A Transmission Control Scheme for Media Access in Sensor Networks

Alec Woo
Department of EECS
Computer Science Division
University of California, Berkeley
awoo@cs.berkeley.edu

David E. Culler
Computer Science Division
University of California, Berkeley
and Intel XIS Lab
culler@cs.berkeley.edu

ABSTRACT

We study the problem of media access control in the novel regime of sensor networks, where unique application behavior and tight constraints in computation power, storage, energy resources, and radio technology have shaped this design space to be very different from that found in traditional mobile computing regime. Media access control in sensor networks must not only be energy efficient but should also allow fair bandwidth allocation to the infrastructure for all nodes in a multihop network. We propose an adaptive rate control mechanism aiming to support these two goals and find that such a scheme is most effective in achieving our fairness goal while being energy efficient for both low and high duty cycle of network traffic.

1. INTRODUCTION

Sensor networks are an important emerging area of mobile computing that presents novel wireless networking issues because of their unusual application requirements, highly constrained resources and functionality, small packet size, and deep multihop dynamic topologies. Although many high-level architectural and programming aspects of this area are still being resolved, the underlying media access control (MAC) and transmission control protocols are critical enabling technology for many sensor network applications. These problems are well-studied for traditional computer networks, however, the different wireless technologies, application characteristics, and usage scenarios create a complex mix of issues that have led to the existence of many distinct solutions. It is natural to expect the low-level protocols to evolve again for this new era.

Application behavior in sensor networks leads to very different traffic characteristics from that found in conventional computer networks. The primary function of a sensor network application is to sample the environment for sensory information, such as temperature, and propagate this data back to the infrastructure, while perhaps performing some in-network processing, such as aggregation or compression. The network tends to operate as a collective structure, rather than supporting many independent point-to-point flows. Traffic tends to be variable and highly correlated. Over lengthy periods there may be little activity or traffic, but for short periods the traffic may be very intense. For example, when an abnormal event, such as a fire, is detected, many devices will initiate communication at once. Often, applications will arrange periodic rendezvous so that data can be communicated over many hops while allowing nodes to turn off their radios for lengthy periods. Even in the simplest case, roughly periodic sampling of the sensor field yields correlated bursts, even when the duty cycle is low.

The data that a networked sensor generates for each sample, such as a temperature value, is relatively small and, given the low bandwidth of the radio, data packets are kept small with a typical size around tens of bytes. Multi-path interference and short, irregular transmission range result in unpredictable cell structure, but bidirectional connectivity can generally be achieved.

Much of the traffic moves through the network over several hops, perhaps with some intermediate processing, to points that are connected to a larger processing infrastructure. The network takes on an ad hoc multihop topology comprising many levels, where the connectivity is determined dynamically by how placement and physical environment influence radio propagation and by the discovery algorithm. Interference effects may overreach useful communication cells. At each hop, traffic originating from the local sensor must be merged with route-thru traffic. Often this merging is application specific, but at the very least every node is both a data source and a router. Generally, the amount of route-thru traffic exceeds that of originating traffic.

The capabilities of sensor devices are also very different from traditional nodes in a computer network. These devices have a very limited amount of storage, processing power, and most importantly, energy resources. These limitations certainly impose constraints in the design of the MAC protocol. On these platforms, a typical low power RF radio delivers moderate bandwidth in a single channel at the ISM band. There is little or no dedicated support for carrier sensing, collision detection, and no specific framing or encoding enforced by the hardware, other than basic DC-balance. Fur-

ACM SIGMOBILE 7/01 Rome, Italy

thermore, there are no specific protocol stacks in place to dictate the MAC protocol design. It is roughly the same cost per unit time to listen as to transmit or receive. Every moment the radio is on, it consumes precious power. Thus, a key requirement is to turn the radio off whenever possible. Also, it is important to sense contention or frame packets with a minimal number of bits and to minimize the number of protocol control packets. Furthermore, there is very little buffering available on the node. Typically, a single packet is moving with only a few bits of buffering.

