

Energy-Efficient, Collision-Free Medium Access Control for Wireless Sensor Networks *

Venkatesh Rajendran, Katia Obraczka, J.J. Garcia-Luna-Aceves

Department of Computer Engineering

Jack Baskin School of Engineering

University of California, Santa Cruz

Santa Cruz, CA 95064

E-mail: {venkat, katia, jj}@soe.ucsc.edu

Feb 11, 2004

Abstract. The traffic-adaptive medium access protocol (TRAMA) is introduced for energy-efficient collision-free channel access in wireless sensor networks. TRAMA reduces energy consumption by ensuring that unicast and broadcast transmissions incur no collisions, and by allowing nodes to assume a low-power, idle state whenever they are not transmitting or receiving. TRAMA assumes that time is slotted and uses a distributed election scheme based on information about traffic at each node to determine which node can transmit at a particular time slot. Using traffic information, TRAMA avoids assigning time slots to nodes with no traffic to send, and also allows nodes to determine when they can switch off to idle mode and not listen to the channel. TRAMA is shown to be fair and correct, in that no idle node is an intended receiver and no receiver suffers collisions. An analytical model to quantify the performance of TRAMA is presented and the results are verified by simulation. The performance of TRAMA is evaluated through extensive simulations using both synthetic- as well as sensor-network scenarios. The results indicate that TRAMA outperforms contention-based protocols (CSMA, 802.11 and S-MAC) and also static scheduled-access protocols (NAMA) with significant energy savings.

Keywords: Sensor networks, energy-efficient scheduling, traffic-adaptive medium access, ad hoc networks

1. Introduction

Sensor networks typically refer to large ensembles of interconnected nodes that, besides processing and communication capabilities, include one or more sensing devices (e.g., thermistors, magnetometers, light detectors). The deployment of such networks is usually done in an ad-hoc manner (e.g., dropping sensors from an aircraft on the field) which implies that sensor-network nodes need to self-organize into a multi-hop wireless ad-hoc network. Many large-scale sensor networks will consist of battery-powered sensor nodes whose battery may be difficult to recharge, or that the nodes themselves may be so cheap that recharging them may not be cost-effective.

* This work was supported in part by the NSF-NGI grant number ANI-9813724 and by the Jack Baskin Chair of Computer Engineering at UCSC.



As the hardware for sensor nodes has become increasingly more affordable and widely available, sensor networks have emerged as an ideal solution to a number of applications in both civilian and military scenarios. Some example applications include monitoring and surveillance of large, remote or inaccessible areas over extended periods of time. However, a major challenge facing the development and eventual deployment of large-scale sensor networks is the scheduling of transmissions among nodes in a way that (i) automatically adapts to changes in traffic, node state, or connectivity; and (ii) prolongs the battery life of each node.

BACKGROUND

There is an extensive body of work on medium access control (MAC) protocols for multihop wireless networks, dating back to DARPA's packet radio program (e.g., [16–18, 7, 8, 11]). These MAC protocols can be categorized as contention- and schedule-based.

The best-known example of contention-based MAC protocol is the distributed coordination function (DCF) of the IEEE 802.11b standard [13]. 802.11's DCF uses the carrier sense multiple access (CSMA) technique combined with a four-way handshake that attempts to avoid collisions of data packets. In terms of energy consumption, a key limitation of traditional contention-based schemes is that nodes consume energy needlessly when they are idle (i.e., not transmitting or receiving) as well as when collisions occur. Until recently, very little work has been reported on contention-based schemes that focus on energy efficiency.

One notable exception is PAMAS [20], one of the earliest contention-based proposals to address power efficiency in channel access. PAMAS saves energy by attempting to avoid overhearing among neighboring nodes. To achieve this, PAMAS uses out-of-band signaling. Woo and Culler [24] address variations of CSMA tailored for sensor networks, and propose an adaptive rate control mechanism to achieve fair bandwidth allocation among sensor network nodes. In the power save (PS) mode in IEEE 802.11 DCF, nodes sleep periodically. Tseng *et al.* [23] investigated three sleep modalities in 802.11 DCF for multi-hop networks. The sensor-MAC protocol [25], or S-MAC, exhibits similar functionality to that of PAMAS and the protocol by Tseng *et al.*. Like the other approaches, S-MAC avoids overhearing and nodes periodically sleep. However, unlike PAMAS, S-MAC uses in-line signaling, and unlike modalities of the PC mode in 802.11 DCF, neighboring nodes can synchronize their sleep schedules.

The probability of collisions of control or data packets in any contention-based scheme increases with the offered load, resulting in degraded channel utilization and reduced battery life. This motivates the need for establishing transmission schedules statically or dynamically to allow nodes to receive

data packets without collisions. The transmission schedule established in a wireless network can be topology-independent or dependent [6, 15, 19, 4]. The scheduled-access MAC protocol described by Sohrabi and Pottie [21] uses a combination of TDMA and FDMA or CDMA for accessing the channel. The main drawback of this scheme is that, like most fixed scheduling mechanisms, time slots are wasted if a node does not have any data to send to the intended receiver.

The Node Activation Multiple Access (NAMA) [4] uses a distributed election algorithm to achieve collision-free transmissions. For each time slot, NAMA selects only one transmitter per two-hop neighborhood and hence all nodes in the one-hop neighborhood of the transmitter are able to receive data collision-free. However, NAMA does not address energy conservation. In fact, very few proposals, if any, implement energy-aware medium access scheduling.

FOCUS

Section 2 introduces the TRAffic-Adaptive Medium Access (TRAMA) protocol, which provides energy-efficient conflict-free channel access in wireless sensor networks. Channel access in TRAMA is energy efficient while maintaining good throughput, acceptable latencies, and fairness. Energy efficiency is attained by (i) transmission schedules that avoid collisions of data packets at the receivers, and (ii) having nodes switch to low power radio mode when there is no data packets intended for those nodes. Adequate throughput and fairness is achieved by means of a transmitter-election algorithm that is inherently fair and promotes channel reuse as a function of the competing traffic around any given source or receiver. TRAMA derives collision-free transmission schedules based on (i) the identifiers of nodes one- and two hops away, (ii) the current time slot, and (iii) traffic information that specifies which node intends to transmit to which other node. Hence, the “sleep schedule” of a node is a direct function of the traffic going through the node and its neighbors, and is synchronized automatically when nodes exchange information about their identifiers and their traffic.

TRAMA is similar to the Node Activation Multiple Access (NAMA) protocol [4] as it provides conflict-free transmission by scheduling access among two-hop neighboring nodes during a particular time slot. However, NAMA does not address energy efficiency, i.e., nodes that are not transmitting switch to receiver mode. In contrast, TRAMA has energy efficiency as one of its main goals: it allows nodes to switch to sleep mode if they are not selected to transmit and are not the intended receivers of traffic for a particular time slot. Furthermore, by building schedules using traffic information, TRAMA, unlike NAMA, adapts to the application at hand. For instance, an event-tracking application will likely generate data only when an event is detected. On the

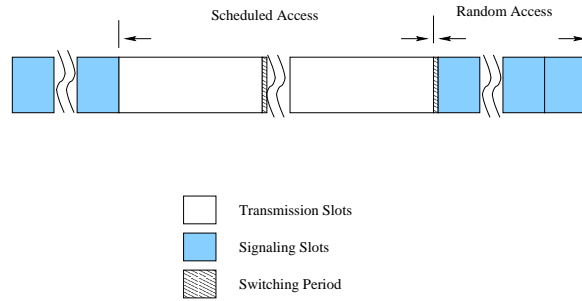


Figure 1. Time slot organization

other hand, monitoring applications may generate data continuously. In either case, TRAMA can adapt its schedules accordingly, delivering adequate performance and energy efficiency.

In contrast to prior MAC protocols proposed for sensor networks, TRAMA provides support for unicast, broadcast, *and* multicast traffic (i.e., transmitting to only a set of one-hop neighbors). TRAMA differs from S-MAC (which also provides explicit energy conservation mechanisms) in two fundamental ways: (i) TRAMA is inherently collision-free as its medium access control mechanism is schedule-based as opposed to S-MAC's which is contention-based; and (ii) TRAMA uses an adaptive, dynamic approach based on current traffic patterns to switch nodes to low power mode, while S-MAC's scheme is static based on a pre-defined duty cycle.

In Section 3, we show that TRAMA is fair and provides transmission schedules in which no collisions, idle listening, or idle senders occur. Section 4 presents an analytical performance comparison of TRAMA against a non-adaptive scheduled access MAC protocol. We evaluate the performance of TRAMA through extensive simulations using the Qualnet network simulator [22]. We compare the performance of TRAMA against three contention-based MAC protocols, namely CSMA, 802.11's DCF, and S-MAC, and also against NAMA, as a representative of collision-free channel access based on dynamic schedules. Section 5 describes our simulation setup, and Section 6 presents simulation results. Our simulation results show that TRAMA exhibits superior end-to-end throughput (around 40% over S-MAC and CSMA and around 20% for 802.11) for both synthetic traffic models and traffic models that are sensor-network specific. This is because TRAMA avoids collisions due to hidden terminals. Our results also show that TRAMA achieves significant energy savings (since nodes can sleep for up to 87% of the time) and higher throughput. Section 7 presents concluding remarks.

