Medium Access Control Issues in Sensor Networks

Muneeb Ali
LUMS, CS Department
Opposite Sector U, DHA
Lahore - 54792, Pakistan
muneeb@sics.se

Umar Saif
MIT, CSAIL
The Stata Center, 32G-780,
Cambridge, MA 02139
umar@mit.edu

Adam Dunkels
SICS
Box 1263
SE-164 29 Kista, Sweden
adam@sics.se

Thiemo Voigt
SICS
Box 1263
SE-164 29 Kista, Sweden
thiemo@sics.se

Kay Römer
ETH Zurich
Department of CS
CH-8092 Zürich, Switzerland
roemer@inf.ethz.ch

Koen Langendoen
TU Delft
Mekelweg 4, 2628 CD Delft
The Netherlands
k.g.langendoen@ewi.tudelft.nl

Joseph Polastre
Moteiv Corporation
55 Hawthorne St, Suite 550
San Francisco, CA 94105
joe@moteiv.com

Zartash Afzal Uzmi
LUMS, CS Department
Opposite Sector U, DHA
Lahore - 54792, Pakistan
zartash@lums.edu.pk

ABSTRACT

Medium access control for wireless sensor networks has been a very active research area for the past couple of years. The sensor networks literature presents an alphabet soup of medium access control protocols with almost all of the works focusing only on energy efficiency. There is much more innovative work to be done at the MAC layer, but current efforts are not addressing the hard unsolved problems. Majority of the works appearing in the literature are "least publishable incremental improvements" over the popular S-MAC [1] protocol. In this paper we present research directions for future medium access research. We identify some open issues and discuss possible solutions.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless communication; C.2.6 [**Internetworking**]: Standards; I.6.5 [**Model Development**]: Modeling methodologies

General Terms

Design, Experimentation, Standardization

Keywords

Wireless Sensor Networks, Medium Access Control

1. INTRODUCTION

Sensor networks (sensor-nets) have emerged as one of the dominant technology trends of this decade (2000-2010). Sensors and actuators, wireless communications, and embedded computing are not new concepts but it is the recent low-cost large-scale integration of computation, communication, and sensing into "wireless sensor networks" that has captured the attention of many researchers. Sensor-nets enable observing the physical world at a granularity level which was unperceived before [2]. Applications of sensor-nets encompass a variety of disciplines and domains, limited only by the imagination of the application developer.

When thinking about future research directions for (wireless) networks it is important to consider which kind of devices would be connected to the network. In the coming years the most common type of devices on the network would be embedded processors, such as sensors and actuators (i.e. sensor-nets), and improving battery life will be more important than improving performance [3]. Energy-consumption by computing chips is falling sharply per unit computation (Moore's Law) whereas energy consumed by radios is determined by laws of physics. Thus, the wireless interface will be the *primary consumer* of energy in any device that combines computation and radios [3].

The Medium Access Control (MAC) layer sits directly on top of the Physical layer and controls the radio. MAC protocols for sensor-nets focus on *energy efficiency* (single most important goal) instead of meeting traditional goals for wireless MAC design such as fairness, delay, and bandwidth utilization [4]. These protocols tradeoff performance (fairness, delay, bandwidth utilization) for energy cost. Main sources of energy wastage at the MAC layer are collisions, idle listening, overhearing, and control packet overhead [4].

Unlike the 802.11 WLAN cards where the MAC is usually included as part of the chipset, in sensor-nets the MAC designer has absolute control on the MAC layer design. This absolute control, and the fact that the wireless interface is the *primary consumer* of battery in energy constrained sensor-nets, has made "medium access for sensor networks" a very active research area. However, recent studies on MAC protocols for sensor-nets observe that there is no clear trend indicating that medium access for sensor-nets is converging towards a unique best solution [5].

There is a tendency of "re-inventing the wheel" in the MAC for sensor-net area with majority of the works claiming a few percent better performance over the popular S-MAC [1] protocol. Instead of focusing only on energy efficiency (which almost all recent works do), it's time to address other unsolved problems at the MAC layer. In this paper we present research directions for design, simulation, and experimentation of MAC protocols for sensor-nets. We discuss technology trends, identify some open issues and discuss possible solutions. A detailed discussion of different MAC protocols for sensor-nets and their specific characteristics is out of the scope of this paper. Readers not familiar with medium access in wireless sensor-nets are encouraged to see Koen and Gertjan's book chapter on MAC for sensor-nets [5] before proceeding to Section 2.