These application and platform characteristics give rise to a new set of metrics. Not only are we interested in high channel utilization, we are interested in communication efficiency in terms of energy consumed per unit of successful communication. Furthermore, fairness is highly desirable. For example, we may want to collect roughly the same amount of temperature data from each deployed sensor in a field to infer the temperature gradient during a fire. Therefore, a fair allocation of bandwidth delivered to the base station from each node over multiple hops is desired. It is not sufficient to share the channel fairly in an individual cell, we would like to achieve a crude level of end-to-end fairness even in a deep and self-organized multihop networks, which may change dynamically and originate data at each intermediate node.

The multihop nature of the network poses four interrelated challenges. First, the originating traffic and the route-thru traffic compete for the same upstream bandwidth. Transmission rate control can potentially be applied to either. Rate control is particularly important around the base station, because traffic from nodes deep in the topology primarily flows to a few "gateway" nodes. Second, a hidden node problem exists, by definition, between every other pair of levels in the network. Thus, it may not be possible to detect contention at the upstream node and a significant loss rate is to be expected. Third, the routing distance and degree of intermediate competition varies widely across the network. Nodes residing farther from the infrastructure face a much higher probability of corruption and possible contention at each hop. Finally, energy is invested in a packet at each hop, so the cost of dropping a packet varies with packet and place. All of these factors make it difficult to achieve fairness through the simple, local algorithms that the platform naturally support. At the very least, fairness is at odds with both energy efficiency and high channel utilization. The easiest way to reduce energy and fill the channel is to only take traffic from the nodes adjacent to the base station, but this hardly provides a valid sample of the overall sensor field.

In short, the characteristics and goals of MAC in sensor networks differ strongly from conventional computer networks. They are dominantly periodic and highly correlated traffic, comprising small packets flowing to base stations in a deep, irregular multihop network with each node seeking to achieve a fair bandwidth allocation to the base station in an energy efficient way over single channel radios.

In this paper we make progress on addressing the array of design tradeoffs for sensor networks by developing an innovative MAC protocol and adaptive transmission rate control scheme for multihop networks in the context of a simplified application scenario on a real low-power networked sensor, as well as in simulation. Section 2 describes the hardware and software platform of our networked sensor development and MAC design. It describes the application scenario, lays out our metrics of fairness and efficiency, and introduces the simulation environment that we use in evaluating our design. Section 3 examines related work on existing MAC protocols and identifies shortcomings relative to the sensor network challenges. Section 4 outlines our proposed MAC and transmission control scheme along with a set of open engineering issues, including the mechanism for carrier sensing, desynchronizing periodic behavior, and backoff scheme for the MAC. For the multihop scenario, we study both the conventional RTS/CTS contention based scheme and a simple adaptive rate control algorithm. Section 5 presents our analysis and evaluation of the different carrier sense multiple access (CSMA) techniques that we study. We conclude that limiting the length of listening, the introduction of random delay in addition to backoff, and phase shift at the application level are necessary. Section 6 compares our adaptive rate control scheme with a conventional contention control scheme in a multihop network scenario. We find that *the adaptive rate control mechanism is the most effective in achieving our fairness goal while being energy efficient for both low and high duty cycle of network traffic.*

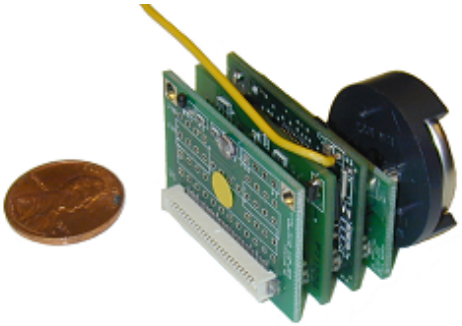

Figure 1: Our low-power networked sensor device prototype.

2. SENSOR NETWORK DESIGN POINT

Our study is grounded in the small, low-power networked sensor device shown in Figure 1 [7]. We believe it be representative of the constraints of limited computation power, storage, and energy supply of the tiny devices that will be deployed into the future [10]. The processor is an ATMEL [3] 4MHz, 8 bit micro-controller with 8K bytes of program memory and 512 bytes of data memory. The radio is a single channel RF transceiver operating at 916MHz and capable of transmitting at 10kbps using on-off-keying encoding. Each radio transition and bit sampling is performed in software. There is no facility for collision detection. A heterogeneous set of sensors such as light, temperature, humidity, pressure, acceleration, and magnetic field can be integrated into these prototypes.

