A Taxonomy of Wireless Micro-Sensor Network Models *

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

In future smart environments, wireless sensor networks will play a key role in sensing, collecting, and disseminating information about environmental phenomena. Sensing applications represent a new paradigm for network operation, one that has different goals from more traditional wireless networks. This paper examines this emerging field to classify wireless micro-sensor networks according to different communication functions, data delivery models, and network dynamics. This taxonomy will aid in defining appropriate communication infrastructures for different sensor network application sub-spaces, allowing network designers to choose the protocol architecture that best matches the goals of their application. In addition, this taxonomy will enable new sensor network models to be defined for use in further research in this area.

I. Introduction

Advances in hardware and wireless network technologies have placed us at the doorstep of a new era where small wireless devices will provide access to information anytime, anywhere as well as actively participate in creating smart environments. One of the applications of smart spaces is *sensor networks*, networks that are formed when a set of small untethered sensor devices that are deployed in an ad hoc fashion cooperate on sensing a physical phenomenon. Sensor networks 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.

To motivate the challenges in designing sensor networks, consider the following scenarios: sensors are rapidly deployed in a remote inhospitable area for a surveillance application; sensors are used to analyze the motion of a tornado; sensors are deployed in a forest for fire detection; sensors are attached to taxi cabs in a large metropolitan area to study the traffic conditions and plan routes effectively; and smart Kindergarten [1] where sensor networks are deployed to create a developmental problem-solving environment for early childhood education.

Clearly, there is a wide range of applications for sensor networks with differing requirements. We believe that a better understanding of micro-sensor network requirements as well as the underlying differences between micro-sensor applications is needed to assist designers. To this end, in this paper we attempt to classify wireless micro-sensor networks. In particular, we classify the aspects of wireless micro-sensor networks that we believe are most relevant to communication. We examine the characteristics and goals of typical micro-sensor networks as well as the different types of communication that are required to achieve these goals. We compare different data delivery models and network dynamics to create a taxonomy of wireless micro-sensor network communication. We believe that this taxonomy will aid network designers in making better decisions regarding the organization of the network, the network protocol and information dissemination models. Furthermore, it will aid in developing realistic sensor network models and benchmarks for use in future sensor network research.

The remainder of this paper is organized as follows. Section II presents some basic definitions and an overview of the characteristics of sensor networks. Section III overviews performance metrics of interest for sensor networks. In Section IV, we describe sensor network architectures. Section V classifies the communication models present in sensor networks and makes the distinction between application and infrastructure related communication. Section VI classifies the data delivery models. In Section VII, the network organization and dynamics are classified. Section VIII presents case studies of existing sensor network protocols, showing how they fit into the taxonomy described in this paper. Finally, Section IX presents a summary and some concluding remarks.

II. Sensor Network Characteristics

In this paper, we use the following terminology:

- *Sensor*: The device that implements the physical

sensing of environmental phenomena and reporting of measurements (through wireless communication). Typically, it consists of five components– sensing hardware, memory, battery, embedded processor, and trans-receiver.

- *Observer*: The end user interested in obtaining information disseminated by the sensor network about the phenomenon. The observer may indicate *interests* (or queries) to the network and receive responses to these queries. Multiple observers may exist in a sensor network.

- *Phenomenon*: The entity of interest to the observer that is being sensed and potentially analyzed/filtered by the sensor network. Multiple phenomena may be under observation concurrently in the same network.

In a sensing application, the observer is interested in monitoring the behavior of the phenomenon under some specified performance requirements (e.g., accuracy or delay). In a typical sensor network, the individual sensors sample local values (*measurements*) and disseminate information as needed to other sensors and eventually to the observer. The measurements taken by the sensors are discrete samples of the physical phenomenon subject to individual sensor measurement accuracy as well as location with respect to the phenomenon.

Sensor networks share many of the challenges of traditional wireless networks, including limited energy available to each node and bandwidth-limited, error-prone channels. However, communication in sensor networks differs from communication in other types of networks in that it is typically not end-to-end [2]. More specifically, the function of the network is to report information regarding the phenomenon to the observer who is not necessarily aware of the sensor network infrastructure and the individual sensors as an end-point of communication. Furthermore, energy is typically more limited in sensor networks than in other wireless networks because of the nature of the sensing devices and the difficulty in recharging their batteries. Studies in the past have shown that 3000 instructions could be executed for the same energy cost as sending a bit 100m by radio [3]. This indicates that the tradeoff between communication and computation in sensor networks should be resolved in favor of computation. In addition, studies have shown that current commercial radio transceivers, for example those used by Bluetooth devices, are unsuitable for sensor network applications because of their energy requirements [4]. Thus sensor networks impose challenges in hardware design as well as in communication protocols.

