Medium Access Control in Wireless Sensor Networks

Wei Ye and John Heidemann

***Abstract*— This paper reviews medium access control (MAC), an enabling technology in wireless sensor networks. MAC protocols control how sensors access a shared radio channel to communicate with neighbors. Battery-powered wireless sensor networks with many nearby nodes challenge traditional MAC design. This paper discusses design trade-offs with an emphasis on energy efficiency. It classifies existing MAC protocols and compares their advantages and disadvantages in the context of sensor networks. Finally, it presents S-MAC as an example of a MAC protocol designed specifically for a sensor network, illustrating one combination of design trade-offs.**

***Index Terms*— Medium access control, wireless sensor networks, energy efficiency**

I. Introduction

A wireless sensor network is a special network with large numbers of nodes equipped with embedded processors, sensors and radios. These nodes collaborate to accomplish a common task such as environment monitoring or asset tracking. In many applications, sensor nodes will be deployed in an ad hoc fashion without careful planning. They must organize themselves to form a multi-hop, wireless communication network.

A common challenge in wireless networks is collision, resulting from two nodes sending data at the same time over the same transmission medium or channel. Medium access control (MAC) protocols have been developed to assist each node to decide when and how to access the channel. This problem is also known as channel allocation or multiple access problem. The MAC layer is normally considered as a sublayer of the data link layer in the network protocol stack.

MAC protocols have been extensively studied in traditional areas of wireless voice and data communications. Time division multiple access (TDMA), frequency division multiple access (FDMA) and code division multiple access (CDMA) are MAC protocols that are widely used in modern cellular communication systems [1]. Their basic idea is to avoid interference by scheduling nodes onto different sub-channels that are divided either by time, frequency or orthogonal codes. Since these sub-channels do not interfere with each other, MAC protocols in this group are largely collision-free. We refer to them as scheduled protocols.

Another class of MAC protocols is based on contention. Rather than pre-allocate transmissions, nodes compete for a shared channel, resulting in probabilistic coordination. Collision happens during the contention procedure in such systems. Classical examples of contention-based MAC protocols include ALOHA [2] and carrier sense multiple access (CSMA) [3]. In ALOHA, a node simply transmits a packet when it is generated (pure ALOHA) or at the next available slot (slotted ALOHA). Packets that collide are discarded and will be retransmitted later. In CSMA, a node listens to the channel before transmitting. If it detects a busy channel, it delays access and retries later. The CSMA protocol has been widely studied and extended; today it is the basis of several widely-used standards including IEEE 802.11 [4].

Sensor networks differ from traditional wireless voice or data networks in several ways. First of all, most nodes in sensor networks are likely to be battery powered, and it is often very difficult to change batteries for all the nodes. Second, nodes are often deployed in an ad hoc fashion rather than with careful pre-planning; they must then organize themselves into a communication network. Third, many applications employ large numbers of nodes, and node density will vary in different places and times, with both sparse networks and nodes with many neighbors. Finally, most traffic in the network is triggered by sensing events, and it can be extremely bursty. All these characteristics suggest that traditional MAC protocols are not suitable for wireless sensor networks without modifications.

This paper reviews MAC protocols for wireless sensor networks. After discussing the attributes of MAC protocols and design trade-offs for sensor networks (Section II), we present TDMA protocols (Section III) and contention-based protocols (Section IV). We then examine S-MAC as a case study of a sensor-net specific MAC protocol (Section V).

II. Trade-offs in MAC Design for Wireless Sensor Networks

This section discusses important attributes of MAC protocols and how design trade-offs can be made to meet the challenges of the sensor network and its applications. Because sensor networks are often battery constrained, we emphasize energy efficiency in MAC design.

A. MAC Attributes and Trade-offs

MAC protocols are influenced by a number of constraints. A protocol designer needs to make trade-offs among different attributes. This section examines MAC attributes and trade-offs in detail, and how their importance varies in the context of wireless sensor networks.

Collision avoidance is the basic task of all MAC protocols. It determines when and how a node can access the medium and send its data. Collisions are not always completely avoided

This work is in part supported by NSF under grant ANI-0220026 as the MACSS project, by DARPA under grant DABT63-99-1-0011 as the SCADDS project, and a grant from the Intel Corporation.

Wei Ye (weiye@isi.edu) and John Heidemann (johnh@isi.edu) are with the Information Sciences Institute, University of Southern California.

in regular operation; contention-based MAC protocols accept some level of collisions. But all MAC protocols avoid frequent collisions.

Energy efficiency is one of the most important attributes for sensor-net MAC protocols. As stated above, with large numbers of battery powered nodes, it is very difficult to change or recharge batteries for these nodes. In fact, some design goals of sensor networks are to build nodes that are cheap enough to be discarded rather than recharged, or that are efficient enough to operate only on ambient power sources. In all cases, prolonging the lifetime of each node is a critical issue. On many hardware platforms, the radio is a major energy consumer. The MAC layer directly controls radio activities, and its energy savings significantly affect the overall node lifetime. We explore energy conservation in more detail below.

Scalability and adaptivity are closely related attributes of a MAC protocol that accommodate changes in network size, node density and topology. Some nodes may die over time; some new nodes may join later; some nodes may move to different locations. A good MAC protocol should accommodate such changes gracefully. Scalability and adaptivity to changes in size, density, and topology are important attributes, because sensor networks are deployed in an ad hoc manner and often operate in uncertain environments.