2.1 Networking Component Stack

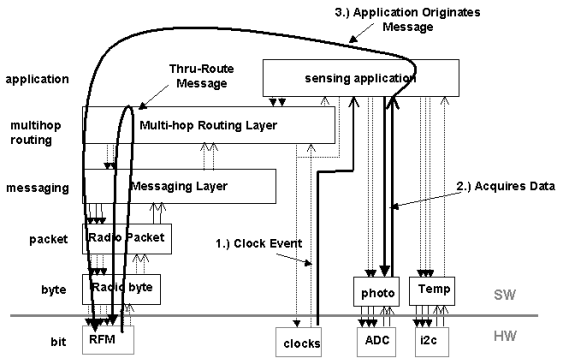

Figure 2: Complete TinyOS application component graph.

TinyOS [7] is an event-based operating system for these devices that provides fine-grained interleaving of event processing and tasks from multiple system components. The complete TinyOS application for our study is shown in Figure 2. There is a component providing an asynchronous interface to each sensor and a stack of components to implement networking over the radio. The lowest layer transmits or receives bytes bit-by-bit over the radio. It provides phase and rate controls to lock on to the packet start symbol and then to spool bits. At this level, the interface is half-duplex - the radio is receiving except during packet transmission. The packet-level component is responsible for spooling incoming bytes and delivering the packet receive event. It is where the media access control mechanisms for transmit reside. (It also performs the encoding and decoding of the byte stream onto the link and error checking: Manchester encoding with an 16-bit CRC.) Packets are short and of a fixed size, typically 30 bytes including an one byte destination field, an one byte handler field, and an application data unit.

The Active Message component delivers tagged packet events to application level components. Here we have two such components. The sensor component periodically receives a clock event, acquires sensor data, and transmits the data toward a base station over the multihop network. The other component is responsible for building the dynamic multihop network and routing traffic. A simple beacon-based discovery protocol maintains a breadth-first spanning tree, such that each node knows a "parent node" closer to the base station. Originating sensor packets are marked for the parent. (All other nodes discard them.) At each hop, the multihop component receives a packet and retransmits it to the upstream level. In general, this component might perform aggregation or statistical analysis. However, we restrict ourselves to the case where it forwards all data to the infrastructure for analysis, as this focuses the work on the media access and transmission control aspects. This component does collect statistics on the number of nodes routing through it. The only buffering in the system is a fixed number of small packet buffers at the application level, one of which is used for the asynchronous receive. Thus, if the radio is busy transmitting or receiving when a packet send is requested, the request will fail back up to the application component. Once the packet component has accepted a packet for transmission, it will work on it until it acquires the channel and transmits it. Thus, the transmission rate control is implemented within the two application components.

2.2 Metric for Evaluation

The metrics for evaluation of a sensor network MAC protocol stress both fairness and energy efficiency. A fair allocation of bandwidth delivered to the base station from each node over multiple hops is desirable. Although aggregate bandwidth is an important metric in evaluating MAC protocols, it can be misleading. High aggregate bandwidth of packets only from nodes around the base station is not desirable.

We evaluate energy efficiency in terms of energy consumed per unit of successful communication or packets received by the base station. The energy consumed is the total energy that the network has invested in propagating data to the base station. A scenario with nodes near the base station filling the channel will perform poorly under this metric. The total energy invested by the network includes energy spent in listening for the channel and all packet transmissions and forwarding. Our experiment is designed such that all nodes are in receive mode even during idle period. Therefore, we only account useful work as energy spent in channel listening and packet transmission under this metric.

2.3 Simulation Environment

Given the difficulty in performing actual measurements in wireless networking, we first evaluate our system through simulation. We have created a simple simulator capable of creating an arbitrary multihop network topology of a group of networked sensors. Each UNIX process represents a networked sensor, and a master process is responsible for synchronizing them to perform bit time simulation. There is a simple radio propagation model in the simulation and we assume bit error rate to be zero, since our main focus is media access control. We use a simple reachability table that specifies whether bidirectional connectivity exists between nodes in the simulated network. The simulator doesn't simulate the actual hardware operating in the TinyOS environment. However, it preserves the event driven semantics and the dynamics of traffic flow shown in Figure 2. All the simulations presented in Section 5 and 6 are collected in this simulation environment.