III. Performance Metrics

We propose using the following metrics to evaluate sensor network protocols.

- *Energy efficiency/system lifetime*. As sensor nodes are battery-operated, protocols must be energy-efficient to maximize system lifetime. System lifetime can be measured by generic parameters such as the time until half of the nodes die or by application-directed metrics, such as when the network stops providing the application with the desired information about the phenomena.

- *Latency*. The observer is interested in knowing about the phenomena within a given delay. The precise semantics of latency are application dependent.

- *Accuracy*. Obtaining accurate information is the primary objective of the observer, where accuracy is determined by the given application. There is a trade-off between accuracy, latency and energy efficiency. The given infrastructure should be adaptive so that the application obtains the desired accuracy and delay with minimal energy expenditure. For example, the application can either request more frequent data dissemination from the same sensor nodes or it can direct data dissemination from more sensor nodes with the same frequency.

- *Fault-tolerance*: Sensors may fail due to surrounding physical conditions or when their energy runs out. It may be difficult to replace existing sensors; the network must be fault-tolerant such that non-catastrophic failures are hidden from the application. Fault-tolerance may be achieved through data replication (e.g., the SPIN protocol [5]). However data replication itself requires energy; there is a trade-off between data replication and energy-efficiency. We suggest that the data replication should be application-specific. The data which have higher priority according to the application might be replicated for fault tolerance and the other data might not be.

- *Scalability*: Scalability for sensor networks is also a critical factor. For large-scale networks, it is likely that localizing interactions through hierarchy and aggregation will be critical for ensuring scalability.

IV. Sensor Network Architecture

A sensor network is a tool for measuring and relaying information about the phenomenon to the observer within the desired performance bound and deployment cost. As such, the organization of the network may be viewed as follows:

1. Infrastructure: The infrastructure consists of the sensors and their current deployment status. More

specifically, the infrastructure is influenced by the characteristics of the sensors (e.g., sensing accuracy, memory size, battery life, transmission range) and deployment strategy (e.g., sensor density, sensor location, sensor mobility).

2. Network Protocol: The network protocol is responsible for creating paths and accomplishing communication between the sensors and the observer(s).

3. Application/Observer: The observer(s) interests in the phenomenon are queries from the observer(s) about the phenomenon as approximated by the distributed data that the sensors are capable of sensing. These queries could be static (the sensors are preprogrammed to report data according to a specific pattern) or dynamic. The network may participate in synthesizing the query (for example, by filtering some sensor data or fusing several measurements into one value); we consider such intelligence to be part of the translation process between observer interests and low-level implementation.

In this work, we focus on classifying issues that influence the second level: the network protocol. We discuss the other two levels only with regard to issues that influence communication. Thus, we do not address the difficult problem of translation between the observer query and the specific low-level interests. This translation could be done by the application software at the observer and/or the sensor nodes, or directly by a human observer. Similarly, we do not discuss the engineering of the infrastructure.

We also note that there is a significant opportunity for optimizations that cut across the three organizational levels. For example, Bhatnagar et al. discuss supporting QoS for sensor networks [6]. More specifically, they discuss discriminating among the type of data that the sensors are reporting and preferentially treating high priority data (for example, by giving it priority in forwarding and using redundancy to increase the chance of its reception). This is an example of an optimization where application-level knowledge provides hints to the network protocol. As another example, consider the case where the deployment of the sensors is chosen to mirror the expected motion pattern of the phenomenon or the interests of the observer. Such a deployment strategy incorporates application knowledge in the infrastructure design.

The network protocol in a sensor network is responsible for supporting all communication, both among sensor nodes as well as between the sensor nodes and the observer(s). The performance of the protocol will be highly influenced by the network dynamics, as well as by the specific data delivery model employed. In order to determine how the network protocol behaves for different scenarios, it is important to classify these features. In the following sections, we classify the different types of communication required in a sensor network and then look at the possible data delivery models and network dynamics.