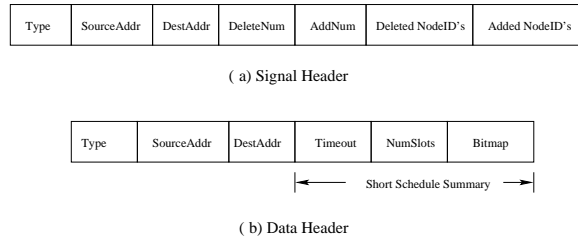


Figure 2. Signaling and data packet header format

2. TRAMA

2.1. PROTOCOL OVERVIEW

TRAMA employs a traffic adaptive distributed election scheme that selects receivers based on schedules announced by transmitters. Nodes using TRAMA exchange their two-hop neighborhood information and the transmission schedules specifying which nodes are the intended receivers of their traffic in chronological order, and then select the nodes that should transmit and receive during each time slot. Accordingly, TRAMA consists of three components: the Neighbor Protocol (NP) and the Schedule Exchange Protocol (SEP), which allow nodes to exchange two-hop neighbor information and their schedules; and the Adaptive Election Algorithm (AEA), which uses neighborhood and schedule information to select the transmitters and receivers for the current time slot, leaving all other nodes in liberty to switch to low-power mode.

TRAMA assumes a single, time-slotted channel for both data and signaling transmissions. Figure 1 shows the overall time-slot organization of the protocol. Time is organized as sections of random- and scheduled-access periods. We refer to random-access slots as *signaling slots* and scheduled-access slots as *transmission slots*. Because the data rates of a sensor network are relatively low, the duration of time slots is much larger than typical clock drifts. For example, for a 115.2 Kbps radio, we use a transmission slot of approximately 46ms to transmit 512-byte application layer data units. Hence, clock drifts in the order of ms can be tolerated, and yet typically clock drifts are in the order of microseconds or even less. This allows very simple timestamp mechanisms (e.g., [10]) to be used for node synchronization. When much smaller clock drifts must be assumed and more expensive nodes can be used, nodes can be time synchronized using techniques such as GPS [9]. Accordingly, in the remainder of our description of TRAMA, we simply assume that adequate synchronization is attained.

NP propagates one-hop neighbor information among neighboring nodes during the random access period using the *signaling slots* to obtain consistent two-hop topology information across all nodes. As the name suggests,

during the random access period, nodes perform contention-based channel acquisition and thus signaling packets are prone to collisions.

Transmission slots are used for collision-free data exchange and also for schedule propagation. Nodes use SEP to exchange traffic-based information, or schedules, with neighbors. Essentially, schedules contain current information on traffic coming from a node, i.e., the set of receivers for the traffic originating at the node. A node has to announce its schedule using SEP before starting actual transmissions. SEP maintains consistent schedule information across neighbors and updates the schedules periodically.

AEA selects transmitters and receivers to achieve collision-free transmission using the information obtained from NP and SEP. This is the case, because electing both the transmitter and the receiver(s) for a particular time slot is a necessity to achieve energy efficiency in a collision-free transmission schedule. Random transmitter selection leads to collisions, and electing the transmitters and not the receivers for a given time slot leads to energy waste, because all the neighbors around a selected transmitter have to listen in the slot, even if they are not to receive any data. Furthermore, selecting a transmitter without regard to its traffic leads to low channel utilization, because the selected transmitter may not have any data to send to the selected receiver. Hence, AEA uses traffic information (i.e., which sender has traffic for which receivers) to improve channel utilization.

The length of a transmission slot is fixed based on the channel bandwidth and data size. Signaling packets are usually smaller than data packets and thus transmission slots are typically set as a multiple of signaling slots to allow for easy synchronization. In our implementation, transmission slots are seven times longer than signaling slots.

2.2. ACCESS MODES AND THE NEIGHBOR PROTOCOL

In sensor networks, nodes may fail (e.g., power drained) or new nodes may be added (e.g., additional sensors deployed). To accommodate topology dynamics, TRAMA alternates between random- and scheduled access.

TRAMA starts in random access mode where each node transmits by selecting a slot randomly. Nodes can only join the network during random access periods. The duty cycle of random- versus scheduled access depends on the type of network. In more dynamic networks, random access periods should occur more often. In more static scenarios, the interval between random access periods could be larger, because topology changes need to be accommodated only occasionally. In the case of sensor networks, there is very little or no mobility, depending on the type of application. Hence, the main function of random access periods is to permit node additions and deletions. Time synchronization could be done during this period. During random access periods, all nodes must be in either transmit or receive state, so they

can send out their neighborhood updates and receive updates from neighbors. Hence, the duration of the random access period plays a significant role in energy consumption.

During random access periods, signaling packets may be lost due to collisions, which can lead to inconsistent neighborhood information across nodes. To guarantee consistent neighborhood information with some degree of confidence, the length of the random access period and the number of retransmissions of signaling packets are set accordingly. In [3], it is shown that, for a network with an average of N two-hop neighbors, the number of signaling packet retransmissions should be 7 and the retransmission interval $1.44 * N$ to guarantee packet delivery of 99%. Thus, the length of the random access period will then be $7 * 1.44 * N$.

NP gathers neighborhood information by exchanging small signaling packets during the random access period. Figure 2(a) shows the format of the header of a signaling packet. Signaling packets carry incremental neighborhood updates and if there are no updates, signaling packets are sent as “keep-alive” beacons. Each node sends incremental updates about its one-hop neighborhood as a set of added and deleted neighbors. These signaling packets are also used to maintain connectivity between the neighbors. A node times out a neighbor if it does not hear from that neighbor for a certain period of time. The updates are retransmitted such that we ensure 0.99 probability of success. Because a node knows the one-hop neighbors of its one-hop neighbors, eventually consistent two-hop neighborhood information makes its way across the network.

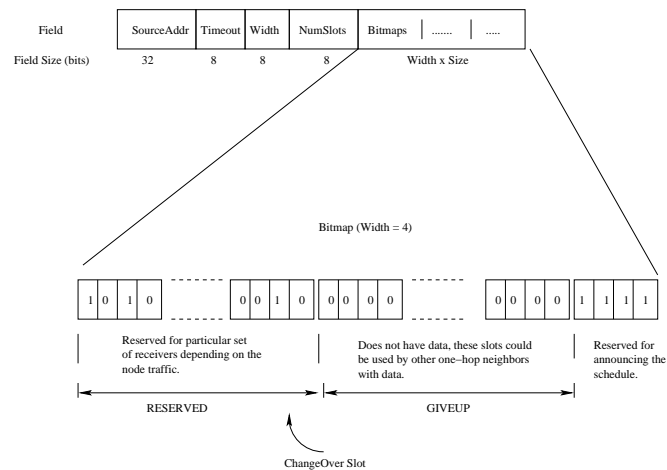
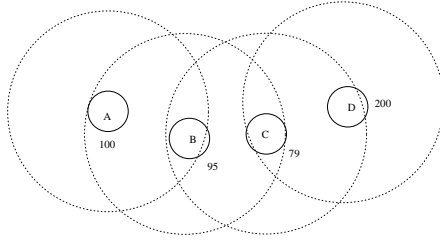


Figure 3. Schedule packet format



Node A: Will be transmitting since it has the highest two-hop priority

Node B: Will see D as the absolute transmitter (if it decides to SLEEP then TX to SLEEP can happen)

Figure 4. Inconsistency problem

2.3. SCHEDULE EXCHANGE PROTOCOL

SEP establishes and maintains traffic-based schedule information required by the transmitter (i.e., slot re-use) and receiver (i.e., sleep state switching) selection. A node's schedule captures a window of traffic to be transmitted by the node. This information is periodically broadcast to the node's one-hop neighbors during scheduled access.

Schedule generation works as follows. Each node computes a *SCHEDULE INTERVAL* based on the rate at which packets are produced by the higher layer application. The *SCHEDULE INTERVAL* of a node represents the number of slots for which the node can announce the schedule to its neighbors according to the current state of its MAC-layer queue. The node then pre-computes the number of slots in the interval $[t, t + \text{SCHEDULE INTERVAL}]$ for which it has the highest priority among its two-hop neighbors (contenders), which we call "winning slots". Because these are the slots for which the node will be selected as the transmitter, the node announces the intended receivers for these slots. Alternatively, if a node does not have enough packets to transmit, it announces that it gives up the corresponding slot(s). Other nodes that have data to transmit can make use of these "vacant" slots. A node's last winning slot in this interval is reserved for broadcasting the node's schedule for the next interval. For example, suppose that node u 's *SCHEDULE INTERVAL* is 100 slots. During time slot 1000, u computes its winning slots between $[1000, 1100]$. Let us assume that these slots are 1009, 1030, 1033, 1064, 1075, and 1098. Node u uses slot 1098, its last winning slot in this interval, to announce its next schedule by looking ahead from $[1098, 1198]$, and so on. The time corresponding to the last winning slot is fixed as the lifetime for the schedule.

