network and application level metrics. Thus, by simply deploying more sensors, we may end up harming the performance of the network. This argues for intelligent management of the infrastructure by the network protocol: a form of congestion avoidance is needed that is significantly different from congestion avoidance in traditional data networks. In particular, the network protocol must balance the offered load to the network against the required accuracy at the observer.

The task of the sensor network may be viewed as a redundant collective communication process from the sensors to the observer. It is redundant in that multiple sensors may report correlated information or information with an accuracy level (e.g., reporting rate) higher than that required by the application. Thus, the congestion avoidance mechanism must converge on a reporting rate/discipline that is just sufficient to meet the performance requirements at the observer. The networking protocol may accomplish this by reducing the reporting rate per sensor, turning some sensors off and/or fusing information to optimize the collective communication operation.

We also investigated the effect of different deployment strategies on the performance of the network. We discovered no appreciable differences between grid-type deployment and random deployment for the scenarios we considered. However, biasing the infrastructure density to the phenomenon movement pattern resulted in significantly higher accuracy. This is an example of using application level information to better architect the infrastructure.

9. REFERENCES

[1] ASADA, G., DONG, M., LIN, T., NEWBERG, F., POTTIE, G., AND KAISER, W. Wireless integrated network sensors: Low power systems on a chip. In *European Solid State Circuits Conference* (Oct. 1998).

[2] BERKELEY/LNBL/ISI, U. The ns-2 network simulator with the cmu mobility extensions, 2002. `http://www.isi.edu/nsnam/ns/`.

[3] BHATNAGAR, S., DEB, B., AND NATH, B. Service differentiation in sensor networks. In *Proc. 4th International Symposium on Wireless Personal Multimedia Communications* (Sept. 2001).

[4] BULUSU, N., HEIDEMANN, J., AND ESTRIN, D. GPS-less low cost outdoor localization for very small devices. *IEEE Personal Communications Magazine* (Oct. 2000), 28–34.

[5] CERPA, A., ELSON, J., ESTRIN, D., GIROD, L., HAMILTON, M., AND ZHAO, J. Habitat monitoring: Application driver for wireless communications technology. In *Proc. ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean* (Apr. 2001).

[6] CHANDRAKASAN, A., AMIRTHARAJAH, A., CHO, S., GOODMAN, J., KONDURI, G., KULIK, J., RABINER, W., AND WANG, A. Design considerations for distributed microsensor systems. In *Proc. of the IEEE 1999 Custom Integrated Circuits Conference (CICC'99)* (May 1999).

[7] ELSON, J., AND ESTRIN, D. Time synchronization for wireless sensor networks. In *Proc. Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing* (Sept. 2001).

[8] ESTRIN, D., GIROD, L., POTTIE, G., AND SRIVASTAVA, M. Instrumenting the world with wireless sensor networks. In *Proceeings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2001)* (May 2001).

[9] HEIDEMANN, J., SILVA, F., INTANAGONWIWAT, C., GOVINDAN, R., ESTRIN, D., AND GANESAN, D. Building efficient wireless networks with low-level naming. In *Proc. 2001 Symposium on Operating Systems Principles* (Oct. 2001), pp. 146–159.

[10] HEINZELMAN, W., CHANDRAKASAN, A., AND BALAKRISHNAN, H. Energy-efficient routing protocols for wireless microsensor networks. In *Proc. 33rd Hawaii International Conference on System Sciences (HICSS '00)* (Jan. 2000).

[11] HEINZELMAN, W., KULIK, J., AND BALAKRISHNAN, H. Adaptive Protocols for Information Dissemination in Wireless Sensor Networks. In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '99)* (Aug. 1999), pp. 174–185.

[12] IMIELINSKI, T., AND GOEL, S. DataSpaces: Querying and monitoring deeply networked colections in physical space. In *Proc. MobiDE 1999* (1999), pp. 44–51.

[13] INTANAGONWIWAT, C., GOVINDAN, R., AND ESTRIN, D. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In *Proc. 6th ACM International Conference on Mobile Computing and Networking (Mobicom'00)* (Aug. 2000).

[14] JOHNSON, D., MALTZ, D., HU, Y., AND JETCHEVA, J. The dynamic source routing protocol for mobile ad hoc networks. Internet Draft, Internet Engineering Task Force, Mar. 2001. http://www.ietf.org/internet-drafts/draft-ietf-manet-dsr-05.txt.

[15] KAHN, J., KATZ, R., AND PISTER, K. Next century challenges: Mobile networking for 'smart dust'. In *Proceedings of the Fifth Annual International Conference on Mobile Computing and Networking* (July 1999), pp. 271–278.