2. RESEARCH ISSUES & DIRECTIONS

2.1 Towards a Sensor Network Architecture

The cross-layer designs in sensor-nets have lead to monolithic, vertically integrated solutions which might work independently but are not really useful for other research groups. One of the early encouraging steps towards a sensor network architecture is Polastre et al. sensornet protocol (SP) [6] which is a flexible alternative to the recent ZigBee standard. SP provides a standardized interface to MAC, with some feedback in both directions and provides an important step to building a larger sensor network architecture. Unlike IP in the Internet, SP is not at the network layer but instead sits between the network and data-link layer (because data-processing potentially occurs at each hop, not just at the end points). Developing a sensor network architecture would be a growing and organic process. Future MAC designers should keep the general goal of *moving towards a sensor network architecture* in mind and try to make use of the services that SP [6] has to offer.

2.2 Standardized Radio Hardware

While SP [6] is the emerging standard on top of the MAC layer in the sensor-net network stack, IEEE 802.15.4 is the emerging standard for lower layers (physical and medium access). IEEE 802.15.4 standard includes both physical (PHY) and medium access control (MAC) specifications. However, we expect that researchers would largely "override" the 802.15.4 MAC (with MAC protocols tailored for specific needs of sensor-nets) and only the 802.15.4 PHY standard would have implications on future sensor-net MAC designs. New and upcoming sensor-net hardware platforms, like Telos Motes (Moteiv), Sun SPOTS platform (Sun Labs), and MicaZ Motes (CrossBow), already use IEEE 802.15.4 compliant radios. By using a standardized radio, the nodes can communicate with any number of devices (possibly from different vendors) while sharing the same physical layer. On the other hand as the IEEE 802.15.4 radio interfaces are packet-based, the developers lose considerable flexibility in software for controlling the radios.

2.3 Peaceful Coexistence

With the anticipated large-scale introduction of sensor-nets into daily lives, a situation will emerge where sensor-nets from different vendors/operators need to operate at a common frequency band in the same physical environment (e.g., 2.4 GHz inside the home). This puts a requirement on the MAC protocols to behave "nice" to each other. Simply running at a low duty-cycle is not enough to warrant coexistence. First, MAC protocols need to be aware of others to guarantee flawless operation; foreign messages need to be filtered out, congestion detection needs to be enhanced to avoid improper back offs, schedules need to be adjusted to avoid overlapping active parts, etc. Second, MAC protocols must consider security issues to protect against eavesdropping and malicious behavior. Although it will be close to impossible to prevent denial-of-service attacks, maybe as brutal as jamming, observing this and signalling it to the routing layer may prevent a total break down of the application. TinySec [7] is a first step in securing MAC protocols, but its reliance on a shared key makes it vulnerable and, therefore more advanced schemes are needed. Heterogeneity, node failures and network extensions are complicating factors that have received little attention in the security arena, but are essential to the successful operation of sensor-nets in the future. Note that hardware solutions, like providing multiple channels at the radio, may alleviate some of the problems of coexistence, but the security issues will still need to be resolved.

2.4 Mobile Sensor Networks

The research community generally ignores mobility at the MAC-layer because sensor-nets were originally assumed to comprise of static nodes. However, recent works like Robo-Mote [8] and Parasitic-Mobility [9] have enabled mobility in sensor-nets. Furthermore, recently there has been an increased interest in medical care and disaster response applications of sensor-nets and these environments make use of mobile sensor nodes e.g. sensors attached to patients, doctors or first responders [10]. The only work, that we are aware of, which explicitly considers mobility at the MAC layer is MMAC [11] and there is much room for research in this area. Mobility evaluations should consider more *realistic* mobility models e.g. CodeBlue [10] only considers the typical movement of a doctor in hospital hallways, and MMAC [11] uses a fairly simple "random" mobility model for simulations. Furthermore, comprehensive *real* mobility traces could prove useful in such mobility evaluations.