Nodes announce their schedule via *schedule packets*. Because nodes have two-hop topology information obtained through NP, there is no need to send receiver addresses in the schedule packet. Instead, nodes convey intended receiver information using a bitmap whose length is equal to the number of one-hop neighbors. Each bit in the bitmap corresponds to one particular

receiver ordered by their identities. The total number of receivers supported by this scheme depends on the size of the data slot and the number of slots for which receivers are announced.¹ For example, a node with four one-hop neighbors with identities 14, 7, 5 and 4 will have a bitmap of size four with first MSB corresponding to node 14, second MSB to node 7. An advantage of using bitmaps is the ease with which broadcast and multicast communication can be supported. To broadcast a packet, all bitmap bits are set to 1, indicating that all one-hop neighbors are intended receivers of the packet. If the packet needs to be multicast to just 14 and 4, then only these bits are set in the bitmap. A node forms the bitmap for the winning slots based on the current traffic information for its queue. If the node's queue size is smaller than the number of bitmaps contained in the schedule, some of the winning slots will go unused. For these "vacant" slots, the node announces a zero bitmap. Slots with zero bitmaps could potentially be used by some other node in the two-hop neighborhood. The slot after which all the winning slots go unused is called *ChangeOver* slot. All unused slots happen contiguously toward the end before the last winning slot, which is reserved for announcing the next schedule. This maximizes the length of sleep periods.

Figure 3 shows the schedule packet format. *SourceAddr* is the address of the node announcing the schedule, *timeout* is the number of slots for which the schedule is valid (starting from the current slot), *width* is the length of the neighbor bitmap (i.e the number of one-hop neighbors), and *numSlots* is the total number of winning slots (i.e., the number of bitmaps contained in the packet). The last winning slot is always reserved for announcing the next schedule.

Additionally, a summary of a node's schedule is sent with every data packet. Schedule summaries help minimize the effects of packet loss in schedule dissemination. As shown in Figure 2(b), the summary includes the schedule's *timeout*, *numSlots*, and a bitmap corresponding to the winning slots in the current interval. The size of the bitmap is *numSlots* and is used to indicate whether the node is transmitting or giving up the corresponding slot. Note that, in order not to incur excessive overhead,² schedule summaries do not carry intended receiver information. They are meant to maintain synchronization³ among one-hop neighbor schedules even in the face of losses. For example by inspecting the values of *numSlots* and the *bitmap* in the schedule summary, the receiving node can update or re-synchronize its stored schedule information. Each schedule has an associated timeout and nodes are not

¹ Assuming the schedule is announced for 16 slots, the scheme can support 256 or 512 neighbors for a 512 or 1024 byte transmission slot size respectively.

² The overhead in the current implementation due to schedule summary is 6 bytes per data packet.

³ As demonstrated in Section 3, TRAMA's correctness is not affected by unsynchronized schedules.

allowed to change the schedule until this timeout expires. This is required to ensure consistency across one-hop neighborhood schedules.

Nodes maintain schedule information for all their one-hop neighbors. The schedule information is consulted whenever a node has the highest two-hop priority to decide if the node will actually transmit (i.e., it has data to send and thus will use the slot) or will give up the slot to another node in the neighborhood. Based on this decision, the schedule information for the node is updated either using the short summary from the data packet (if the node is receiving), or assuming transmissions (if the node is sleeping since it is not the intended receiver of transmitter). In the latter case, its schedule is in an unsynchronized state until the node verifies or updates it based on the schedule summary piggybacked in a future data packet from that transmitter.

All nodes listen during the *ChangeOver* slot of the transmitter to synchronize their schedule. For instance, if a node u keeps assuming transmissions for a particular neighbor tx at different timeslots and the neighbor does not transmit any packets due to a contender that is hidden from u , then the schedule at node u for node tx will be unsynchronized. If node u does not listen during the *ChangeOver* slot, which is the last slot in the current schedule interval that will be used by node tx for transmission, it may assume that the node tx is transmitting the data corresponding to the *ChangeOver* slot and update the corresponding schedule. From now on until the schedule announced by node tx expires, node u considers that node tx gives up winning slots for re-use. This can lead to collisions if node u tries to reuse the winning slot of node tx that is actually used for transmission. Hence, schedules among neighbors could be unsynchronized only until the *ChangeOver* slot and a node has to listen during the *ChangeOver* slot of the transmitter.

It is possible for a node to get some extra slots for transmissions in addition to the original winning slots computed while announcing the schedule. To prevent inconsistencies and collisions when transmitting the schedule packet, a node should always send the schedule packet only on the previously announced timeout. Because all the slots after the *Changeover* slot are assumed as give-up slots by the neighbors, the schedules might be unsynchronized. Hence, transmitting a schedule before the timeout can potentially cause collisions with neighbors. In the following section, we describe how these schedule information is used to adaptively decide the node state.

2.4. ADAPTIVE ELECTION ALGORITHM

In the original NCR algorithm [4], a node is selected to transmit if it has the highest priority among its *contending set*. Node u 's *contending set* is the set of all nodes that are in u 's two-hop neighborhood. Node u 's priority at time

slot t is defined as the pseudo-random hash of the concatenation of node u 's identity and t , or

$$prio(u,t) = hash(u \oplus t) \quad (1)$$

Assuming that node identities are unique and nodes are synchronized, all nodes compute the same priority value at any given time slot. However, if the selected node does not have any data to send, then the slot is wasted. Furthermore, nodes are free to transmit to any one-hop neighbor, because there is no sleep state in NCR.

For energy efficiency, TRAMA switches nodes to sleep state whenever possible, and attempts to re-use slots that are not used by the selected transmitter for bandwidth efficiency. A selected node may give up its transmission slot if it does not have any packets to send; this slot could then be used by another node. Nodes exchange current traffic information with their neighbors to make effective use of low-power, idle radio mode and accomplish slot re-use.

At any given time slot t during the scheduled access period, the state of a given node u is determined based on u 's two-hop neighborhood information and the schedules announced by u 's one-hop neighbors. Possible states are: transmit (TX), receive (RX), and sleep (SL).

At any given slot t , a node u is in the TX state if: (1) u has the highest priority, i.e., $prio(u,t)$ among its *contending set* and (2) u has data to send.

A node is in the RX state when it is the intended receiver of the current transmitter. Otherwise, the node can be switched off to the SL state, because it is not participating in any data exchange. This means that, if a node is not the selected transmitter, it will decide whether it needs to be in RX state by consulting the schedule sent out by the selected transmitter. If the transmitter does not have traffic destined for that node in the current slot, the node can then sleep.

Each node executes AEA to decide its current state (TX , RX , or SL) based on current node priorities (within its two-hop neighborhood) and also on the announced schedules from one-hop neighbors. The algorithm's pseudo-code is provided in Figure 15. Table I lists some basic terminology and notation used in the description of AEA.

The state of a node depends on the *Absolute Winner* and the announced schedules from its one-hop neighbors. From node u 's point of view, the *Absolute Winner* at any given time slot t can be: (1) node u itself, (2) node v that lies in the two-hop neighborhood of node u in which case the *Alternate Winner* $atx(u)$ needs to be accounted for if hidden from node v , or (3) a node w that lies in node u 's one-hop neighborhood.

Whenever a node becomes an *Absolute Winner* for a particular timeslot and has announced a non-zero bitmap for this slot, it knows that no other node in its two-hop neighborhood will be transmitting in this slot. Thus, the node can transmit collision-free to its intended receiver(s). When a node is

Table I. Notations and terminologies

$\mathbf{N2}(u)$	Set of neighbors of node u which are two-hops away.
$\mathbf{N1}(u)$	Set of neighbors of node u which are one-hop away.
$\mathbf{CS}(u)$	u 's <i>Contending Set</i> is the set of nodes in u 's two-hop neighborhood such that $\{u \cup \mathbf{N1}(u) \cup \mathbf{N2}(u)\}$.
$tx(u)$	<i>Absolute Winner</i> is the node with the highest priority in $\mathbf{CS}(u)$.
$atx(u)$	<i>Alternate Winner</i> is the node which has the highest priority among u 's one-hop neighbors, i.e., over the set $\{u \cup \mathbf{N1}(u)\}$.
$\mathbf{PTX}(u)$	<i>Possible Transmitter Set</i> is the set of all nodes in $\{u \cup \mathbf{N1}(u) - atx(u)\}$ that satisfy the condition given in Equation 2.
$\mathbf{NEED}(u)$	<i>Need Contender Set</i> is the set of nodes in $\{\mathbf{PTX}(u) \cup u\}$ that are in need of additional transmission slots.
$ntx(u)$	<i>Need Transmitter</i> is the node with the highest priority among the set of nodes $\mathbf{NEED}(u)$ containing valid synchronized schedule.

not an *Absolute Winner*, it is not certain who the actual transmitter for a particular slot is. For example, consider the topology shown in Figure 4. Let D be the node with highest priority in node B 's two-hop neighborhood in a given time slot and let A be the highest 2-hop priority node in node A 's two-hop neighborhood. Both A and D could transmit in the time slot because they are *Absolute Winners*; the *Absolute Winner* to node B is node D . Therefore, if B looks at its schedule information for D and finds out that it is not D 's intended receiver for the current slot, it will decide to switch to *SL* mode. However, if it happens to be A 's intended receiver, it will end up missing A 's transmission. Hence, before switching to *SL* mode, a node must also account for the *Alternate Winner*. This potential inconsistency occurs only if the *Alternate Winner* is hidden from the *Absolute Winner* i.e., they are three hops away.