Channel utilization reflects how well the entire bandwidth of the channel is utilized in communications. It is also referred to as bandwidth utilization or channel capacity. It is an important issue in cell phone systems or wireless local area networks (LANs), since the bandwidth is the most valuable resource in such systems and service providers want to accommodate as many users as possible. In contrast, the number of active nodes in sensor networks is primarily determined by the application. Channel utilization is normally a secondary goal in sensor networks.

Latency refers to the delay from when a sender has a packet to send until the packet is successfully received by the receiver. In sensor networks, the importance of latency depends on the application. In applications such as surveillance or monitoring, nodes will be vigilant for long time, but largely inactive until something is detected. These applications can often tolerate some additional messaging latency, because the network speed is typically orders of magnitude faster than the speed of a physical object. The speed of the sensed object places a bound on how rapidly the network must react. During a period that there is no sensing event, there is normally very little data flowing in the network. Sub-second latency for an initial message after an idle period may be less important than potential energy savings and longer operational lifetime. By contrast, after a detection, low-latency operation becomes more important.

Throughput (often measured in bits or bytes per second) refers to the amount of data successfully transfered from a sender to a receiver in a given time. Many factors affect the throughput, including efficiency of collision avoidance, channel utilization, latency and control overhead. As with latency, the importance of throughput depends on the application. Sensor applications that demand long lifetime often accept longer latency and lower throughput. A related attribute is *goodput*, which refers to the throughput measured only by data received by the receiver without any errors.

Fairness reflects the ability of different users, nodes, or applications to share the channel equally. It is an important attribute in traditional voice or data networks, since each user desires an equal opportunity to send or receive data for their own applications. However, in sensor networks, all nodes cooperate for a single common task. At any particular time, one node may have dramatically more data to send than some other nodes. Thus, rather than treating each node equally, success is measured by the performance of the application as a whole, and per-node or per-user fairness becomes less important.

In summary, the above attributes reflects the characteristics of a MAC protocol. For wireless sensor networks, the most important factors are effective collision avoidance, energy efficiency, scalability and adaptivity to densities and numbers of nodes. Other attributes are normally secondary.

B. Energy Efficiency in MAC Protocols

Energy efficiency is one of the most important issues in wireless sensor networks. To design an energy-efficient MAC protocol, we must consider the following question: what causes energy waste from the MAC perspective? The following sources are major causes of energy waste.

Collision is a first source of energy waste. When two packets are transmitted at the same time and collide, they become corrupted and must be discarded. Follow-on retransmissions consume energy too. All MAC protocols try to avoid collisions one way or another. Collision is a major problem in contention protocols, but is generally not a problem in scheduled protocols.

A second source is *idle listening*. It happens when the radio is listening to the channel to receive possible data. The cost is especially high in many sensor network applications where there is no data to send during the period when nothing is sensed. Many MAC protocols (such as CSMA and CDMA protocols) always listen to the channel when active, assuming that the complete device would be powered off by the user if there is no data to send.

The exact cost of idle listening depends on radio hardware and mode of operation. For long-distance radios (0.5km or more), transmission power dominates receiving and listening costs. By contrast, several generations of short-range radios show listening costs of the same order of magnitude as receiving and transmission costs, often 50–100% of the energy required for receiving. For example, Stemm and Katz measure that the power consumption ratios of idle:receiving:transmission are 1:1.05:1.4 [5] on the 915MHz Wavelan card, while the Digitan wireless LAN module (IEEE 802.11/2Mbps) specification shows the ratios are 1:2:2.5 [6]. On the Mica2 mote [7], the ratios for radio power draw are 1:1:1.41 at 433MHz with RF signal power of 1mW in transmission mode. Most sensor networks are designed to operate over long time, and the nodes will be in idle state for long time. In such cases, idle listening is a dominant factor of radio energy consumption.

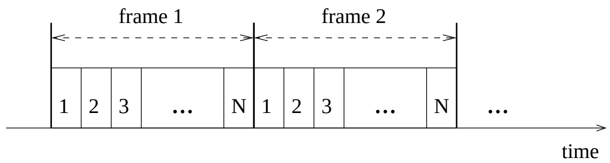

Fig. 1. TDMA divides the channel into N time slots.

A third source is *overhearing*, which occurs when a node receives packets that are destined to other nodes. Overhearing unnecessary traffic can be a dominant factor of energy waste when traffic load is heavy and node density is high.

The last major source that we consider is *control packet overhead*. Sending, receiving, and listening for control packets consumes energy. Since control packets do not directly convey data, they also reduce the effective goodput.

A MAC protocol achieves energy savings by controlling the radio to avoid or reduce energy waste from the above sources. Turning off the radio when it is not needed is an important strategy for energy conservation. A complete energy management scheme must consider all sources of energy consumption, not just the radio. In laptop computers, for example, display back-lighting can dominate costs [8].

On a tiny sensor node such as the Berkeley mote [9], the radio and the CPU are two major energy consumers. For example, on the Mica2 mote, the 433MHz radio consumes 22.2mW [10] when idle or receiving data, about the same power draw as the CPU when active [11], and is much higher than other components. From a system point-of-view, MAC energy control must be integrated with control of the CPU and other components.

III. Scheduled Protocols

According to the underlying mechanism for collision avoidance, MAC protocols can be broadly divided into two groups: scheduled and contention-based. Among protocols in the first group, TDMA has attracted attentions of sensor network researchers.