3. RELATED WORK

In designing the MAC protocol, we first examine existing mechanisms and determine whether or not they apply in the regime of sensor networks.

Many variations of the CSMA [8, 17] strategy can be found in the literature. Listening to the channel before transmission to exploit information about other users is a very common approach found in almost all CSMA schemes, except pure ALOHA [15]. Another approach is to use explicit positive or negative acknowledgments to signal collision and perform necessary random delay before retransmit. ALOHA takes this approach. IEEE 802.11 [2] uses it in addition to listening. Other CSMA schemes rely on time synchronized slotted channel, such as Slotted ALOHA [11]. Finally,

CSMA with collision detection [9], which is widely used, including wired Ethernet. Though many CSMA schemes exist, they all lean toward a fundamental assumption that packet transmissions occur with a stochastic distribution, that is very different from the correlated traffic found in sensor networks. Furthermore, they aim to support many independent point-to-point flows while the network in this new regime tends to operate as a collective structure. As a result, re-exploring CSMA strategies with a different fundamental assumption is extremely relevant.

Energy consumption is an important metric in the design of media access control. PAMAS [13, 14] is a power aware media access protocol which powers off radio when not actively transmitting or receiving packets. Our work does not explore this power scheduling aspect. Instead, we focus on the energy efficiency in basic media access control schemes, and the overall bandwidth, energy, and fairness tradeoff in a multihop network.

The IEEE 802.11 [2] standard aims to provide a wireless Ethernet illusion. The design is based on an assumption of a single cell scenario, with mobile stations always in range of at least one base station, with a hand-off when migrating from one cell to another. As a result, there is no multihop scenario. Furthermore, the large transmission range and the spread spectrum capability of the radio make it feasible to create a grid of overlapping cells to cover a field without need for taking several hops to reach the base station. The ad hoc aspect of the protocol assumes peer-to-peer communications rather than many-to-one data propagation scenario as found in sensor network. Nevertheless, the primary mechanisms of carrier sensing and contention control scheme are fundamentals that we will study and evaluate.

Bluetooth [1] is an emerging standard for many future wireless devices. Its usage model is to create a "wireless cable" illusion for applications like connecting a cellular phone or speaker to a notebook computer with voice or data streaming among these devices. The primary media access control is a centralized Time Division Multiple Access (TDMA) protocol within a piconet which is a relatively static ad hoc network supporting a small number of nodes within a single cell. Overlapping cells is also feasible due to the spread spectrum radios. Bluetooth assumes no multihop scenario. The centralized TDMA protocol and the tight requirement of time synchronization between each node in the piconet make it inappropriate for sensor networks.

MACAW [4] shares a single channel radio similar to ours, but its main focus is for single hop base station interaction within a cell. In fact, in one of its communication scenarios resembling a hidden node problem in a multihop network, it explicitly states that the scenario cannot be solved by contention control protocols unless time synchronization information is present for contention period to be known. Nevertheless, we will study the most primitive contention control protocol in the context of multihop network, and in Section 4, we will discuss how hidden node problem can be addressed.

In large, dense packet radio networks [12], a collision free channel access scheme is accomplished using locally generated and published transmit and receive schedules. A station can send if its sending slot overlaps with the receiver's receiving slot and the amount of overlap is long enough for data transmission. To prevent synchronization between neighboring stations, they produce random or pseudo-random schedules. This is a distributed TDMA scheme with the assumption that the network topology is relatively static. Given the spontaneous nature of sensor networks where the receiver or the next hop to the base station may change at any time, such a scheme may be very inefficient in publishing different schedules.

Prior work shows that an adaptive rate control algorithm is very effective in achieving proportional fairness of media access for packet radio [16]. Their design goal is to have a fair sharing of the channel among local competing neighbors. Our proposed adaptive rate control scheme builds upon this work, but our goal is to have media access control assist in achieving fair bandwidth delivery to the base station for nodes in a multihop network.