To avoid wasting slots when the *Absolute Winner* has no data to send, TRAMA keeps track of nodes that could use extra slots to send their data. It first computes the set of nodes that can possibly transmit at the current time slot. They are kept in the *Possible Transmitter Set*, which contains all nodes in the one-hop neighborhood that can possibly transmit without any collision. A node can transmit without collisions only if it has the highest priority in the two-hop neighborhood. Hence, a node checks for possible transmitters in the one-hop neighborhood using the available information. Because a node cannot know the entire two-hop neighborhood of its one-hop neighbors, it can only check if this neighbor has the highest priority among the nodes that are

known to be the neighbor's two-hop neighbor. In other words, for a one-hop neighbor of node u , say node y , the following condition should be satisfied to be in the $\mathbf{PTX}(u)$:

$$prio(y) > prio(x) \forall x, x \in \mathbf{N1}(\mathbf{N1}(y)) \text{ and } x \notin \mathbf{N1}(y) \quad (2)$$

The *Need Contender Set* is the subset of the *Possible Transmitter Set* and contains only those nodes that have data to send. Nodes for which node u does not have (valid) schedules are also included in this set as node u does not know whether these nodes have data to send.

The *Absolute Winner* is the assumed transmitter for a node, unless the *Alternate Winner* is hidden from *Absolute Winner* and it belongs to the *Possible Transmitter Set*. In the latter case, the *Alternate Winner* is the assumed transmitter. Whenever the assumed transmitter gives up, the *Need Contender Set* is checked and the node with the highest priority within this set is selected as the *Need Transmitter* $ntx(u)$. Nodes that are not in the schedule listed by the assumed transmitter can switch to *SL* mode to save energy. This is especially beneficial in scenarios in which only a few nodes generate data at a time and data are destined to small subset of receivers.

3. TRAMA Correctness

TRAMA is correct if it avoids collisions and transmissions to a sleeping node, both of which can cause packet losses. A node can, however, assume that some of its neighbors is transmitting when the transmission does not actually happen. Though this will lead to increased energy consumption because nodes may be in receive mode unnecessarily, it does not affect the correctness of the algorithm.

Arguing for collision freedom is simple. The only two ways in which a node u can be a transmitter is through line 4 in the pseudo-code, where the node is the *Absolute Winner* and line 34 in the pseudo-code, where the node is the *Need Transmitter*. In both cases, there cannot be a node two-hops away from node u and transmitting. This follows from Equation 2 and by the definition of *Absolute Winner*. Hence, there can be no collisions due to transmissions from two-hop neighbors. Assuming that schedules are synchronized (which allows a node to know exactly whether the elected one-hop neighbor uses the slot or gives it up for re-use) and by virtue of the election mechanism, no other node that is one-hop away from node u can transmit. Hence, there can be no collisions due to a neighbor transmitting at the same time and the protocol maintains collision freedom at all the times.

To show that TRAMA never loses a packet due to an invalid state assignment (i.e., a node transmitting to a sleeping node), it is enough to show that, whenever a node u goes to sleep assuming that some node v is transmitting

in its one-hop neighborhood (given that node u is not the intended receiver of node v), then no other node except node v can transmit in the one-hop neighborhood. A node always considers the transmitter to be either the *Absolute Winner*, the *Alternate Winner*, or the *Need Transmitter*. Because a node can receive only from a node that is one-hop away, the *Absolute Winner* should be a one-hop neighbor. Hence, if a node assumes that a neighbor is transmitting, it is either the *Alternate Winner* or the *Need Transmitter*.

Consider the case in which node u decides to sleep during a time slot t assuming that node v is the transmitter. Hence, node v has the highest priority among the two-hop neighbors of node v known to node u and among the one-hop neighbors of node u that have data to send. Let node w be the actual transmitter for the time slot t in the one-hop neighborhood of node u . This means that node w has the highest priority among the two-hop neighbors of node w and it has the highest priority among the one-hop neighbors of node w that have data to send. Node w can either be a one-hop neighbor to node v or a two-hop neighbor to node v . If node w is a one-hop neighbor of node v , then node w should have higher priority than node v . Because, node v and node w are neighbors of node u and the schedules are synchronized, it contradicts the fact that node v has the highest priority among the one-hop neighbors of node u that have data to send. Hence, node w cannot be the actual transmitter. When node w is a two-hop neighbor to node v , it should have higher priority than node v , which contradicts the fact that node v has the highest priority among the two-hop neighbors of node v known to node u . Hence, there cannot be a transmitter, node w in the one-hop neighborhood of node u .

Schedules can get unsynchronized for different reasons, and TRAMA is also correct in the face of unsynchronized schedules. For example when a node assumes that node v is transmitting and decides to sleep, node v may not transmit. This will make the schedules unsynchronized and the *ChangeOver* slot for node v is reached earlier. The requirement that all the nodes should listen during the *ChangeOver* slot of the neighbors, prevents a node going to sleep or transmit state assuming that the neighbor is giving up and ensures correctness. Whenever a node assumes that a neighbor is transmitting a data to it, the schedules are updated only after receiving the data from the neighbor using the schedule summary. Hence, packet losses due to transmission errors can cause the schedules to be unsynchronized and forces a node to listen whenever the unsynchronized neighbor is elected for transmission. This continues until the node receives a data packet from the unsynchronized neighbor, and also prevents invalid state assignment. Hence, TRAMA is correct even when the schedules are not synchronized.

4. Analytical Model

In this section we present an analytical model for the delay performance of scheduling-access protocols (viz. NAMA and TRAMA) and use it to validate our simulation results.

The main goal of TRAMA's traffic-adaptive channel access is to achieve energy savings by making use of low-power, stand-by radio mode. The price to pay is the overhead involved in schedule propagation and also the latency inherent to scheduled-access MACs. Network latency consists of propagation, transmission and queuing delay; in the case of scheduled-access MACs, the latter component overshadows the other two. Scheduled-access using random priority for transmitter selection introduces a fixed delay irrespective of the traffic pattern. This *node activation* latency is mainly dependent on the number of contenders, which in turn depends on node density. In TRAMA, packets could be sent only after announcing the schedule. Hence, packets that arrive after the schedule announcement have to wait until the current schedule expires and a new schedule is announced. This introduces an additional *scheduling latency*, which, on average, is equal to $\frac{SCHEDULE_INTERVAL}{2}$. The schedules are transmitted during a transmission slot in the scheduled access period. This increases the latency for data packets because a transmission slot is wasted for announcing the schedule. On the other hand, adaptive slot reuse allocates slots more frequently than plain node activation schemes (e.g., NAMA). This increases the effective channel access probability for TRAMA when compared to plain node activation.

We assume that nodes are uniformly placed over the given area and all nodes have equal number of contenders. The channel access probability is the probability that a node wins a contention context and is given by,

$$q = \frac{1}{\text{Number of contenders}} \quad (3)$$

We also assume that channel access probability is uniform across nodes, packet arrivals follow a Poisson distribution, and all nodes generate equal length packets. The objective of the model is to determine a packet's average waiting time (before it gets serviced). We first derive this waiting time for scheduled access protocols and then extend it to NAMA and TRAMA.

A node can transmit only if it wins the contention context and this happens with probability q , the channel access probability, as given in Equation 3. A packet arriving at a node has to wait until the node succeeds in getting the channel. This can be modeled as the service time for a data packet. This service time could be modeled by a geometrically distributed random variable X with parameter q , which is the channel access probability. The random variable X is the number of the first winning slot. The mean and second

moment of X are given by,

$$\bar{X} = \frac{1}{q}, \quad \bar{X}^2 = \frac{(2-q)}{q^2} \quad (4)$$

Once the node services all the packets in the queue and goes idle, new arrivals to the idle system have to wait till the next winning slot. In other words, the node takes a vacation until the next winning slot. The interval until the next winning slot is again a geometrically distributed random variable V with parameter q . The random variable V is the number of failed slots before the first success. Thus, the mean and second moment of the vacation interval are:

$$\bar{V} = \frac{1-q}{q}, \quad \bar{V}^2 = \frac{2-3q+q^2}{q^2} \quad (5)$$

This system can be modeled using a M/G/1 service model with vacations [5]. Both the service time and vacation time follow a geometric distribution. The waiting time can be readily obtained [5] as:

$$W = \frac{\bar{X}^2 \cdot \lambda}{2(1-\lambda \cdot \bar{X})} + \frac{\bar{V}^2}{2\bar{V}} \quad (6)$$

The values of $\bar{X}, \bar{X}^2, \bar{V}, \bar{V}^2$ can be substituted in Equation 6 to get the average waiting time in the queue.

For node activation (i.e., NAMA), the channel access probability depends only on the contending set which is the entire two-hop neighborhood i.e., $q_{NAMA} = \frac{1}{N_1+N_2+1}$. In TRAMA, the channel access probability is also affected by adaptive slot reuse and schedule transmission. Adaptive slot reuse depends on the traffic in the one-hop neighborhood because we reallocate slots that are not used by winning nodes. These additional slots gained by slot reuse reduce the contender size by eliminating the set of one-hop nodes that do not have any data to send.