TDMA divides the channel into N time slots, as shown in Figure 1. In each slot, only one node is allowed to transmit. The N slots comprises a frame, which repeats cyclically. TDMA is frequently used in cellular wireless communication systems, such as GSM [1]. Within each cell, a base station allocates time slots and provides timing and synchronization information to all mobile nodes. Typically, mobile nodes communicate only with the base station; there is no direct, peer-to-peer communications between mobile nodes. The major advantage of TDMA is its energy efficiency, because it directly supports low-duty-cycle operations on nodes.

However, TDMA has some disadvantages that limits its use in wireless sensor networks. TDMA normally requires nodes to form clusters, analogous to the cells in the cellular communication systems. One of the nodes within the cluster is selected as the cluster head, and acts as the base station. This hierarchical organization has several implications. Nodes are normally restricted to communicate with the cluster head within a cluster; peer-to-peer communication is not directly supported. (If nodes communicate directly, then they must listen during all slots, reducing energy efficiency.) Inter-cluster communications and interference need to be handled by other approaches, such as FDMA or CDMA. More importantly, TDMA protocols have limited scalability and adaptivity to the changes on number of nodes. When new nodes join or old nodes leave a cluster, the base station must adjust frame length or slot allocation. Frequent changes may be expensive or slow to take effect. Also, frame length and static slot allocation can limit the available throughput for any given node, and the the maximum number of active nodes in any cluster may be limited. Finally, TDMA protocols depend on distributed, fine-grained time synchronization to align slot boundaries.

Many variations on this basic TDMA protocol are possible. Rather than scheduling slots for node transmissions, slots may be assigned for reception with some mechanism for contention within each slot. The base station may dynamically allocate slot assignments on a frame-by-frame basis. In ad hoc settings, regular nodes may assume the role of base station, and this role may rotate to balance energy consumption.

A. Examples of Scheduled Protocols

This subsection shows some examples of scheduled protocols for sensor networks. (We do not consider cellular communication systems here. Interested readers can refer to [1])

Sohrabi and Pottie proposed a self-organization protocol for wireless sensor networks [12]. The protocol assumes that multiple channels are available (via FDMA or CDMA), and any interfering links select and use different sub-channels. During the time that is not scheduled for transmission or reception, a node turns off its radio to conserve energy. Each node maintains its own time slot schedules with all its neighbors, which is called a superframe. Time slot assignment is only decided by the two nodes on a link, based on their available time. It is possible that nodes on interfering links will choose the same time slots. Although the superframe looks like a TDMA frame, it does not prevent collisions between interfering nodes, and this task is actually accomplished by FDMA or CDMA. This protocol supports low-energy operation, but a disadvantage is the relatively low utilization of available bandwidth. A sub-channel is dedicated to two nodes on a link, but is only used for a small fraction of time, and it cannot be re-used by other neighboring nodes.

LEACH (Low-Energy Adaptive Clustering Hierarchy), proposed by Heinzelman *et al.* [13] is an example of utilizing TDMA in wireless sensor networks. LEACH organizes nodes into cluster hierarchies, and applies TDMA within each cluster. The position of cluster head is rotated among nodes within a cluster depending on their remaining energy levels. Nodes in the cluster only talk to their cluster head, which then talks to the base station over a long-range radio. LEACH is an example that directly extends the cellular TDMA model to sensor networks. The advantages and disadvantages of LEACH are summarized above.

Bluetooth [14], [15] is designed for personal area networks (PAN) with target nodes as battery-powered PDAs, cell phones and laptop computers. Its design for low-energy operation and inexpensive cost make it attractive for use in wireless sensor networks. As with LEACH, Bluetooth also organizes nodes into clusters, called *piconets*. Frequency-hopping CDMA is adopted to handle inter-cluster communications and interference. Within a cluster, a TDMA-based protocol is used to handle communications between the cluster head (master) and other nodes (slaves). The channel is divided into time slots for alternate master transmission and slave transmission. The master uses *polling* to decide which slave has the right to transmit. Only the communication between the master and one or more slaves is possible. The maximum number of active nodes within a cluster is limited to eight, an example of limited scalability. Larger networks can be constructed as *scatternets*, where one node bridges two piconets. The bridge node can temporarily leave one piconet and join another, or operate two radios.

B. Energy Conservation in Scheduled Protocols

Scheduled protocols such as TDMA are very attractive for applications in sensor networks because of their energy efficiency. Since slots are pre-allocated to individual nodes, they are collision-free. There is no energy wasted on collisions due to channel contention. Second, TDMA naturally supports low-duty-cycle operation. A node only needs to turn on its radio during the slot that it is assigned to transmit or receive. Finally, overhearing can be easily avoided by turning off the radio during the slots of other nodes.

In general, scheduled protocols can provide good energy efficiency, but they are not flexible to changes in node density or movement, and lack of peer-to-peer communication.

IV. Contention-Based Protocols

Unlike scheduled protocols, contention protocols do not divide the channel into sub-channels or pre-allocate the channel for each node to use. Instead, a common channel is shared by all nodes and it is allocated on-demand. A contention mechanism is employed to decide which node has the right to access the channel at any moment.

Contention protocols have several advantages compared to scheduled protocols. First, because contention protocols allocate resources on-demand, they can scale more easily across changes in node density or traffic load. Second, contention protocols can be more flexible as topologies change. There is no requirement to form communication clusters, and peer-to-peer communication is directly supported. Finally, contention protocols do not require fine-grained time synchronizations as in TDMA protocols.