V. Communication Models

There are multiple ways for a sensor network to achieve its accuracy and delay requirements; a well designed network meets these requirements while optimizing the sensor energy usage and providing fault tolerance. By studying the communication patterns systematically, the network designer will be able to choose the infrastructure and communication protocol that provide the best combination of performance, robustness, efficiency and deployment cost.

Conceptually, communication within a sensor network can be classified into two categories: *application* and *infrastructure*. The network protocol must support both these types of communication. Application communication relates to the transfer of sensed data (or information obtained from it) with the goal of informing the observer about the phenomena. Within application communication, there are two models: cooperative and non-cooperative. Under the cooperative sensor model, sensors communicate with other sensors to realize the observer interest. This communication is beyond the relay function needed for routing. For example, in a clustering protocol a cluster-head and the sensor nodes communicate with each other for information dissemination related to the actual phenomenon. In-network data processing [5, 7, 8] is an example of co-operative sensors. Non-cooperative sensors do not cooperate for information dissemination.

Infrastructure communication refers to the communication needed to configure, maintain and optimize operation. More specifically, because of the ad hoc nature of sensor networks, sensors must be able to discover paths to other sensors of interest to them and to the observer regardless of sensor mobility or failure. Thus, infrastructure communication is needed to keep the network functional, ensure robust operation in dynamic environments, as well as optimize overall performance. We note that such infrastructure communication is highly influenced by the application interests since the network must reconfigure itself to best satisfy these interests. As infrastructure communication represents the overhead of the protocol, it is important to minimize this communication while ensuring that the network can support efficient application communication.

In sensor networks, an initial phase of infrastructure communication is needed to set up the network. Furthermore, if the sensors are energy-constrained, there will be additional communication for reconfiguration. Similarly, if the sensors are mobile or the observer interests

dynamic, additional communication is needed for path discovery/reconfiguration. For example, in a clustering protocol, infrastructure communication is required for the formation of clusters and cluster-head selection; under mobility or sensor failure, this communication must be repeated (periodically or upon detecting failure). Finally, infrastructure communication is used for network optimization. Consider the Frisbee model, where the set of active sensors follows a moving phenomenon to optimize energy efficiency [9]. In this case, the sensors wake up other sensors in the network using infrastructure communication.

Sensor networks require both application and infrastructure communication. The amount of required communication is highly influenced by the networking protocol used. Application communication is optimized by reporting measurements at the minimal rate that will satisfy the accuracy and delay requirements given sensor abilities and the quality of the paths between the sensors and the observer. The infrastructure communication is generated by the networking protocol in response to application requests or events in the network. Investing in infrastructure communication can reduce application traffic and optimize overall network operation.

VI. Data Delivery Models

Ideally, the observer interest is specified in terms of the phenomenon, allowing the observer to be oblivious to the underlying sensor network infrastructure and protocol. The query is implemented as one or more specific low-level interests (e.g., requesting a specific sensor to report a specific measurement at some specific interval). Sensor networks can be classified in terms of the data delivery required by the application (observer) interest as: *continuous*, *event-driven*, *observer-initiated* and *hybrid*. These models govern the generation of the application traffic. In the continuous model, the sensors communicate their data continuously at a prespecified rate. The authors in [8] showed that clustering is most efficient for static networks where data is continuously transmitted. For dynamic sensor networks, depending upon the degree of mobility, clustering may be applicable as well. In the event-driven data model the sensors report information only if an event of interest occurs. In this case, the observer is interested only in the occurrence of a specific phenomenon or set of phenomena. In the observer-initiated (or request-reply) model, the sensors only report their results in response to an explicit request from the observer (either directly, or indirectly through other sensors). Finally, the three approaches can coexist in the same network; we refer to this model as the hybrid model.

Thus far, we have only discussed data delivery from the application perspective, and not the actual flow of data packets between the sensors and the observer; this is a routing problem subject to the network protocol. For any of the above-mentioned models, we can classify the routing approach as: flooding (broadcast-based), unicast, or multicast/other. Using a flooding approach, sensors broadcast their information to their neighbors, who rebroadcast this data until it reaches the observer. This approach incurs high overhead but is immune to dynamic changes in the topology of the network. Research has been conducted on techniques such as data aggregation that can be used to reduce the overhead of the broadcast [2, 5, 8]. Alternatively, the sensors can either communicate to the observer directly (possibly using a multi-hop routing protocol) or communicate with a cluster-head using one-to-one unicast. Finally, in a multicast approach, sensors form application-directed groups and use multicast to communicate among group members. The observer could communicate with any member of the group to obtain the desired data. A major advantage of flooding or broadcast is the lack of a complex network layer protocol for routing, address and location management; existing sensor network efforts have mostly relied on this approach (e.g., [2, 5]). However, the overhead of a broadcasting approach may be prohibitive.