Thus, the effective contender set for calculating TRAMA's channel access probability is $(N_2(u) + N_1(u) + u - \{\text{onehop nodes without traffic}\})$. We represent this effective channel access probability taking into account channel reuse by q_e .

The reduction in the channel access probability due to schedule transmission can be quantified as:

$$q_{TRAMA} = \frac{SCHEDULE_INTERVAL \cdot q_e - 1}{SCHEDULE_INTERVAL}, \quad (7)$$

where q_{TRAMA} is the actual channel access probability. Equation 7 based on the observation that one slot is wasted for schedule propagation every $SCHEDULE_INTERVAL$. As it is difficult to calculate q_e analytically, we use the original channel access probability q_{NAMA} to get a lower bound for TRAMA's

channel access probability. The total waiting time for TRAMA also involves the *scheduling delay*. It is given as:

$$W = \frac{\bar{X}^2 \cdot \lambda}{2(1 - \lambda \cdot \bar{X})} + \frac{\bar{V}^2}{2\bar{V}} + \frac{SCHEDULE_INTERVAL}{2} \quad (8)$$

The term $\frac{SCHEDULE_INTERVAL}{2}$ accounts for the additional delay faced by a packet as it has to wait until the schedule is announced.

Figure 5 shows the average delay (in number of slots) for NAMA and an upper bound⁴ for the average delay for TRAMA for different two-hop neighborhood sizes. The number next to the protocol in the legend represents the total number of one-hop and two-hop neighbors (i.e. $N1 + N2$). The *SCHEDULE_INTERVAL* for TRAMA is fixed as 100 slots. As can be observed, TRAMA's average queuing delay is higher than NAMA. This is mainly due to the overhead introduced by the adaptive scheduling mechanism. In the next section we validate our analytical model by comparing to simulation results. We also show the improvement due to the adaptive slot reuse mechanism.

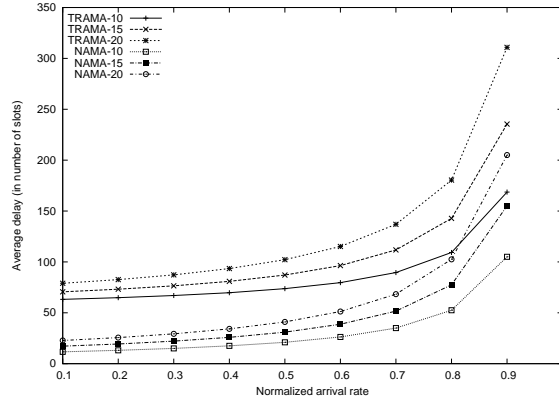


Figure 5. Average queuing delay for NAMA and TRAMA

4.1. SIMULATION RESULTS

The results from the analytical model are verified by simulation using the Qualnet [22] network simulator. The simulation parameters are set to satisfy the assumptions made in deriving the model. To achieve a uniform channel access probability across nodes (i.e., all nodes have equal number of contenders, which is an assumption we make in our model), 100 nodes are deployed in a 650x650m square grid topology. The transmission range of the

⁴ The delay is the upper bound as we assume there is no slot reuse i.e. $q_e = q_{NAMA}$.

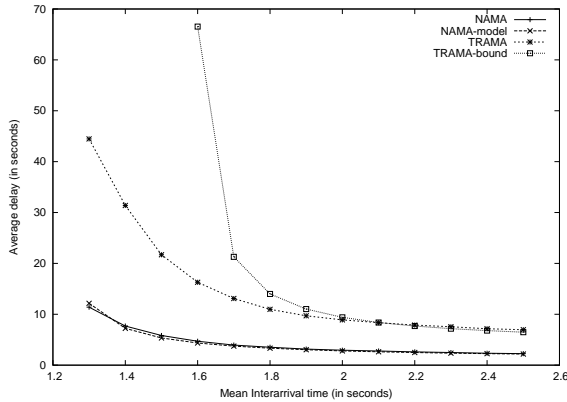


Figure 6. Comparison between average queuing delay obtained from analytical model and simulation for NAMA and TRAMA

nodes is set to 104m^5 and the grid unit is 65m. On average, nodes in the center of the grid have a contender size of 25 neighbors and this results in a channel access probability of $q_{NAMA} = 0.04$ for NAMA. The *SCHEDULE_INTERVAL* is fixed at 100 slots for TRAMA. The nodes in the edge have lesser number of neighbors and hence higher channel access probability. Therefore, these nodes are not considered for gathering the delay measurement. Traffic is generated based on Poisson arrivals and the queuing delay is measured at the MAC layer queue. Traffic is generated throughout the simulation period of 600 seconds. Results are averaged over multiple runs and are compared results from our analytical model.

Figure 6 shows the average delay for different inter-arrival times for TRAMA and NAMA. The analytical results coincide with the simulation results for NAMA. The actual delay for TRAMA based on simulation is less than the delay obtained from the analytical model. This is consistent with the fact that the analytical model provides an upper bound for TRAMA's average delay and does not account for slot-reuse. We can clearly see the improvement in delay due to the adaptive slot reuse in TRAMA. For lower traffic rates, delay is mainly dominated by the *scheduling delay*. Hence, there is no improvement in the delay due to slot reuse. For higher traffic loads, the delay due to the contention for channel access dominates. Here the slot reuse mechanism improves the channel access probability and the total observed delay is less than the modeled delay.

⁵ This value is set based on the type of radio we assume for sensor networks. This is discussed in more detail in Section 5.

5. Experimental Setup

Through simulations, we evaluate TRAMA and compare its performance against both contention- and scheduling-based protocols. While we consider Carrier-Sense Multiple Access (CSMA) [16], IEEE802.11 DCF [13] and S-MAC [25] as example contention-based protocols, we use Node Activation Multiple Access (NAMA) [4] as example scheduling-based protocol.

We used Qualnet [22] as our simulation platform and we present the results for a variety of scenarios. The underlying physical layer model used for all the experiments was based on the TR1000, a typical radio used in sensor networks. The TR1000 [2], the radio used by the UC Berkeley Motes [1], are short range, low data-rate (a maximum of 115.2KBPS) radios with built-in support for low-power sleep state. The average power consumption in transmit, receive and sleep modes is $24.75mW$, $13.5mW$ and $15\mu W$, respectively. The maximum transition time for switching is $20\mu S$. The modulation type used in the physical layer is ASK and the receiver threshold is $-75dBm$. Fifty nodes are uniformly distributed over a 500m x 500m area in all the experiments. The transmission range of each node is 100m and the topology is such that the nodes have 6 one-hop neighbors on average. The average size of the two-hop neighborhood for this network is 17 nodes. Two different types of traffic load are considered in our study. We used a scenario in which node traffic is statistically generated based on an exponentially distributed inter-arrival time. We chose this to stress-test protocol performance for different arrival rates. We also test TRAMA's performance when driven by data gathering applications, which are considered typical of sensor networks. Below, we describe these traffic scenarios as well as other simulation parameters in detail.

5.1. PROTOCOL PARAMETERS

In both scenarios we fixed up the *SCHEDULE_INTERVAL* to be 100 transmission slots for TRAMA. The maximum size of a signaling packet is fixed at 128 bytes which gives to a slot period of $6.82ms$ with guard time to take care of switching. Transmission slots are seven times longer than the signaling slots supporting a maximum data fragment size 896 bytes. The random access period is fixed to 72 transmission slots and is repeated once every 10000 transmission slots.

S-MAC is a contention-based channel access protocol and it uses periodic sleep intervals to conserve energy. Sleep schedules are established using *SYNC* packets which are exchanged once every *SYNC_INTERVAL*. The duty cycle determines the length of the sleep interval.

We set *SYNC_INTERVAL* as $10sec$ and we varied the duty cycle (10% and 50%). All the nodes are time synchronized and hence we favored S-MAC by

allowing the listen and sleep periods synchronized across the entire network. We also observed that S-MAC needed larger time to set up the listen/sleep schedules with the neighbor. This is because S-MAC does not have a proper neighbor discovery protocol, it has to rely on the *SYNC* packets for doing this. *SYNC* packets are transmitted only once and are transmitted unreliably. Hence, a large warmup time of *20sec* is allowed for the neighbor information to settle down. Because the queuing delay for the scheduling-type MAC's is higher, we allowed some more time for delivering the queued packets before ending the simulation. The simulation is run for *400sec* and the results are averaged over multiple runs.