The major disadvantage of a contention protocol is its inefficient usage of energy. It normally has all the sources of energy waste we discussed in Section II: nodes listen at all times and collisions and contention for the media can waste energy. Overcoming this disadvantage is required if contention-based protocols are to be applied to long-lived sensor networks.

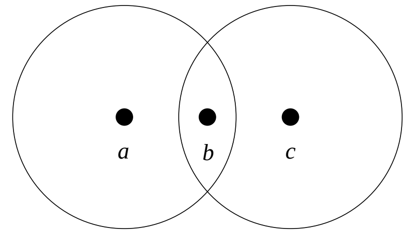

Fig. 2. Hidden terminal problem: nodes a and c are hidden to each other.

A. Examples of Contention Protocols

As mentioned in Section I, CSMA [3] is an important contention protocol. Its central idea is listening before transmitting. The purpose of listening is to detect if the medium is busy, also known as carrier sense. There are several variants of CSMA, including non-persistent, 1-persistent, and p-persistent CSMA. In non-persistent CSMA, if a node detects an idle medium, it transmits immediately. If the medium is busy, it waits a random amount of time and start carrier sense again. In 1-persistent CSMA, a node transmit if the medium is idle. Otherwise it continues to listen until the medium becomes idle, and then transmits immediately. In p-persistent CSMA, a node transmits with probability p if the medium is idle, and with probability $(1-p)$ to back-off and restart carrier sense. Woo and Culler examined the performance of CSMA with various configurations when it is used in wireless sensor networks [16].

In a multi-hop wireless network, however, CSMA alone is not sufficient due to the hidden terminal problem [17]. Figure 2 illustrates the hidden terminal problem on a two-hop network with three nodes. Suppose nodes a, b and c can only hear from their immediate neighbors. When node a is sending to b, node c is not aware of this transmission, and its carrier sense still indicates that the medium is idle. If c starts transmitting now, b will receive collided packets from both a and c.

CSMA/CA, where CA stands for collision avoidance, was developed to address the hidden terminal problem, and is adopted by the wireless LAN standard, IEEE 802.11 [4]. The basic mechanism in CSMA/CA is to establish a brief handshake between a sender and a receiver before the sender transmits data. The handshake starts from the sender by sending a short Request-to-Send (RTS) packet to the intended receiver. The receiver then replies with a Clear-to-Send (CTS) packet. The sender starts sending data after it receives the CTS packet. The purpose of RTS-CTS handshake is to make an announcement to the neighbors of both the sender and the receiver. In the example of Figure 2, although node c cannot hear the RTS from a, it can hear the CTS from b. If a node overhears an RTS or CTS destined to other nodes, it should back-off without sending its own packet. CSMA/CA does not completely eliminate the hidden terminal problem, but now the collisions are mainly on RTS packets. Since the RTS packet is very short, the cost of collisions is greatly reduced.

Based on CSMA/CA, Karn proposed MACA [18], which added a duration field in both RTS and CTS packets indicating the amount of data to be transmitted, so that other nodes know how long they should back-off. Bharghavan *et* *al.* further improved MACA in their protocol MACAW [19]. MACAW proposed several additions to MACA, including use of an acknowledgment (ACK) packet after each data packet, allowing rapid link-layer recovery from transmission errors. The transmission between a sender and a receiver follows the sequence of RTS-CTS-DATA-ACK.

IEEE 802.11 adopted all these features of CSMA/CA, MACA and MACAW in its distributed coordination function (DCF), and made various enhancement, such as virtual carrier sense, binary exponential back-off, and fragmentation support [4]. DCF is designed for ad hoc networks, while the point coordination function (PCF, or infrastructure mode) adds support where designated access points (or base-stations) manage wireless communication.

Woo and Culler proposed a MAC protocol for wireless sensor networks [16], which combined CSMA with an adaptive rate control mechanism. This protocol is based on a specific network setup where there is a base station that tries to collect data equally from all sensors in the field. The major problem faced by the network is that nodes that are closer to the base station carry more traffic, since they have to forward more data from nodes down to the network. The MAC protocol aims to fairly allocate bandwidth to all nodes in the network. Each node dynamically adjusts its rate of injecting its original packets to the network: linearly increases the rate if it successfully injects a packet; otherwise multiplicatively decreases the rate. This protocol does not use RTS and CTS packets to address the hidden terminal problem. Instead, a node relies on overhearing the transmissions of the next-hop node and longer back-off time in CSMA to reduce the effect of the hidden terminal problem.

B. Energy Conservation in Contention Protocols

Various techniques have been proposed to improve energy consumption of contention-based protocols for sensor networks. The basic approach is to put the radio into sleep state when it is not needed. For example, the Chipcon radio used on a Mica2 mote only consumes 15μW in sleep mode [10], three orders of magnitude less than that in idle/receive mode.

However, uncoordinated sleeping can make it difficult for adjacent nodes to communicate with each other. TDMA protocols provide structure by scheduling when nodes can communicate. Contention-based MAC protocols have explored similar but less restrictive sleep/wake schedules to improve energy consumption. Some examples are described in this subsection.

Piconet is a low-power ad hoc wireless network developed by Bennett *et al.* [20]. (It is not the same piconet in Bluetooth.) The basic MAC protocol used in Piconet is the 1-persistent CSMA protocol. To reduce energy consumption, each node sleeps autonomously. Since nodes do not know when their neighbors are listening, they beacon their ID each time they wake up. Neighbors with data for a particular destination must listen until they hear the destination's beacon. They then coordinate using CSMA.