Our adaptive rate control uses loss as collision signal to adjust transmission rate in a manner similar to the congestion control used in TCP [5, 6]. While TCP's congestion control is end-to-end over a network with many independent flows, our proposed adaptive rate control works collectively at every node in the network, since each node is both a router and a sender, and routing is done at the application level.

4. DESIGN

We discuss how media access control for sensor network should be done differently in this section. First, we explore what type of listening mechanism is appropriate for the case where all nodes can hear each other. Second, we discuss how backoff should be implemented in a sensor network. Third, we present two mechanisms which we will study their effectiveness in the context of a multihop network. The first scheme is a conventional RTS/CTS contention control scheme, and the second one is our proposed adaptive transmission control scheme, and finally, a mechanism that all schemes can leverage off for avoiding some cases of hidden node problems in multihop network.

4.1 Listening Mechanism

Carrier Sense Multiple Access (CSMA) and the Collision Detection (CD) scheme found in Ethernet are examples of listening mechanisms. Listening is very effective when all nodes can hear each other, (i.e. without hidden nodes). Unfortunately, collision detection is not possible in wireless network technology without additional circuitry. Though listening is simple, it does come with an energy cost, because the radio must be on to listen. To conserve energy, it is important to shorten the length of carrier sensing. Many protocols such as IEEE 802.11 require sensing the channel even during backoff. However, CSMA for sensor networks should take this opportunity to turn the radio off.

The highly synchronized nature of the traffic imposes a new criteria for CSMA. Given there are no hardware mechanism for detecting collisions, nodes that happen to send at the same time will corrupt each other. If the traffic pattern of each node is independent, this situation is not likely to repeat. However, detection of one common physical event will

synchronize these nodes and lead them to send at the same time, which repeats periodically. The result is no packet transfer at all. The solution is to introduce random delay for transmission to unsynchronize the nodes. Section 5 will discuss and evaluate various ways in introducing randomness for CSMA.

4.2 Backoff Mechanism

Backoff is a widely used mechanism in media access control to reduce contention. The idea of backoff is to restrain a node from accessing the channel for a period of time and hopefully, the channel will become free after the backoff period. In the case of sensor networks where the traffic is a superposition of different periodic streams, backoff should not just restrain a node from sending for the backoff period. In fact, the backoff period should be applied as a phase shift to the periodicity of the application so that synchronization among periodic streams of traffic can be broken.

4.3 Contention Based Mechanism

Explicit contention control schemes, which are widely used in many MAC protocols, e.g., IEEE 802.11 [2] and MACAW [4], require the use of control packets, such as Request to Send (RTS) and Clear to Send (CTS). Acknowledgments (ACKs) serve a different purpose in IEEE 802.11; they indicate lack of collision. For computer networks where packets are large, these small control packets impose very little overhead. However, for sensor networks where packet size is small, they can constitute a large overhead. A RTS-CTS-DATA-ACK handshake series in transmitting a packet can constitute up to 40% overhead in our platform. (Each control packet is 3 bytes long (type,destination,source) and the packet is 30 bytes long.) This can be extremely costly, since energy has to be spent in CSMA, transmitting, and receiving each control packet. One advantage of a bidirectional multihop network is that acknowledgments are free when the receiving node (your parent in the multihop topology) routes the packet to its parent. This eliminates an explicit ACK control packet. If the receiver performs some kind of application specific aggregation before routing the packet, the originator of the packet may still be capable of detecting the success of the transmission.

A contention control scheme for sensor networks should use a minimum number of control packets. The most basic types are RTS and CTS. Though it may be effective in solving the hidden node problem in a multihop network, such a scheme should only be used if the amount of traffic is high while a simple CSMA scheme is actually adequate for low traffic since the probability of corruption due to collision is very small.

For the contention scheme that we study in Section 6, only RTS and CTS packets are used for handshakes. A node wishing to transmit first sends a RTS packet to its parent and waits for a CTS reply. If no CTS is received for a timeout period (2 CTS packet times), the node will enter backoff with a binary exponential increasing backoff window. Similarly, if it receives a CTS not destined to it, it will also go into backoff. If no CTS has been received after five retries, the transmission will be dropped. Furthermore, if a node hears a CTS before any of its own transmission, it will defer transmission for one packet time to avoid corrupting the traffic.