5.2. SYNTHETIC DATA GENERATION

The objective of this experiment is to measure the performance of TRAMA when all the nodes in the network generate traffic based on some statistical distribution. We used exponential inter-arrival for generating data and varied the rate from 0.5 to 2.5 seconds. A neighbor is randomly selected as a next-hop every time a node transmits a packet. We tested both unicast and broadcast data generation separately. The performance metrics are:

- **Average Packet Delivery Ratio:** It is the ratio of number of packets received to the number of packets sent averaged over all the nodes. For broadcast traffic a packet is counted to be received only if it is received by all the one-hop neighbors.
- **Percentage Sleep Time:** It is the ratio of the number of sleeping slots to the total number of slots averaged over the entire network.
- **Average queuing Delay:** Average delay for the packet to be delivered to the receiver
- **Average Sleep Interval:** This is the average length of sleeping interval. This measures the number of radio mode switching involved. Frequent

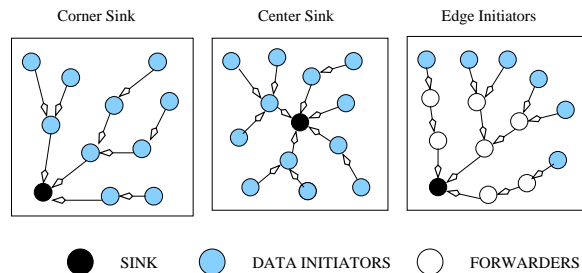


Figure 7. Data gathering application

switching can waste energy due to the transient power consumption involved in switching.⁶

5.3. DATA-GATHERING APPLICATION

We assume a sink is collecting data from all the sensors for these experiments. The sink sends out a broadcast query requesting data from all the sensors. The sensors respond back with the data, which are generated periodically to the sink. We implemented a simple reverse-path routing to forward the data from the sensors to the sink. Figure 7 shows the three different scenarios considered for this study. Data-collection node or sink is placed in the corner for the first case and in the middle for the second case.

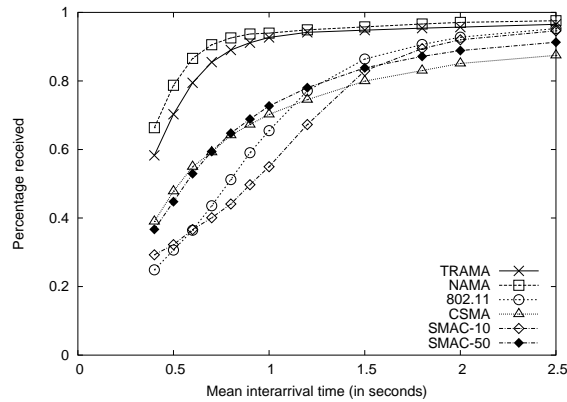
All the sensors respond with periodically generated data in both cases. Because data aggregation [14] or grouping data to minimize traffic, are advantageous, we also emulated data aggregation in a third case. Here, only the nodes at the edge generate traffic and we assume that the nodes do data aggregation and appends its reading to the parent node. To measure performance in these experiments we use the metrics defined for the synthetic case. The average packet delivery ratio is measured as the ratio of total number of data received by the sink to the total number of data sent by all the sensors, unlike the per-hop delivery ratio used for synthetic traffic generation.

6. Simulation Results

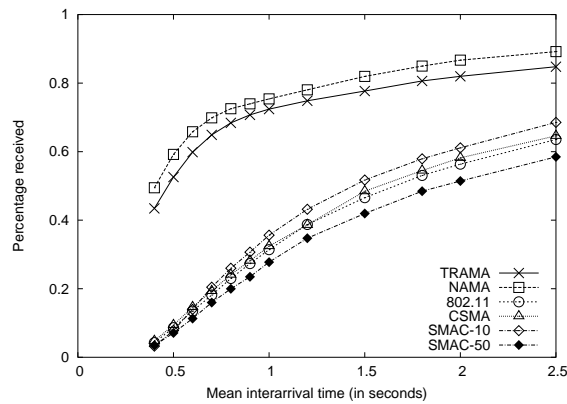
6.1. SYNTHETIC TRAFFIC

The packet delivery ratio, average queuing delay, percentage of sleep time, and average length of sleep intervals for synthetic traffic scenarios are shown in Figure 8, Figure 9, Figure 10(a) and Figure 10(b), respectively. We present results for S-MAC using two different duty cycles, namely 50% and 10%. Our results indicate that, in general, schedule-based MACs based on NCR achieve better delivery than IEEE802.11, CSMA and S-MAC. The main reason for the improvement in delivery is the collision freedom guaranteed at all times during data transmission. The effect is more noticeable when all nodes generate broadcast traffic. In CSMA, S-MAC and IEEE802.11, broadcasting is unreliable and susceptible to hidden-terminal collisions. This reduces broadcast delivery significantly when we increase the load as our results indicate. For IEEE802.11 and S-MAC, the RTS/CTS/DATA/ACK exchange for unicast traffic improves delivery when compared to broadcast traffic because it reduces hidden-terminal collisions by doing collision avoidance.

⁶ Measurements for 802.11 based radios are available in [12].



(a) Unicast Traffic

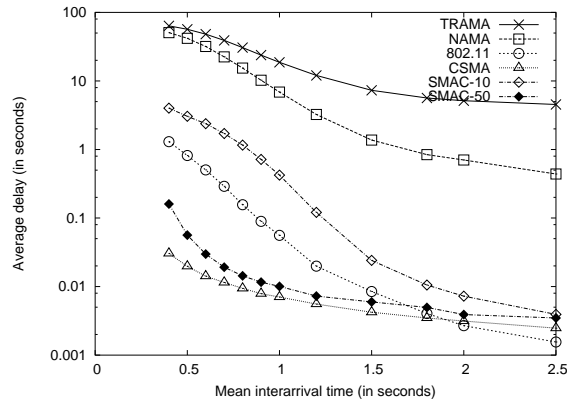


(b) Broadcast Traffic

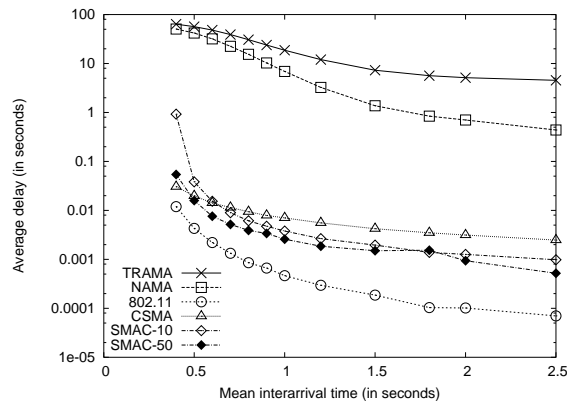
Figure 8. Average packet delivery ratio for synthetic traffic

Schedule-based MACs, on the other hand, incur higher average queuing delays. We should point out that when measuring average delay, we account for the delay of packets successfully delivered. However, TRAMA and NAMA deliver more packets than contention-based MACs and this will reduce the retransmissions at the higher layers. Hence, the end-to-end delay perceived by the application will be comparable to that of contention-based protocols.

The average queuing delay for TRAMA is higher than that of NAMA due to the overhead involved in propagating scheduling information. Once every *SCHEDULE INTERVAL*, a transmission slot is used for announcing sched-



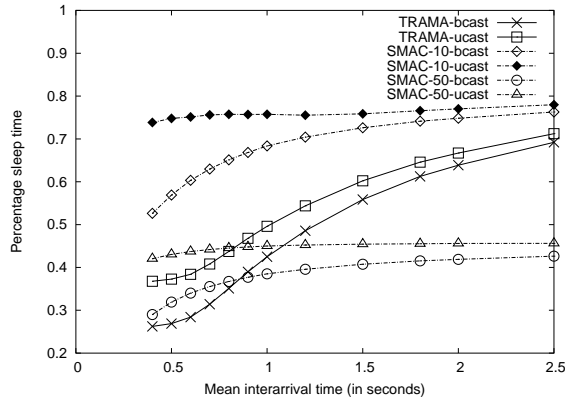
(a) Unicast Traffic



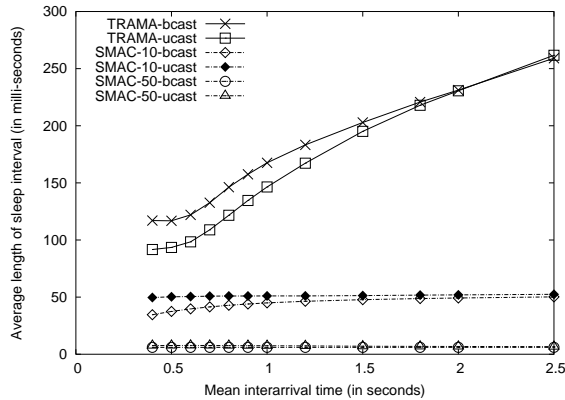
(b) Broadcast Traffic

Figure 9. Average queuing delay for synthetic traffic

ules. This decreases the effective channel access probability for data transmission. This scenario is not a favorable scenario for traffic-adaptive elections, because the traffic is homogeneous across the network and all nodes periodically generate traffic. The throughput of TRAMA is comparable with that of NAMA and is significantly better than contention-based protocols for both unicast and broadcast traffic scenarios. The performance of the only other energy-efficient protocol, S-MAC is comparable with IEEE802.11 in terms of throughput. However, the delay is slightly higher than that of IEEE802.11 or CSMA due to the sleep periods, and it increases as the duty cycle of the listen periods is decreased. For a duty cycle of 50%, the delay of S-MAC is smaller



(a) Percentage Energy Savings



(b) Average Sleep Interval

Figure 10. Energy savings and average sleep interval for synthetic traffic

than that of IEEE802.11 for the unicast-data generation scenario. This is because S-MAC frequently switches between sleep and listen modes (average sleep interval plotted in Figure 10(b) reflects this) for that duty cycle. This is equivalent to a node being awake most of the time, and the delay is less because S-MAC does not have any contention resolution algorithm.