In IEEE 802.11, both PCF and DCF have power-save (PS) modes that allow nodes to periodically sleep to conserve energy. Figure 3 shows the diagram of the PS mode in DCF. A basic assumption here is that all nodes can hear each other—the network consists of only a single hop. One node periodically broadcasts a beacon to synchronize all nodes clocks. All nodes participate in beacon generation, and if one node sends it out first, others will suppress their transmissions. Following each beacon, there is an ATIM (ad hoc traffic indication message) window, in which all nodes are awake. If a sender wants to transmit to a receiver in power save-mode, it first sends out an ATIM packet to the receiver. After the receiver replies to the ATIM packet, the sender starts sending data.

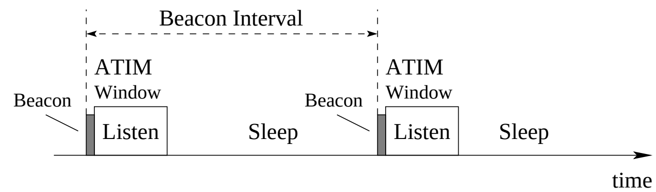

Fig. 3. The power-save (PS) mode in IEEE 802.11 DCF.

The above PS mode in 802.11 DCF is designed for a single-hop network. Generalizing it to a multi-hop network is not easy, since problems may arise in clock synchronization, neighbor discovery and network partitioning, as pointed out by Tseng *et al.* [21]. They designed three sleep patterns to enable robust operation of 802.11 power-saving mode in a multi-hop network. Their schemes do not synchronize the listen time of each node. Instead, the three sleep patterns guarantee that the listen intervals of two nodes periodically overlap. Thus it resolves the problems of 802.11 in multi-hop networks. The cost is the increased control overhead and longer delay. For example, to send a broadcast packet, the sender has to explicitly wake up each individual neighbor before it sends out the actual packet. Without synchronization, each node has to send beacons more frequently than the original 802.11 PS mode to prevent long-term clock drift.

Both Piconet and the 802.11 PS mode try to save energy by reducing the time of idle listening. They do not address the overhearing problem. PAMAS, proposed by Singh and Raghavendra [22], avoids overhearing by putting nodes into sleep state when their neighbors are in transmission. PAMAS uses two channels, one for data and one for control. All control packets are transmitted in the control channel. After a node wakes up from sleep, it also probes in the control channel to find any possible ongoing transmissions and their durations. If any neighbor answers the probe, the node will sleep again for the specified duration. Probing in the control channel avoids interfering a neighbor's transmission in the data channel, and the neighbor is able to answer the probe in the control channel without interrupting its data transmission. However, the requirement of separate control and data channels makes PAMAS more difficult to deploy, since multiple channels require multiple radios or additional complex channel allocation. Also, PAMAS does not reduce idle listening.

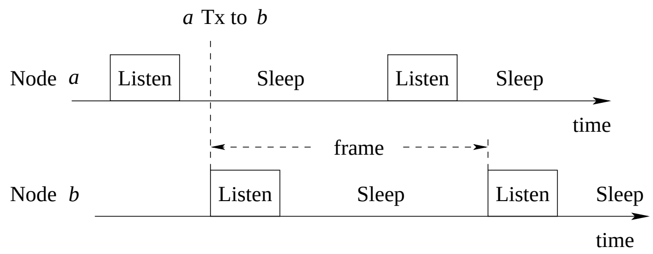

Fig. 4. Node a transmit to b when b starts listening.

V. Case Study: S-MAC

S-MAC is a MAC protocol specifically designed for wireless sensor networks, proposed by Ye *et al.* [23], [24]. Building on contention-based protocols like 802.11, S-MAC strives to retain the flexibility of contention-based protocols while improving energy efficiency in multi-hop networks. S-MAC includes approaches to reduce energy consumption from all the major sources of energy waste described in Section II: idle listening, collision, overhearing and control overhead.

A. S-MAC Design Approaches

At a high-level, S-MAC uses a coarse-grained sleep/wakeup cycle to allow nodes to spend most of their time asleep, as shown in Figure 4. We call a complete listen/sleep cycle a frame, after the TDMA frame. Each frame begins with a listen period for nodes that have data to send to coordinate. A sleep period follows, during which nodes sleep if they have no data to send or receive, and nodes remain awake and exchange data if they are involved in communication. We briefly describe how S-MAC establishes schedules in a multi-hop network, how nodes contend for the channel during listen periods, and how several optimizations improve throughput.

Scheduling: The first technique in S-MAC is to establish low-duty-cycle operation on nodes in a multi-hop network. For long-lived sensor networks, we expect duty cycles of 1–10%. The basic scheme is similar to the 802.11 PS mode, but without assuming all nodes can hear each other, or a designated base-station.

In S-MAC, all nodes are free to choose their own listen/sleep schedules. They share their schedules with their neighbors so that communication between all nodes is possible. Nodes then schedule transmissions during the listen time of their destination nodes. For example, nodes a and b in Figure 4 follow different schedules. If a wants to send to b, it just wait until b is listening. S-MAC enables multi-hop operation by accommodating multiple schedules in the network.