4.4 Rate Control Mechanism

The tension between originating traffic and route-thru traffic has a direct impact in achieving our fairness goal. Media access control must assist in balancing this tension for the channel. Specifically, the MAC should control the rate of originating data of a node in order to allow route-thru traffic to access the channel and reach the base station. Similarly, some kind of progressive signalling mechanism should exist for route-thru traffic, such that back pressure can propagate deep down into the network for those nodes to lower their rate of originating data. This in turn will decrease the aggregate route-thru traffic and open up the channel for nodes closer to the base station to originate data. Such a rate control mechanism should only use distributed local algorithms. We propose an implicit mechanism which passively adapts the rate of transmission of both original and route-thru traffic without the use of any MAC control packets.

The adaptive rate control idea is very simple and can be explained with an analogy of metering traffic onto a freeway where the route-thru traffic is like traffic on the freeway and each node originating data is like cars trying to enter. Periodically, a node attempts to inject a packet. If the packet is successfully injected, it becomes part of the route-thru traffic. As it is routed by the node's parent, it signals that the road still has capacity for more traffic and thus, the node can increase its transmission rate. However, if the injection of the packet wasn't successful, it signals that the road is jammed and the node decreases its rate of originating data and backoff to achieve a phase change effect.

The above explains how the originating data rate adapts to the route-thru traffic. Route-thru traffic will adapt to the traffic of original data using a similar mechanism. If a node injects lots of original traffic into the freeway, the route-thru traffic will be hindered and thus, the rate of transmitting route-thru traffic will decrease (cars at the back have to slow down for cars in front once they are on the bridge). It has a domino effect in propagating this back pressure deep down into the network which ultimately decreases the amount of aggregate route-thru traffic.

The metering effect discussed in the analogy above can be set by a global schedule. Given that each node has an omniscient knowledge of the total number of nodes N in the entire network, each node can meter its own rate by $ChannelCapacity/N$ to achieve our fairness goal. However, the spontaneous ad hoc nature of sensor networks make such a global knowledge impractical. Therefore, we propose an adaptive scheme attempting to approximate it.

Our rate control mechanism uses a linear increase and multiplicative decrease approach to control the transmission rate of the application. Given the application transmission rate is S, the actual rate of originating data is $S*p$ where $p \in [0, 1]$. This rate control is probabilistic, where p is the probability of transmission. To linearly increase the rate, simply increment p by a constant α. To multiplicatively decrease the rate, multiply p by a factor β where $0 < \beta < 1$. This probabilistic mechanism of rate control is based on the work done

in [16]. In general, a large α tends to be aggressive in competing for the channel. β controls the penalty given a failure of transmission. The choice of α and β will be investigated and discussed in Section 6.

Dropping route-thru traffic is considered a waste of network resource, and preference is given to route-thru traffic by making its penalty to be 50% less (i.e., $\beta_{route} = 1.5 * \beta_{originate}$). Furthermore, a node should give a fair proportion of its bandwidth to each node routing through it. If a node has route-thru traffic from n children, the bandwidth for its original data should be $1/(n+1)$ in order to achieve our fairness goal. Since a node can estimate n by monitoring the route-thru traffic, $\alpha_{orignate}$ can be set to equal to $\alpha_{route}/(n+1)$. With this choice, the adaptive scheme has only two parameters, α and β.

The amount of computation for this adaptive scheme is small and within networked sensor's computation capability. It requires a simple pseudo random number generator and a few addition and divide operations. Furthermore, the scheme is totally computational, which is much cheaper in energy cost than operations on the radio.

4.5 Multihop Hidden Node Problem

The adaptive transmission control scheme attempts to avoid hidden node problem without explicit control packets by constantly tuning the transmission rate and performing phase changes, so that the aggregate periodic streams of traffic will not repeatedly collide with each other. The contention control RTS/CTS scheme can solve the hidden problem to some degree. However, as discussed in Section 3, MACAW has suggested a scenario where multihop hidden node problem cannot be solved due to lack of synchronized information of knowing when is the contention period between a child node and its grandparent's node.