It is likely that the interaction between the data delivery model from the application and the routing model employed by the network protocol will significantly impact the performance of the network. Consider a scenario where a sensor network is deployed for intrusion detection. In this case, the data delivery model is event driven – the event being an intruder entering the area. If the network level routing model is flooding based, in such a case physically co-located sensors will in general sense the intruder at the same time and try to send data to the observer simultaneously. These concurrent communications in the neighborhood will contend with each other for the use of the communication medium, raising: (1) the probability of loss of critical information; and (2) the latency in event reporting. A similar problem is studied by Woo and Culler [10].

VII. Network Dynamics Models

A sensor network forms a path between the phenomenon and the observer. The goal of the sensor network protocol is to create and maintain this path (or multiple paths) under dynamic conditions while meeting the application requirements of low energy, low latency, high accuracy, and fault tolerance. Without loss of generality, this discussion assumes a single observer. Multiple observers can be supported as multiple instances of a single observer. More sophisticated protocols could also capitalize on the presence of multiple observers to merge related interests

and/or optimize communication.

The problem of setting up paths for information dissemination is similar to the problem of routing in ad hoc networks [11]. However, there are a few critical differences, including: (i) the sensors are not generally addressed individually; rather, the interest is in the set of sensors that are *in a position to contribute to the active observer interests*. The sensors could be addressed by attributes of the sensor (e.g., their capabilities) and/or the phenomenon (e.g., the sensors close to a lion in a habitat monitoring scenario). The mapping between the observer interest and a set of sensors is influenced by the network dynamics and the application; and (ii) nodes along the path can take an active role in the information dissemination and processing. In this respect, sensor networks are similar to Active Networks [12] whereas ad hoc networks are traditional "passive" networks.

There are several approaches to construct and maintain a path between observer and phenomenon. These will differ depending on the network dynamics, which we classify as: *static sensor networks* and *mobile sensor networks*. We focus on mobility because it is the most common source of dynamic conditions; other sources include sensor failure and changes in observer interests.

Static Sensor Networks

In static sensor networks, there is no motion among communicating sensors, the observer and the phenomenon. An example is a group of sensors spread for temperature sensing. For these types of sensor networks, previous studies have shown that localized algorithms can be used in an effective way [2, 8]. The sensors in localized algorithms communicate with nodes in their locality. An elected node relays a summary of the local observations to the observer, perhaps through one on more levels of hierarchy. Such algorithms extend the lifetime of the sensor network because they trade-off local computation for communication [8]. In this type of network, sensor nodes require an initial set-up infrastructure communication to create the path between the observer and the sensors with the remaining traffic exclusively application communication[1].

Dynamic Sensor Networks

In dynamic sensor networks, either the sensors themselves, the observer, or the phenomenon are mobile. Whenever any of the sensors associated with the current path from the observer to the phenomenon moves, the path may fail. In this case, either the observer or the concerned

---

[1]Note that if energy is limited among the nodes, the network will require infrastructure communication to maintain a path between the observer and the phenomenon as nodes run out of energy.

sensor must take the initiative to rebuild a new path. During initial set-up, the observer can build multiple paths between itself and the phenomenon and cache them, choosing the one that is the most beneficial at that time as the current path. If the path fails, another of the cached paths can be used. If all the cached paths are invalid then the observer must rebuild new paths. This observer-initiated approach is a *reactive* approach, where path recovery action is only taken after observing a broken path.

Another model for rebuilding new paths from the observer to the phenomenon is a sensor-initiated approach. In a sensor-initiated path recovery procedure, path recovery is initiated by a sensor that is currently part of the logical path between the observer and the phenomenon and is planning to move out of range. The sensor might perform some local patching procedure to build a new path by broadcasting a *participation request* for a given logical flow to all its neighboring sensors. Any one of the neighboring sensors can send a *participation reply* message to the given initiator sensor indicating willingness to participate and become a part of the requested path. If none of the neighboring sensors respond, the sensor can default to sending a path invalidation request to the observer so that the observer can start building the path. This is similar to soft hand-off in traditional Mobile IP based networks [13]. This sensor-initiated approach is a *proactive* approach where path recovery operations are begun in anticipation of a future broken path.