2.5 New Optimization Criteria

To date, the primary design goal for sensor networks in general and MAC in particular has been energy efficiency. However, as new applications of sensor networks emerge, other optimization criteria (or Quality-of-Service parameters) such as latency and compliance with real-time constraints (e.g., monitoring and control in industrial environments), or reliable data delivery (e.g., medical applications) may gain importance. So far, little attention has been paid to them in the context of sensor networks. One particular issue is that many applications need to be optimized for multiple, conflicting criteria (e.g., energy-efficient and reliability). Hence, applications need a way to implement particular trade-offs between these conflicting goals. Here, MAC protocols would be required that provide "turning knobs" to provide application-specific trade-offs.

2.6 Understanding and Exploiting Traffic Patterns

Most existing MAC protocols for sensor networks are rather general purpose in that they support arbitrary communication patterns. While this allows the implementation of arbitrary distributed protocols, the price for flexibility is often limited efficiency (e.g., in terms of energy consumption).

However, many applications of sensor networks exhibit a few rather specific traffic patterns such as broadcast from the sink to all sensor nodes (tasking, query distribution, etc.) and convergecast from all nodes towards the sink (data collection). Also, traffic generation shows specific patterns in sensor networks. While sensor networks for continuous monitoring exhibit (aggregated) data streams from nodes to the sink, event-based applications generate traffic only upon the occurrence of certain interesting events. Since real-world events are often concurrently observed by many nodes in the neighborhood of the event, such event-based applications often result in highly correlated (both temporally and spatially) generation of traffic in the network. In addition, certain traffic patterns often appear in phases such as occasional query distributions followed by long phases of data collection or a burst of event reports followed by long idle phases.

With respect to the efficiency of MAC protocols, much could be gained by better understanding and exploiting these traffic patterns in the MAC design. Optimal solutions for specific patterns could then be integrated into traffic-adaptive MAC protocols that learn the current patterns and adapt their behavior accordingly.

2.7 Simulations Considered Harmful

Assumptions made in most simulation environments (a radio's transmission area is circular, all radios have equal range etc.) do not necessarily reflect the real-world conditions. In order to fully understand the complexity of designing a MAC protocol and to develop solutions which work in real life, it is necessary to not only model or simulate but also to implement and test on real world systems.

Axiom 0:	The world is flat.
Axiom 1:	A radio's transmission area is circular.
Axiom 2:	All radios have equal range.
Axiom 3:	If I can hear you, you can hear me.
Axiom 4:	If I can hear you at all, I can hear you perfectly.
Axiom 5:	Signal strength is a simple function of distance.

Table 1: Mistaken axioms of wireless-network research

It is important to revisit Kotz's mistaken axioms of wireless-network research [12] (see Table 1) to understand why MAC protocols that yield extremely accurate results in simulation fail in real life deployments (see [13] for experiences from a real sensor-net deployment). Kotz et al. surveyed MobiCom proceedings from 1995 to 2002 and classified the simulation radio models used in the works as: *Flat Earth*, *Simple*, and *Good* [12]. "Flat earth" models assume that two nodes could "perfectly" communicate if they are within some distance, say d, of each other. "Simple models" are ns-2 models which are more realistic than the "Flat Earth" models but are still fairly limited in emulating *real* radio propagation (ns-2 simulations are popular in sensor-net research as well). "Good models" are mainly used by the cellular telephony community and concentrate on exact mechanisms of RF propagation (takes into account factors like terrain, tree density, 3-D antenna location, foliage types, wavelength, etc.). Kurkowski et al. [14] surveyed MobiHoc proceedings from 2000-2005 and showed that published MANET simulation results lack believability. Sensor-net researchers need to conduct studies similar to Kotz et al. [12] and Kurkowski et al. [14] on the sensor-net literature and identify the "mistaken axioms" in the radio models used by popular sensor-net simulation environments. In Section 2.8, we discuss Software Defined Radios as possible means of bridging the gap between simulation and real world performance of MAC protocols.

2.7.1 Current Simulation Environments

Choice of the simulator makes an impact on the validity of the results. The radio models provided in the standard distribution of ns-2 (*free-space model* and *two-ray ground reflection model*) do not reflect the channel propagation typical of sensor-net environments. *Shadowing model* is the latest addition (2000) to ns-2 radio models and it takes into account the effects of indoor obstructions and outdoor shadowing. However, there is still a need for better RF propagation treatment in ns-2.