[16] LI, J., BLAKE, C., DE COUTO, D., LEE, H., AND MORRIS, R. Capacity of ad hoc wireless networks. In *Proceedings of the 2001 ACM Mobile Computing and Networking Conference (Mobicom'01)* (July 2001), pp. 61–69.

[17] LI, J., JANNOTTI, J., COUTO, D., KARGER, D., AND MORRIS, R. A scalable location service for geographic ad hoc routing. In *Proceedings of the International Conference on Mobile Computing and Networks (MobiCom'00)* (Aug. 2000), pp. 120–130.

[18] LINDSEY, S., RAGHAVENDRA, C., AND SIVALINGAM, K. Data gathering in sensor networks using energey-delay metric. In *Proc. International Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing* (Apr. 2001).

[19] MEGUERDICHIAN, S., KOUSHANFAR, F., QU, G., AND POTKONJAK, M. Exposure in wireless ad-hoc sensor networks. In *The Seventh Annual International Conference on Mobile Computing and Networking 2001* (July 2001), pp. 139–150.

[20] MEGUERDICHIAN, S., SLIJEPCEVIC, S., KARAYAN, V., AND POTKONJAK, M. Localized algorithms in wireless ad hoc networks: Location discovery and sensor exposure. In *Proc. MobiHoc 2001* (2001).

[21] NICULESCU, D., AND NATH, B. Ad hoc positioning system (aps). In *Proc. GLOBECOM 2001* (2001).

[22] PERKINS, C., ROYER, E., AND DAS, S. Ad hoc on-demand distance vector (aodv) routing. Internet Draft, Internet Engineering Task Force, Mar. 2001. http://www.ietf.org/internet-drafts/draft-ietf-manet-aodv-08.txt.

[23] PERKINS, C. E., AND BHAGWAT, P. Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. *ACM Computer Communications Review 24*, 4 (Oct. 1994), 234–244. SIGCOMM '94 Symposium.

[24] POTTIE, G., AND KAISER, W. Wireless integrated network sensors. *Communications of the ACM 43*, 5 (May 2000), 551–558.

[25] RABAEY, J., AMMER, M., DA SILVA, J., PATEL, D., AND ROUNDY, S. PicoRadio supports ad hoc ultra-low power wireless networking. *IEEE Computer 33*, 7 (July 2000).

[26] Rockwell science center sensor network project, 2002. (Available on the web at: `http://www.rsc.rockwell.com/wireless_systems/sensorware`).

[27] RUTGERS UNIVERSITY, C. S. D. Wireless sensor networks bibliography website, 2002. `http://www.cs.rutgers.edu/~mini`.

[28] SAVVIDES, A., HAN, C., AND STRIVASTAVA, M. Dynamic fine-grained localization in ad hoc networks of sensors. In *Proc. 7th ACM International Conference on Mobile Computing and Networking* (July 2001).

[29] SCHWEIBERT, L., GUPTA, S., AND WEINMANN, J. Research challenges in wireless networks of biomedical sensors. In *The Seventh Annual International Conference on Mobile Computing and Networking 2001* (July 2001).

[30] SHIH, E., CHO, S., ICKES, N., MIN, R., SINHA, A., WANG, A., AND CHANDRAKASAN, A. Physical layer driven protocol and algorithm design for energy-efficient wireless sensor networks. In *Proceedings of the Seventh Annual International Conference on Mobile Computing and Networking* (July 2001), pp. 272–287.

[31] SINHA, A., AND CHANDRAKASAN, A. Dynamic power management in wireless sensor networks. *IEEE Design and Test of Computers 18*, 2 (Mar. 2001).

[32] SOHRABI, K., GAO, J., AILAWADHI, V., AND POTTIE, G. Protocols for self-organization of a wireless sensor architecture. *IEEE Personal Communications 7*, 5 (Oct. 2000), 16–27.

[33] SRIVASTAVA, M., MUNTZ, R., AND POTKONJAK, M. Smart Kindergarten: Sensor-based Wireless Networks for Smart Developmental Problem-solving Environments. In *The 7th Annual International Conference on Mobile Computing and Networking 2001* (July 2001), pp. 132–138.

[34] TILAK, S., ABU-GHAZALEH, N., AND HEINZELMAN, W. A taxonomy of wireless micro-sensor network communication models. *ACM Mobile Computing and Communication Review* (Apr. 2002).

[35] WOO, A., AND CULLER, D. A transmission control scheme for media access in sensor networks. In *Proceedings of the Seventh Annual International Conference on Mobile Computing and Networking* (July 2001).

[36] ZHAO, Y., GOVINDAN, R., AND ESTRIN, D. Residual energy scans for monitoring wireless sensor networks. In *IEEE Wireless Communications and Networking Conference (WCNC'02)* (Mar. 2002).