Dynamic sensor networks can be further classified by considering the motion of the components. This motion is important from the communications perspective since the degree and type of communication is dependent on network dynamics. We believe that each of the following require different infrastructures, data delivery models, and protocols:

- *Mobile observer*. In this case the observer is mobile with respect to the sensors and phenomena. An example of this paradigm is sensors deployed in an inhospitable area for environment monitoring. For example, a plane might fly over a field periodically to collect information from a sensor network. Thus the observer, in the plane, is moving relative to the sensors and phenomena on the ground.

- *Mobile sensors*. In this case, the sensors are moving with respect to each other and the observer. For example, consider traffic monitoring implemented by attaching sensors to taxis. As the taxis move, the attached sensors continuously communicate with each other about their own observations of the traffic conditions. If the sensors are co-operative, the communication paradigm imposes additional constraints such as detecting the link layer addresses of the neighbors and constructing localization and in-

formation dissemination structures. From previous work [2], we know that the overhead of maintaining a globally unique sensor ID in a hierarchical fashion like an IP address is expensive and not needed. Instead, these sensors should communicate only with their neighbors with the link layer MAC address. In such networks, the above-mentioned proactive algorithm with local patching for repairing a path can be used so that the information about the phenomenon is always available to the observer regardless of the mobility of the individual sensors.

- *Mobile phenomena.* In this case, the phenomenon itself is moving. A typical example of this paradigm is sensors deployed for animal detection. In this case the infrastructure level communication should be event-driven. Depending on the density of the phenomena, it will be inefficient if all the sensor nodes are active all the time. Only the sensors in the vicinity of the mobile phenomenon need to be active. The number of active sensors in the vicinity of the phenomenon can be determined by application specific goals such as accuracy, latency, and energy efficiency. A model that is well-suited to this case is the Frisbee model [9].

It is important to note that the effect of mobility in sensor networks is fundamentally different than that in traditional wireless networks. Mobility in ad hoc networks has been addressed from the point of view of mobility of one or more of the communicating nodes during communication. However, since the sensors themselves are of no interest to the observer, their mobility is not necessarily of interest; rather, the sensor network must adapt its operation to continue to reflect the observer interests in the presence of mobility. Thus, the mobility of the sensing nodes themselves should be handled in a different way than for ad hoc networks; for example, a node that is moving away from a phenomenon may choose to hand-off the responsibility of monitoring to a closer node as it drifts away.

VIII. Case Studies and Related Work

In this section we consider several existing protocols for sensor networks and analyze them in the context of our taxonomy.

Ad hoc routing protocols may be used as the network protocol for sensor networks. However, such protocols will generally not be good candidates for sensor networks because of the following reasons: (i) sensors have low battery power and low available memory; (ii) the routing table size scales with the network size; (iii) these networks are designed for end to end communication and react inappropriately to mobility; (iv) their addressing requirements may be inappropriate for sensor networks [7]; and (v) ad hoc routing protocols do not support cooperative dissemination. More specifically, multihop routing protocols such as DSR [14] and AODV [15] support the creation and maintenance of paths to route packets from source to destination. Sensor network studies have shown that application specific in-network data processing is essential to maximize the performance of the sensor-network [7, 8]. As ad hoc routing protocols do not inherently support data aggregation or fusion, they will not perform well in sensor network applications.

From an operational perspective, it is interesting to see the parallel between ad hoc routing protocol and the sensor network taxonomy. It appears that proactive protocols such as DSDV [16] are more appropriate to continuous data delivery since they proactively maintain paths throughout the network. In fact, one can think of the link state update function in these protocols as a form of continuous data delivery. Similarly, reactive protocols such as DSR [14] appear better suited for event-driven or query based information dissemination. In addition, a similar distinction can be made based on the network dynamics: the more dynamic the network, the better the reactive appraoches.

LEACH is an energy efficient protocol for sensor networks designed for sensor networks with continuous data delivery mechanism and no mobility [8]. LEACH uses a clustering architecture where member nodes send their data to the local cluster-head. Cluster-heads aggregate the data from each sensor and then send this information to the observer node. LEACH uses rotation of the cluster-head in order to evenly distribute the energy load. Once clusters are formed, cluster members uses TDMA to communicate with the cluster-head. Thus LEACH is suitable for networks where every node has data to send at regular intervals. However, it needs to be extended for event-driven models as well as for mobile sensors.