If using OMNeT++ discrete event simulator for sensor-net MAC simulations it is better to use MAC Simulator 0.2.2 [15] that has code for comparison protocols already available. The OMNeT++ MAC Simulator 0.2.2 implements radio characteristics of EYES nodes radio. However, it suffers from some of the mistaken axioms (Table 1) e.g. it assumes that the radio signals are circular.

TOSSIM [16] enables the MAC developers to choose the accuracy and complexity of the radio model as necessary for their simulations as the models are external to the simulator. However, like ns-2 and MAC Simulator 0.2.2, the available radio models in TOSSIM are fairly simple.

Emstar [17] helps the sensor-net developers to easily move from simulation to prototype to deployment. The important point to note about Emstar is that it provides an interface to *real* low-power radios instead of a simulated channel. Emstar provides support for multiple underlying radio types and drivers for radio link hardware (IEEE 802.11 and several flavors of Mica Motes) are already implemented in Emstar.

2.8 Embrace Software Radios

Software defined radios might be as revolutionary and transforming for wireless communication as packet switching was to circuits [3]. In software radios characteristics like medium access control, data encoding, frequency usage etc. are programmable instead of hard-wired and introducing new media access rules becomes a matter of a simple software change.

We argued in Section 2.7 that the performance results of MAC protocols in simulations should be taken with a grain of salt because of possible unrealistic assumptions made in radio models of the simulation environment. A possible solution around this problem could be to start designing for software defined radio systems and move from simulation to prototypes. Software radios allow network researchers to use inexpensive off-the-shelf networking cards to experiment with new MAC protocols. Neufeld et. al. have developed a software architecture (software defined radio) *SoftMAC*, on top of the radio subsystem (IEEE 802.11 family network cards) that permits network researchers to easily construct and deploy experimental dynamic MAC layers [18]. Developing a software architecture like SoftMAC [18] on top of the IEEE 802.15.4 compliant radios is crucial for future medium access research in sensor networks.

2.9 Real-World Experiments

While software defined radios are one direction towards obtaining more realistic insight into MAC layer performance, a further step is real-world experiments. Long-term experiments such as the one at Great Duck Island [19] are the most realistic ones, but require a lot of engineering effort before research questions can be answered [13]. Controlled indoor experiments are not as expressive as large-scale deployments but since real radio hardware is used, they do not suffer from the mistaken axioms listed in Table 1. Ritter et al. [20] have proposed to replace the batteries with high-capacity capacitors, so-called GoldCaps, for the purpose of experimental validation of lifetime bounds for wireless sensor networks. This approach can also be used to compare the energy-efficiency of MAC protocols without waiting for the batteries of the sensor nodes to drain. While this approach has inherent drawbacks, e.g. it does not take the battery relaxation effect into account, results obtained using this approach are more realistic than simulations.

3. CONCLUSIONS

When reviewing the open issues in the MAC for sensor-net area we find that there is no clear single direction in which future efforts should be directed. Some general recommendations, however, could be made:

- Energy efficiency might be the primary design goal for MAC protocols in sensor-nets but it should not be the *only* design goal. Optimization criteria such as latency, reliable data delivery, and compliance with real-time constraints may gain importance in the future.
- Applications of sensor-nets exhibit a few specific traffic patterns and much could be gained by better understanding and exploiting these traffic patterns in the MAC design (instead of supporting arbitrary communication patterns).
- MAC protocols should consider security issues to protect against eavesdropping and malicious behavior.
- Different MAC protocols for sensor-nets (possibly from different vendors/operators) operating at a common frequency band in the same physical environment should be able to peacefully coexist with each other.
- Existing MAC protocols for (static) sensor-nets fail to provide acceptable performance when applied to sensor-nets with mobile sensor nodes. MAC protocols for sensor-nets should explicitly address the effects of mobile sensor nodes in the protocol design.
- In order to obtain more realistic insight into MAC layer performance, sensor-net researchers should move from simulation to prototypes (software defined radios) or real-world experiments.

4. REFERENCES

[1] W. Ye, J. Heidemann, and D. Estrin, "An energy-efficient mac protocol for wireless sensor networks," in *Proceedings of the IEEE Infocom*. New York, NY, USA: IEEE, June 2002, pp. 1567–1576.

[2] J. Elson and D. Estrin, *Wireless Sensor Networks*. Kluwer Academic Publishers, 2004, ch. Wireless Sensor Networks: A bridge to the Physical World.