To prevent timing errors due to long-term clock drift, each node periodically broadcasts its schedule in a SYNC packet, which provides simple clock synchronization. The period for a node to send a SYNC packet is called a synchronization period. Combined with relatively long listen time and short guard time in waking up, S-MAC does not require tight clock synchronization of a TDMA protocol.

On the other hand, to reduce control overhead, S-MAC encourages neighboring nodes to adopt identical schedules. When a node first configures itself, it listens for a synchronization period and adopts the first schedule it hears. In addition, nodes periodically perform neighbor discovery, listening for an entire frame, allowing them to discover nodes on different schedules that may have moved within range.

Data transmission: The collision avoidance mechanism in S-MAC is similar to that in the IEEE 802.11 DCF [4]. Contention only happens at a receiver's listen interval. S-MAC uses both virtual and physical carrier sense. Unicast packets combine CSMA with an RTS-CTS-DATA-ACK exchange between the sender and the receiver, while broadcast packet use only CSMA procedure.

S-MAC puts a duration field in each packet, which indicates the time needed in the current transmission. If a neighboring node receives any packet from the sender or the receiver, it knows how long it needs to keep silent. In this case, S-MAC puts the node into sleep state for this amount of time, avoiding energy waste on overhearing. Ideally the node goes to sleep after receiving a short RTS or CTS packet destined to other nodes, and it avoids overhearing subsequent data and ACK packets. Compared with PAMAS, S-MAC only uses in-channel signaling for overhearing avoidance.

An important feature of wireless sensor networks is the in-network data processing, since data aggregation or other techniques can greatly reduces energy consumption by largely reducing the amount of data to be transmitted [25], [26], [27]. In-network processing requires store-and-forward processing of application-level *messages*, not just individual MAC-layer packets or fragments. While traditional MAC protocols emphasize fairness and interleave communication from concurrent senders, S-MAC utilizes *message-passing*, an optimization that allows multiple fragments from a message to be sent in a burst. It reduces message-level latency by disabling fragment-level interleaving of multiple messages.

In message passing, only one RTS and one CTS are used to reserve the medium for the time needed to transmit all fragments. Each fragment is separately acknowledged (and retransmitted if necessary). Besides RTS and CTS, each fragment or ACK also includes the duration of the remaining transmission, allowing nodes that wake up in the middle of the transmission to return to sleep. This is unlike 802.11's fragmentation mode where each fragment only indicates the presence of an additional fragment, not all of them.

With the low-duty-cycle operation, nodes must delay sending a packet until the next listen period of a destination, which increases latency. In addition, by limiting the opportunity to content for the channel, throughput can be reduced to one message per frame. These costs can accumulate at each hop of a multi-hop network. As an optimization to reduce this delay, S-MAC uses *adaptive listening*. Rather than waiting until the next scheduled listen interval after an RTS-CTS-DATA-ACK sequence, neighbors wake up immediately after the exchange completes. This allows immediate contention for the channel, either by another node with data to send, or for the next hop in a multi-hop path. With adaptive listen, the overall multi-hop latency can be reduced by at least half [24].

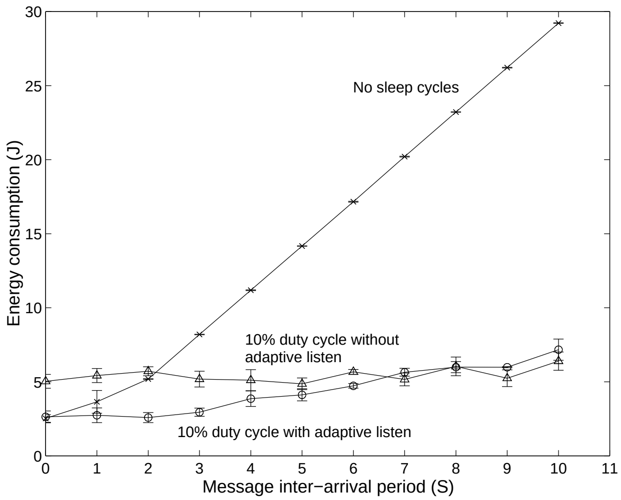

Fig. 5. Aggregate energy consumption on radios in the entire 10-hop network using three S-MAC modes. (From [24], ©2004 IEEE)

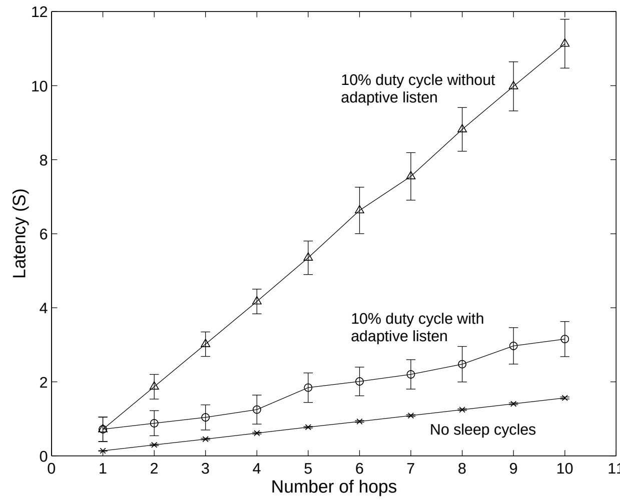

Fig. 6. Mean message latency on each hop under the lightest traffic load. (From [24], ©2004 IEEE)

B. S-MAC Performance

S-MAC has been implemented on Berkeley motes [28], [7]. Motes use an 8-bit embedded CPU and short-range, low-power radios: either an RFM TR1000 [29] or TR3000 [30], or a Chipcon CC1000 [10]. The following measurements use Mica motes with RFM TR3000 and 20kb/s bandwidth. An attractive feature of the mote for MAC research is that it provides very low-level access to the radio.