Directed Diffusion (DD) is a data-centric protocol, where nodes are not addressed by their addresses but by the data they sense [2]. Data is named by attribute-value pairs. In directed diffusion the interest is expressed by observer nodes in term of a query which diffuses through the network using local interactions. Once a sensor node that satisfies the query (source node) is reached, that node starts transmitting data to the sink node, again using local interactions. The absence of a notion of a global id (e.g., IP address) makes directed diffusion efficient for networks with mobility as well. Directed diffusion is applicable for event-driven and query-driven networks as defined in our taxonomy. The localized interactions allow the protocol to scale to large networks; DD scales as a function of the

number of active interests present in the network.

The Publish/Subscribe model has been proposed for mobile networks by Huang and Gracia-Molina [17]. In this model, communication is typically anonymous, inherently asynchronous and multicasting in nature. From an application perspective, it also appears that the publish/subscribe model captures the relationship between the observer and phenomenon for some applications. More specifically, this model has desirable properties from the perspective of sensor networks; since the communication is not end-to-end, anonymous communication with application-specific multicast group formation is a viable approach. From an implementation perspective, asynchronous communication helps to preserve energy and increase the life-time of the network.

Ratnasamy et al. [18] present an alternative classification of sensor networks based on the data dissemination model. They propose that data dissemination can be done in at least three ways: (1) external storage - pass all the data to the observer and let them process this information; (2) local storage - information about the event is stored locally by the sensors; and (3) data-centric storage - data is stored by name and queries are directed by that name to the corresponding sensor. Clearly, the choice of the model will influence the communication patterns within the network. We view this as an application level decision.

IX. Conclusion

The overall communication behavior in a wireless micro-sensor network is application driven. We believe that it is useful to decouple the application communication used for information dissemination from the infrastructure communication used to configure and optimize the network. This separation will aid network designers in selecting the appropriate sensor network architecture that will best match the characteristics of the communication traffic of a given application. This will allow the network protocol to achieve the application-specific goals of energy-efficiency, low latency, and high accuracy in the sensing application. We also believe that a sensor-initiated proactive path recovery approach with local patching will be beneficial in efficient information dissemination in wireless micro-sensor networks.

We plan to study the behavior of various communication protocols for the different application sub-spaces described in this paper. This will be done through analysis and simulation to determine the advantages and disadvantages of existing approaches, such as DSR (Dynamic Source Routing) [19], directed diffusion [2], and LEACH [8]. We hope that the taxonomy we have presented will be helpful in designing and evaluating future network protocols for wireless micro-sensor networks.

Often, it is possible to implement a sensor network for a specific phenomenon in a number of different ways. Consider the problem of monitoring a tornado. One option would be to fly airplanes to sense the tornado (mobile phenomenon; mobile sensors; continuous data delivery). Another would be to have a sensor grid statically placed on the ground and report data as the tornado passes through (mobile phenomenon; static sensors; continuous data delivery). Yet another would be to release lightweight sensors into the tornado (static phenomenon; mobile sensors; continuous data delivery). The primary concern here is the ability of the sensor network to report the desired level of accuracy under latency constraints within an acceptable deployment cost. The accuracy is a function of the sensing technology of the sensors and their distance from the phenomenon. However, since the performance is measured at the observer end, it is also a function of the performance of the communication model. We hope that this taxonomy will assist in developing relevant simulation models to enable empirical study of the performance of the different sensor network organizations and assist in making design and deployment decisions.

References

[1] Mani Srivastava, Richard Muntz, and Miodrag Potkonjak, "Smart Kindergarten: Sensor-based Wireless Networks for Smart Developmental Problem-solving Environments," in *The Seventh Annual International Conference on Mobile Computing and Networking 2001*, July 2001, pp. 132 – 138.

[2] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks," *Proc. 4th ACM International Conference on Mobile Computing and Networking (Mobicom'98)*, Aug. 2000.

[3] J. Pottie and W. J. Kaiser, "Embedding the internet wireless integrated network sensors," *Communicaitons of the ACM*, vol. 43, no. 5, pp. 51–58, May 2000.