The energy savings of TRAMA depend mainly on the traffic pattern, while the energy savings of S-MAC are dependent on the duty cycle. The total energy savings depend on both percentage sleep time and average length of sleep interval. Percentage sleep time metric does not account for the performance loss that is possible due to frequent radio mode switching. The

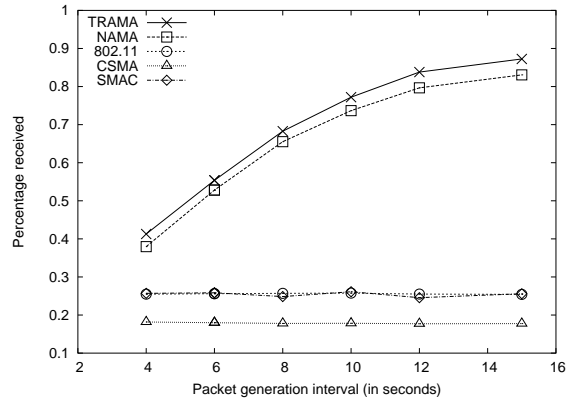
average length of sleep interval quantifies the amount of radio-mode switching involved. A higher value of average sleep length is preferred because this implies less radio-mode switching and hence more savings. The results indicate that the percentage of sleep time is less for broadcast traffic when compared to unicast traffic, which is intuitive. The percentage sleep time of S-MAC increases as the duty cycle decreases. The price paid for decreasing the duty cycle is an increase in latency. The throughput decreases more steeply for a lower duty cycle as traffic increases. For the broadcast packets, a 50% duty cycle achieves less throughput than 10% duty cycle. Broadcast is done by plain carrier sensing and is more prone to hidden terminal collisions. The timing structure for 512 bytes data that favors low duty cycle as it reduces the channel contention and hence the collisions are reduced. Note that a broadcast packet is counted as delivered ONLY if it is delivered to all the neighbors.

Compared with TRAMA, S-MAC with 10% duty cycle exhibits higher percentage of sleeping time. But the average length of sleep intervals is much lower for S-MAC when compared to TRAMA. This reduces the overall energy savings due to the overhead involved in mode switching. This is the case even though S-MAC is being favored by assuming that synchronized listen/sleep schedules are established across the nodes due to the simulation setup. The performance of TRAMA is not affected by nodes joining at discrete intervals because it does not require any synchronization of listen/sleep schedules.

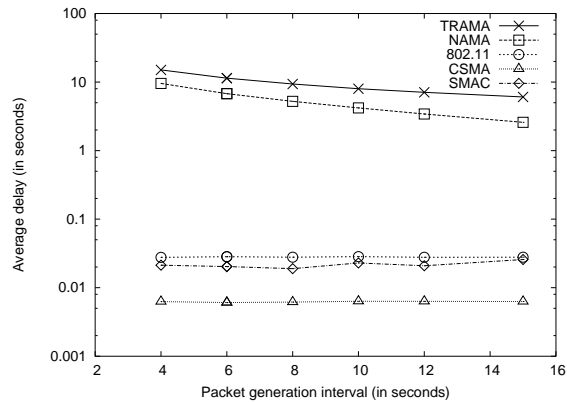
In the subsequent experiments, we only consider S-MAC with 10% duty cycle, which has better performance than with a 50% duty cycle.

6.2. SENSOR NETWORK APPLICATION

We tested the protocols using a sensor network data gathering application. One of the nodes in the network is designated as the sink and the sink starts sending a broadcast query. All nodes receiving a non-duplicate query add the sender of the query as the next hop for data forwarding, establishing a reverse-shortest path tree with the sink node as the root. Figures 11(a), 12(a), and 13(a) show the average packet delivery ratio for the corner sink, center sink, and edge sink scenarios respectively. Schedule-based MAC protocols outperform the contention-based MAC protocols in all the cases. The delivery is highest for the scenario in which the edge nodes are generating traffic. This is because the overall load in the network is low and well within the capacity of the protocols. Delivery to the center sink is slightly higher than when corner sink is used because the packets need to go through fewer number of hops to reach the sink. TRAMA performs much like NAMA and the decrease in throughput due to scheduling overhead is overcome by TRAMA's adaptive scheduling approach. The average delivery ratio is nearly constant for contention-based protocols as the variation in the offered load is not high.



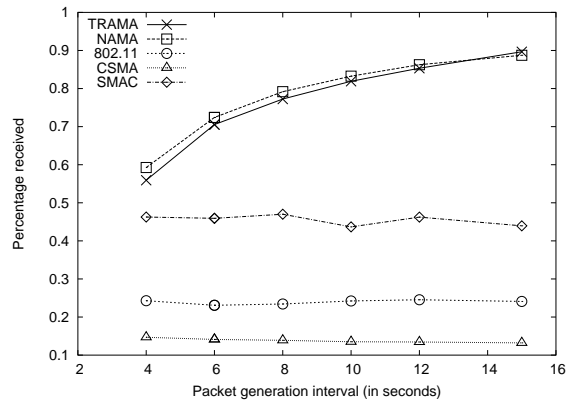
(a) Average delivery ratio



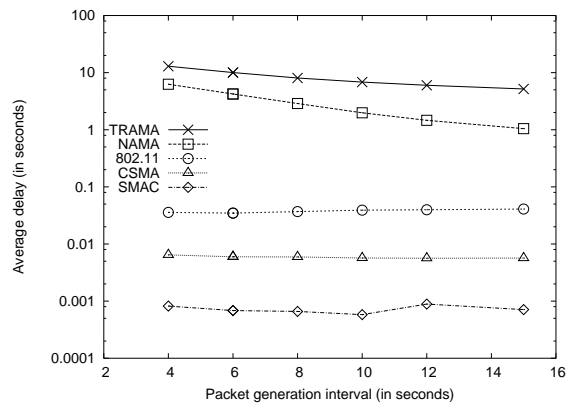
(b) Average queuing delay

Figure 11. Corner Sink

Figures 11(b), 12(b), and 13(b) show the average per-hop delay for all the protocols for the sensor scenarios. Contention-based protocols outperform scheduling-based protocols in terms of delay. This is due to the latency introduced by random scheduling. Finally we show the percentage sleep time achieved by TRAMA and S-MAC in Figure 14(a). The percentage of time nodes can be put to sleep increases with decrease in traffic load. The percentage sleep time is quite high (as high as 85%) for the edge sink scenario which has the lowest load. Again the average length of the sleep interval is also the highest for this case. This clearly shows the benefit of TRAMA's traffic adaptability when compared to S-MAC. The average sleep interval



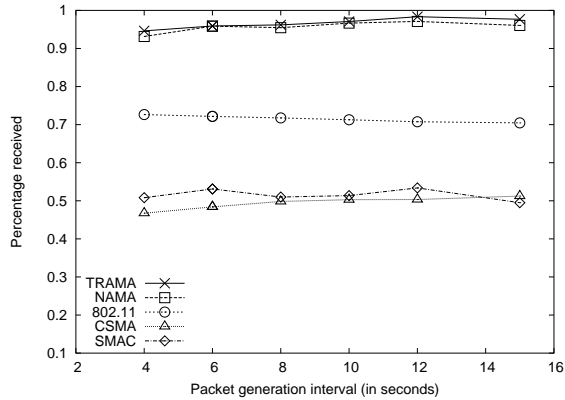
(a) Average delivery ratio



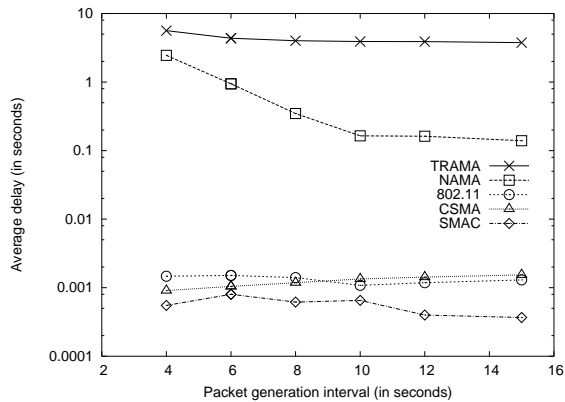
(b) Average queuing delay

Figure 12. Center Sink

of TRAMA is significantly higher than that of S-MAC. When compared to the edge sink scenario, the percentage sleep time is less for the center and corner sink scenarios due to the increased load. In the corner sink case, data forwarded by the nodes which are closer to the sink is heavier than data forwarded by nodes farther away. This reduces sleep time for these nodes and hence the overall percentage sleep time is lesser than the case where the sink is in the center. This also applies to the average length of sleep period shown in Figure 14(b).