S-MAC implementation allows a user to configure it into different modes. This subsection shows some measurement results of S-MAC over Mica motes with the following configurations:

- 10% duty cycle without adaptive listen
- 10% duty cycle with adaptive listen
- No sleep cycles (100% duty cycle), but with overhearing avoidance

The topology in the measurement is a linear network of 11 nodes with the first node as the source and the last node as the sink. For complete details of these experiments, see [24].

1) Energy consumption: Energy consumption is measured in the ten-hop network with S-MAC configured in the above modes. In each test, the source node sends a fixed amount of data, 20 messages of 100-bytes each. Figure 5 shows how energy consumption on all nodes in the network changes as traffic load varies from heavy (on the left) to light (on the right).

Figure 5 shows that, at light load, operating at a low duty cycle can save significant amounts of energy compared to not sleeping, a factor of about 6 in this experiment. It also shows the importance of adaptive listening when traffic becomes heavy. Without adaptive listening, a 10% duty cycle consumes *more* energy than always listening because fewer opportunities to send require a longer time to send the same amount of data. By contrast, adaptive sending allows S-MAC to be as efficient as a MAC that never sleeps, even when there is always data to send.

2) Latency: A disadvantage of S-MAC is that the latency of sending a message can be increased. In this example, latency is measured by the time a message takes to travel over a varying number of hops when there is only one message in the network at a time.

Figure 6 shows the measured latency as a function of distance. In all three S-MAC modes, the latency increases linearly with the number of hops. However, S-MAC at 10% duty cycle without adaptive listen has much higher latency than the other two. The reason is that each message has to wait for one sleep cycle on each hop. In comparison, the latency of S-MAC with adaptive listen is very close to that of the MAC without any periodic sleep, because adaptive listen often allows S-MAC to immediately send a message to the next hop. On the other hand, for either low-duty-cycle mode, the variance in latency is much larger than that in the fully active mode, and it increases with the number of hops. The reason is that messages can miss sleep cycles in the path, and different parts of the network may be on different schedules.

3) Energy vs. Latency and Throughput: Now we look at the trade-offs that S-MAC has made on energy, latency and throughput. S-MAC reduces energy consumption, but it increases latency, and thus has a reduced throughput. To evaluate the overall performance, we compare the combined effect of energy consumption and reduced throughput by calculating the per-byte cost of energy and time to pass data from the source to the sink.

Figure 7 shows the results under different traffic loads. We can see that when traffic load is very heavy (on the left), adaptive listen and the no-sleep mode both show statistically equivalent performances that are significantly better than sleeping without adaptive listen. Without adaptive listen, the sleep delay at each hop lowers overall energy-time efficiency. At lighter traffic load, the energy-time cost without sleeping quickly exceeds the cost of sleep modes.

In summary, periodic sleep provides excellent performance at light traffic load. Adaptive listen is able to adjust to traffic,

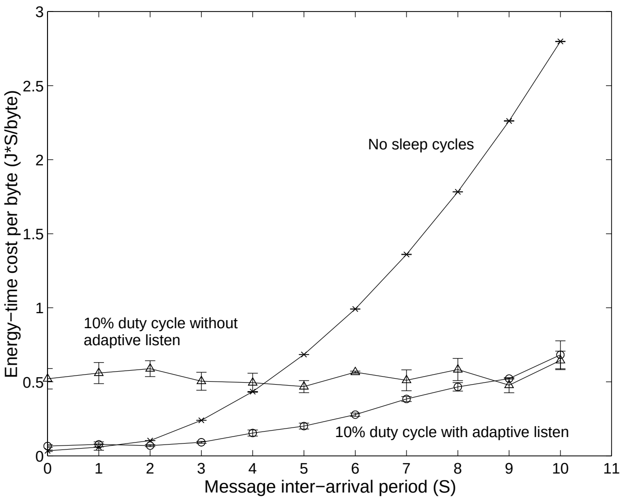

Fig. 7. Energy-time cost per byte on passing data from source to sink under different traffic load. (From [24], ©2004 IEEE)

and achieves performance as good as the no-sleep mode at heavy load. Therefore, S-MAC with adaptive listen is a good choice for sensor networks.

VI. SUMMARY

This paper reviews MAC protocols for wireless sensor networks. Large scale, battery powered wireless sensor networks put serious challenges to the MAC design. We have discussed important MAC attributes and possible design trade-offs, with an emphasis on energy efficiency. It described both scheduled and contention-based MAC protocols and evaluated their advantages and disadvantages when applied to sensor networks. Finally, we presented S-MAC as an example of sensor-net MAC protocols, illustrating the design trade-offs for energy conservation.

REFERENCES

[1] T. S. Rappaport, *Wireless Communications, Principles and Practice*, Prentice Hall, 1996.

[2] Norman Abramson, "Development of the ALOHANET," *IEEE Transactions on Information Theory*, vol. 31, no. 2, pp. 119–123, Mar. 1985.

[3] Leonard Kleinrock and Fouad Tobagi, "Packet switching in radio channels: Part I - carrier sense multiple access modes and their throughput delay characteristics," *IEEE Transactions on Communications*, vol. 23, no. 12, pp. 1400–1416, Dec. 1975.