[3] D. Clark, C. Partridge, R. Braden, B. Davie, S. Floyd, V. Jacobson, D. Katabi, G. Minshall, K. Ramakrishnan, T. Roscoe, I. Stoica, J.Wroclawski, and L. Zhang, "Making the world (of communications) a different place," *ACM SIGCOMM Computer Communication Review*, July 2005.

[4] W. Ye and J. Heidemann, *Wireless Sensor Networks*. Kluwer Academic Publishers, 2004, ch. Medium Access Control in Wireless Sensor Networks, Also a similar version is available as ISI-TR-580.

[5] K. Langendoen and G. Halkes, *Embedded Systems Handbook*. CRC Press, Aug. 2005, ch. Energy-Efficient Medium Access Control, Available online at: http://pds.twi.tudelft.nl/∼koen/.

[6] J. Polastre, J. Hui, P. Levis, J. Zhao, D. Culler, S. Shenker, and I. Stoica, "A unifying link abstraction for wireless sensor networks," in *Proc. ACM SenSys'05*, Nov. 2005.

[7] C. Karlof, N. Sastry, and D. Wagner, "TinySec: a link layer security architecture for wireless sensor networks," in *Proc. SenSys'04*, Nov. 2004, pp. 162–175.

[8] K. Dantu, M. Rahimi, H. Shah, S. Babel, A. Dhariwal, and G. S. Sukhatme, "Robomote: Enabling mobility in sensor networks," in *IEEE/ACM Fourth International Conference on Information Processing in Sensor Networks (IPSN/SPOTS)*, Apr. 2005.

[9] M. Laibowitz and J. A. Paradiso, "Parasitic mobility for pervasive sensor networks," in *Third International Conference on Pervasive Computing (PERVASIVE 2005)*, Munich, Germany, May 2005.

[10] V. Shnayder, B. rong Chen, K. Lorincz, T. R. F. Fulford-Jones, and M. Welsh, "Sensor networks for medical care," Harvard University, Tech. Rep. Technical Report TR-08-05, April 2005.

[11] M. Ali, T. Suleman, and Z. A. Uzmi, "MMAC: A mobility-adaptive, collision-free mac protocol for wireless sensor networks," in *Proc. 24th IEEE IPCCC'05*, Phoenix, Arizona, USA, April 2005.

[12] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott, "Experimental evaluation of wireless simulation assumptions," in *Proceedings of the ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, October 2004, pp. 78–82.

[13] K. Langendoen, A. Baggio, and O. Visser, "Murphy loves potatoes: Experiences from a pilot sensor network deployment in precision agriculture," in *14th Int. Workshop on Parallel and Distributed Real-Time Systems (WPDRTS)*, Rhodes, Greece, April 2006.

[14] S. Kurkowski, T. Camp, and M. Colagrosso, "MANET simulation studies: the incredibles," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 9, no. 4, pp. 50–61, 2005.

[15] "Mac simulator 0.2.2." [Online]. Available: http://www.consensus.tudelft.nl/software.html

[16] P. Levis, N. Lee, M. Welsh, and D. Culler, "TOSSIM: accurate and scalable simulation of entire tinyos applications," in *ACM Sensys'03*. New York, NY, USA: ACM Press, 2003, pp. 126–137.

[17] L. Girod, T. Stathopoulos, N. Ramanathan, J. Elson, D. Estrin, E. Osterweil, and T. Schoellhammer, "A system for simulation, emulation, and deployment of heterogeneous sensor networks," in *ACM Sensys'04*. New York, NY, USA: ACM Press, 2004, pp. 201–213.

[18] M. Neufeld, J. Fifield, C. Doerr, A. Sheth, and D. Grunwald, "SoftMAC—flexible wireless research platform," in *Proc. HotNets-IV*, College Park, Maryland, USA, Nov. 2005.

[19] R. Szewczyk, J. Polastre, A. Mainwaring, and D. Culler, "Lessons from a sensor network expedition," in *Proc. EWSN 2004*, Berlin, Germany, Jan. 2004.

[20] H. Ritter, J. Schiller, T. Voigt, A. Dunkels, and J. Alonso, "Experimental evaluation of lifetime bounds for wireless sensor networks," in *Proc. EWSN 2005*, Istanbul, Turkey, Jan. 2005.