(a) Average delivery ratio



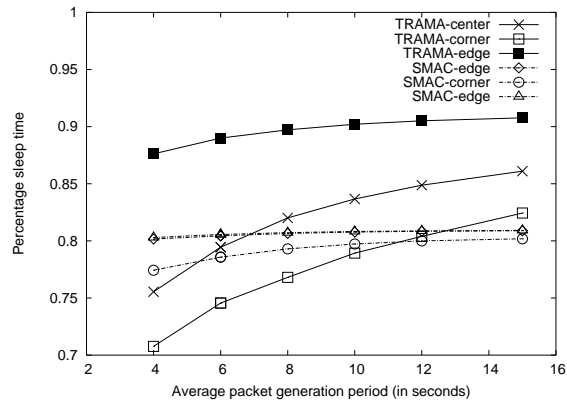
(b) Average queuing delay

Figure 13. Edge Sink

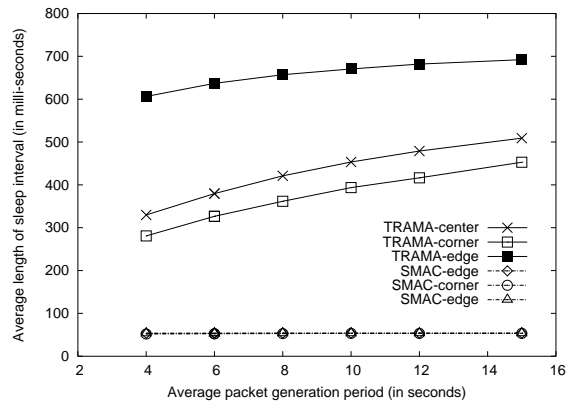
7. Conclusion and Future Work

In this paper we presented TRAMA, a new energy-aware channel access protocol for sensor networks. TRAMA uses traffic-based scheduling to avoid wasting slots when nodes do not have data to send and to switch nodes to a low-power standby radio mode when they are not intended receivers of traffic.

Through extensive simulations, we compared TRAMA's performance against a number of contention- and a scheduled-based MACs. It is evident from the simulation results that significant energy savings (since nodes can sleep for up to 87% of the time) can be achieved by TRAMA depend-



(a) Percentage Energy Savings



(b) Average Sleep Interval

Figure 14. Energy savings and average sleep interval for sensor scenarios

ing on the offered load. TRAMA also achieves higher throughput (around 40% over S-MAC and CSMA and around 20% over 802.11) when compared to contention-based protocols because it avoids collisions due to hidden terminals.

In general, scheduled-based MACs exhibit higher delays than contention-based MACs. In the case of TRAMA, the delay is higher than random-selection protocols (e.g., NAMA) due to the scheduling overhead. We presented an analytical model to quantify the delay for scheduling-based MACs and verified the model by simulations. TRAMA is well suited for applications that are not delay sensitive but require high delivery guarantees and energy efficiency.

A typical example is sensor networks used for periodic data collection and monitoring applications. Future work focuses on adaptive traffic predictions to improve TRAMA's delay performance.

```

1 Compute  $tx(u)$ ,  $atx(u)$  and  $ntx(u)$ 
2 if ( $u = tx(u)$ ) then
3   if ( $u.isScheduleAnnouncedForTx = TRUE$ ) then
4     let  $u.state = TX$ 
5     let  $u.receiver = u.reported.rxId$ 
6     Transmit the packet and update the announced schedule
7   else if ( $u.giveup = TRUE$ ) then
8     call HandleNeedTransmissions
9   endif
10 else if ( $tx(u) \in N1(u)$ ) then
11   if ( $tx(u).announcedScheduleIsValid = TRUE$  AND  $tx(u).announcedGiveup = TRUE$ ) then
12     call HandleNeedTransmissions
13   else if ( $tx(u).announcedScheduleIsValid = FALSE$  OR  $tx(u).announcedReceiver = u$ ) then
14     let  $u.mode = RX$ 
15   else
16     let  $u.mode = SL$ 
17     Update schedule for  $tx(u)$ 
18   endif
19 else
20   if ( $atx(u)$  hidden from  $tx(u)$  AND  $atx(u) \in \mathbf{PTX}(u)$ ) then
21     if ( $atx(u).announcedScheduleIsValid = TRUE$  AND  $atx(u).announcedGiveup = TRUE$ ) then
22       call HandleNeedTransmissions
23     else if ( $atx(u).announcedScheduleIsValid = FALSE$  OR  $atx(u).announcedReceiver = u$ ) then
24       let  $u.mode = RX$ 
25     else
26       let  $u.mode = SL$ 
27       Update schedule for  $atx(u)$ 
28     endif
29   else
30     call HandleNeedTransmissions
31   endif
32 procedure HandleNeedTransmissions
33 if ( $ntx(u) = u$ ) then
34   let  $u.state = TX$ 
35   let  $u.receiver = u.reported.rxId$ 
36   Transmit the packet and update the announced schedule
37 else if ( $ntx(u).announcedScheduleIsValid = FALSE$  ||  $ntx(u).announcedReceiver = u$ ) then
38   let  $u.mode = RX$ 
39 else
40   let  $u.mode = SL$ 
41   Update the schedule for  $ntx(u)$ 
42 endif

```

Figure 15. Pseudo-code description of AEA

References

1. 'http://www.cs.berkeley.edu/~awoo/smartdust/'.
2. 'Product Specification, <http://www.rfm.com/products/data/tr1000.pdf>'.
3. Bao, L. and J. Garcia-Luna-Aceves: 2002, 'Hybrid Channel Access Scheduling in Ad Hoc Networks'. *Proc. IEEE Tenth International Conference on Network Protocols (ICNP)*.
4. Bao, L. and J. J. Garcia-Luna-Aceves: 2001, 'A new approach to channel access scheduling for Ad Hoc networks'. In: *The seventh annual international conference on Mobile computing and networking 2001*. pp. 210–221.
5. Bertsekas, D. and R. Gallager: 1992, *Data Networks*. Prentice Hall, 2 edition.
6. Chlamtac, I. and A. Farago: February 1994, 'Making Transmission Schedules Immune to Topology Changes in Multi-Hop Packet Radio Networks'. *IEEE/ACM Transactions on Networking* **2**(1), 23–29.
7. Chlamtac, I. and A. Lerner, 'Fair algorithms for Maximal Link Activation in Multihop Radio Networks'. *IEEE Transactions on Communications* **35**(7), 739–746.
8. Cidon, I. and M. Sidi: October 1989, 'Distributed Assignment Algorithms for Multihop Packet Radio Networks'. *IEEE Transactions on Computers* **38**(10), 1236–1361.
9. E.D.Kaplan: 1996, *Understanding GPS: Principles and Applications*. Artech House.
10. Elson, J. and D. Estrin, 'Time Synchronization for Wireless Sensor Networks'. pp. 186–186.
11. Ephremides, A. and T. Truong: April 1990, 'Scheduling Broadcasts in Multihop Radio Networks'. *IEEE Transactions on Communications* **38**(4), 456–460.
12. Feeney, L. M. and M. Nilsson: 2001, 'Investigating the Energy Consumption of a Wireless Network Interface in an Ad Hoc Networking Environment'. In: *IEEE INFOCOM*.
13. IEEE: 1999, 'Wireless LAN Medium Access Control (MAC) and Physical Layer Specifications'. ANSI/IEEE Standard 802.11, 1999 Edition.
14. Intanagonwiwat, C., D. Estrin, R. Govindan, and J. Heidemann: 2001, 'Impact of network density on data aggregation in wireless sensor networks'.
15. Ju, J. and V. Li: June 1998, 'An Optimal Topology-Transparent Scheduling Method in Multihop Packet Radio Networks'. *IEEE/ACM Transactions on Networking* **6**(3), 298–306.
16. Kleirock, L. and F. Tobagi, 'Packet Switching in Radio Channels, Part 1: Carrier Sense Multiple-Access Models and their Throughput-delay Characteristics'. *IEEE Transactions on Communications* **23**(12), 1400–1416.
17. Kleirock, L. and F. Tobagi: 1975, 'Packet Switching in Radio Channels, Part 2: Hidden-terminal Problem in Carrier Sense Multiple Access and the Busy-tone Solution'. *IEEE Transactions on Communications* **23**(12), 1417–1433.
18. Lam, S.: 1980, 'A Carrier Sense Multiple Access Protocol for Local Networks'. *Computer Networks* **4**, 21–32.
19. Ramanathan, S.: 1999, 'A Unified Framework and Algorithm for Channel Assignment in Wireless Networks'. *Wireless Networks* **5**(2), 81–94.
20. Singh, S. and C. Raghavendra: 1999, 'PAMAS: Power Aware Multi-Access protocol with Signaling for Ad Hoc Networks'.
21. Sohrabi, K. and G. Pottie: 1999, 'Performance of a Novel Self-Organization Protocol for Wireless Ad Hoc Sensor Networks'. *IEEE 50th. Vehicular Technology Conference* pp. 1222–1226.
22. title, 'Scalable Networks, <http://www.scalable-solutions.com>'.
23. Tseng, Y.-C., C.-S. Hsu, and T.-Y. Hsieh: 2002, 'Power-saving protocols for IEEE 802.11-based multi-hop ad hoc networks'. In: *Proceedings of the IEEE Infocom*.

24. Woo, A. and D. Culler: 2001, 'A Transmission Control Scheme for Media Access in Sensor Networks'. *ACM/IEEE International Conference on Mobile Computing and Networking (Mobicom) 2001*.
25. Ye, W., J. Heidemann, and D. Estrin: June 2002, 'An Energy-Efficient MAC Protocol for Wireless Sensor Networks'. *IEEE Infocom 2002* pp. 1567–1576.