[4] LAN MAN Standards Committee of the IEEE Computer Society, *Wireless LAN medium access control (MAC) and physical layer (PHY) specification*, IEEE, New York, NY, USA, IEEE Std 802.11-1999 edition, 1999.

[5] Mark Stemm and Randy H Katz, "Measuring and reducing energy consumption of network interfaces in hand-held devices," *IEICE Transactions on Communications*, vol. E80-B, no. 8, pp. 1125–1131, Aug. 1997.

[6] Oliver Kasten, *Energy Consumption*, `http://www.inf.ethz.ch/~kasten/research/bathtub/energy_consumption.html`, Eldgenossische Technische Hochschule Zurich.

[7] Crossbow Technology Inc., `http://www.xbow.com/`.

[8] Brian Marsh and Bruce Zenel, "Power measurements of typical notebook computers," Tech. Rep. MITL-TR-110-94, Matsushita Information Technology Laboratory, May 1994.

[9] Jason Hill, Robert Szewczyk, Alec Woo, Seth Hollar, David Culler, and Kristofer Pister, "System architecture directions for networked sensors," in *Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems*, Cambridge, MA, USA, Nov. 2000, pp. 93–104, ACM.

[10] Chipcon Inc., `http://www.chipcon.com/`, *Chipcon CC1000 Data Sheet*.

[11] Atmel Corporation, `http://www.atmel.com/`, *AVR Microcontroller ATmega128L Reference Manual*.

[12] Katayoun Sohrabi and Gregory J. Pottie, "Performance of a novel self-organization protocol for wireless ad hoc sensor networks," in *Proceedings of the IEEE 50th Vehicular Technology Conference*, 1999, pp. 1222–1226.

[13] Wendi Rabiner Heinzelman, Anantha Chandrakasan, and Hari Balakrishnan, "Energy-efficient communication protocols for wireless microsensor networks," in *Proceedings of the Hawaii International Conference on Systems Sciences*, Jan. 2000.

[14] Bluetooth SIG Inc., "Specification of the Bluetooth system: Core," `http://www.bluetooth.org/`, 2001.

[15] Jaap C. Haartsen, "The Bluetooth radio system," *IEEE Personal Communications Magazine*, pp. 28–36, Feb. 2000.

[16] Alec Woo and David Culler, "A transmission control scheme for media access in sensor networks," in *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, Rome, Italy, July 2001, pp. 221–235, ACM.

[17] Fouad Tobagi and Leonard Kleinrock, "Packet switching in radio channels: Part II - the hidden terminal problem in carrier sense multiple access and the busy-tone solution," *IEEE Transactions on Communications*, vol. 23, no. 12, pp. 1417–1433, Dec. 1975.

[18] Phil Karn, "MACA: A new channel access method for packet radio," in *Proceedings of the 9th ARRL Computer Networking Conference*, London, Ontario, Canada, Sept. 1990, pp. 134–140.

[19] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, "MACAW: A media access protocol for wireless lans," in *Proceedings of the ACM SIGCOMM Conference*, London, UK, Sept. 1994, pp. 212–225.

[20] Frazer Bennett, David Clarke, Joseph B. Evans, Andy Hopper, Alan Jones, and David Leask, "Piconet: Embedded mobile networking," *IEEE Personal Communications Magazine*, vol. 4, no. 5, pp. 8–15, Oct. 1997.

[21] Yu-Chee Tseng, Chih-Shun Hsu, and Ten-Yueng Hsieh, "Power-saving protocols for IEEE 802.11-based multi-hop ad hoc networks," in *Proceedings of the IEEE Infocom*, New York, NY, June 2002, pp. 200–209.

[22] S. Singh and C.S. Raghavendra, "PAMAS: Power aware multi-access protocol with signalling for ad hoc networks," *ACM Computer Communication Review*, vol. 28, no. 3, pp. 5–26, July 1998.

[23] Wei Ye, John Heidemann, and Deborah Estrin, "An energy-efficient mac protocol for wireless sensor networks," in *Proceedings of the IEEE Infocom*, New York, NY, June 2002, pp. 1567–1576.

[24] Wei Ye, John Heidemann, and Deborah Estrin, "Medium access control with coordinated, adaptive sleeping for wireless sensor networks," Tech. Rep. ISI-TR-567, USC Information Sciences Institute, Jan. 2003, To appear in the IEEE/ACM Transactions on Networking, April, 2004.

[25] Gregory J. Pottie and William J. Kaiser, "Embedding the internet: wireless integrated network sensors," *Communications of the ACM*, vol. 43, no. 5, pp. 51–58, May 2000.

[26] Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin, "Directed diffusion: A scalable and robust communication paradigm for sensor networks," in *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, Boston, MA, USA, Aug. 2000, pp. 56–67, ACM.

[27] John Heidemann, Fabio Silva, Chalermek Intanagonwiwat, Ramesh Govindan, Deborah Estrin, and Deepak Ganesan, "Building efficient wireless sensor networks with low-level naming," in *Proceedings of the Symposium on Operating Systems Principles*, Lake Louise, Banff, Canada, Oct. 2001, pp. 146–159.

[28] `http://webs.cs.berkeley.edu/tos`.

[29] RF Monolithics Inc., `http://www.rfm.com/`, *ASH Transceiver TR1000 Data Sheet*.

[30] RF Monolithics Inc., `http://www.rfm.com/`, *ASH Transceiver TR3000 Data Sheet*.