[4] Eugene Shih, Seong-Hwan Cho, Nathan Ickes, Rex Min, Amit Sinha, Alice Wang, and Anantha Chandrakasan, "Physical layer driven protocol and algorithm design for energy-efficient wireless sensor networks," in *The Seventh Annual International Conference on Mobile Computing and Networking 2001*, July 2001, pp. 272 – 287.

[5] W. Heinzelman, J. Kulik, and H. Balakrishnan, "Adaptive Protocols for Information Dissemination in Wireless Sensor Networks," in *Proceedings of*

the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (Mobi-Com '99), Aug. 1999, pp. 174–185.

[6] Sudeept Bhatnagar, Budhaditya Deb, and Badri Nath, "Service Differentiation in Sensor Networks," in *The Fourth International Symposium on Wireless Personal Multimedia Communications, September 2001.*, Sept. 2001.

[7] J. Heidemann, F. Silva, C. Intanagonwiwat, R. Govindan, D. Estrin, and D. Ganesan, "Building efficient wireless sensor networks with low-level naming.," in *Proceedings of the Eighteenth ACM Symp. on Operating Systems Principles [21]*, Oct. 2001, pp. 146–159.

[8] W. Heinzelman, *Application-Specific Protocol Architectures for Wireless Networks*, Ph.D. thesis, Massachusetts Institute of Technology, 2000.

[9] A. Cerpa, J. Elson, D. Estrin, L. Girod, M. Hamilton, and J. Zhao, "Habitat Monitoring: Application Driver for Wireless Communications Technology," in *Proc. ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean*, Apr. 2001.

[10] Alec Woo and David Culler., "A Transmission Control Scheme for Media Access in Sensor Networks.," in *Mobicom 2001*, 2001.

[11] Internet Engineering Task Force MANET Working Group, "Mobile ad hoc networks (MANET) charter," http://www.ietf.org/html.charters/manet-charter.html.

[12] D. Tennenhouse, J. Smith, W. Sincoskie, D. Wetherall, and G. Minden, "A survey of active network research," *IEEE Communications Magazine*, vol. 35, no. 1, pp. 80–86, Jan. 1997.

[13] "IETF MobileIP Working Group Internet Draft," http://www.ietf.org/rfc/rfc2002.txt, 1996.

[14] D. Johnson, D. Maltz, Y-C. Hu, and J. Jetcheva, "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] C. Perkins, E. Royer, and S. Das, "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.

[16] Charles Perkins and Pravin Bhagwat, "Highly dynamic destination-sequenced distance-vector routing (dsdv) for mobile computers," in *ACM SIGCOMM'94 Conference on Communications Architectures, Protocols and Applications*, Aug. 1994, pp. 234–244.

[17] Hector Garcia-Molina Yongqiang Huang, "Publish/Subscribe in a mobile enviroment," in *International Workshop on Data Engineering for Wireless and Mobile Access*, 2001, pp. 27–34.

[18] Sylvia Ratnasamy, Deborah Estrin, Ramesh Govindan, Brad Karp, Scott Shenker, Li Yin, and Fang Yu, "Data-centric storage in Sensornets," in *Submitted for review to SIGCOMM '02*, Feb. 2002.

[19] "IETF MANET Working Group Internet Draft– Dynamic Source Routing Protocol for Mobile Ad Hoc Networks," http://www.ietf.org/internet-drafts/draft-ietf-manet-dsr-05.txt, 2001.

Author Biographies

Sameer Tilak is pursuing an MS degree in computer science at the Computer Science Department at SUNY–Binghamton. He received his B.E. degree in computer engineering from Pune Institute of Computer Technology (Poona University, India) in 1999. His research interests lie in protocols for wireless and mobile networks, particularly sensor networks, and operating systems.

Nael Abu-Ghazaleh is an Assistant Professor at the Computer Science Department at the State University of New York, Binghamton. He received his MS and PhD degrees in computer engineering from the University of Cincinnati in 1994 and 1997 respectively. His research interests are in mobile computing and networking, computer architecture, and parallel processing.

Wendi Heinzelman is an Assistant Professor in the Department of Electrical and Computer Engineering at the University of Rochester. She received the B.S. degree in Electrical Engineering from Cornell University in 1995 and the M.S. and Ph.D. degrees in Electrical Engineering and Computer Science from MIT in 1997 and 2000 respectively. Her current research interests are in ad-hoc wireless protocol architectures, sensor networks, and multimedia communication. She is a member of Sigma Xi, the IEEE, and the ACM.