If we assume that packets will be routed after some processing time x, a child node is able to avoid a potential hidden node problem with its grandparent. The idea is that if a child node hears the end of its parent's transmission at time t, it should expect that its grandparent will route its parent's packet starting at time $t + x$. Therefore, if the child node can restrain from transmitting from time t to $t + x + PACKETTIME$, the hidden node problem can be reduced. In fact, if the child node detects such situation it should perform a backoff to change its phase such that it will not encounter the same situation the next time it transmits.

5. ANALYSIS OF CSMA SCHEMES

In this section, we use both simulated and empirical measurements to explore the various ways of performing CSMA, and study their performance based on our energy efficiency metric, as well as traditional metrics, such as channel utilization and fairness.

Traditional CSMA schemes have two basic design parameters: the carrier sense (or listening) mechanism and the backoff mechanism. In sensor networks we optionally add a random delay prior to listening to avoid repeated collision due to synchronized behavior. With this new delay component, the CSMA algorithms that we study can be expressed using a regular expression like syntax as $DELAY$ $[LISTEN(busy)BACKOFF]* LISTEN(idle)TRANSMIT$, where $*$ means zero or more occurrences.

CSMA Schemes	Random Delay	Listening Time	Backoff Mechanism
ND_RAND	No	Random	None
ND_RAND_FIX	No	Random	Fixed Window
ND_RAND_EXP	No	Random	Exp Increase
ND_RAND_REVEXP	No	Random	Exp Decrease
ND_CONST_FIX	No	Constant	Fixed Window
ND_CONST_EXP	No	Constant	Exp Increase
ND_CONST_REVEXP	No	Constant	Exp Decrease
D_CONST_FIX	Yes	Constant	Fixed Window
D_CONST_EXP	Yes	Constant	Exp Increase
D_CONST_REVEXP	Yes	Constant	Exp Decrease

Table 1: Different CSMA schemes with variations on the amount of delay period, listening period, and ways of backoff. All exponential backoffs are binary exponential.

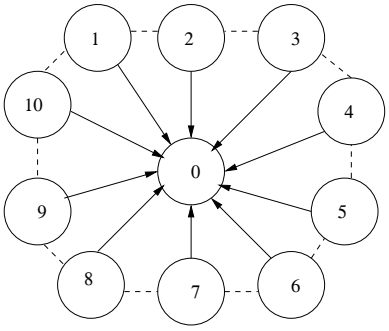

Figure 3: Network topology for the evaluating the CSMA scheme.

Table 1 shows the $DELAY$, $LISTEN$, and $BACKOFF$ components we consider. The listening period can be random over a fixed interval or constant. Backoff time is random drawn from a fixed window, binary exponentially increasing window, or binary exponentially decreasing window. We evaluate all these schemes and compare them with our tiny implementation of the well known 802.11 CSMA/ACK protocol over a single cell scenario shown in Figure 3. All nodes in the cell are able to hear each other, with node 0 being the base station.

5.1 Simulation Settings

To make the simulation approximate our real platform, the packet size is set to 30 bytes, which is the actual packet size used in many networked sensor applications on our prototype. With 30 byte packets in Manchester Encoding, the 10kbps channel capacity can deliver at most 20.8 packet/s. We use a 16-bit CRC error detection mechanism to check for corrupted packets. The specific the values of all the necessary parameters for the CSMA schemes in Table 1 are given in Table 2.

5.2 Delivered Bandwidth under Simulation

The average aggregate bandwidth received at the base station for each scheme in Table 1 is shown in Figure 4. The simulation is set up to have each node attempting to send

Parameter	Value
Constant Listen Window	7 bit time
Random Listen Window	64 bit time
Random Delay Window	64 bit time
Fixed Backoff Window	2400 bit time
Binary Exponential Backoff Window	480 - 7680 bit time
802.11 SIFS	7 bit time
802.11 DIFS	14 bit time
802.11 Backoff Window	480 - 7680 bit time

Table 2: Values for the parameters of different CSMA scheme used in simulation. Since the radio's raw bandwidth is 10kbps, one bit time is 100us.

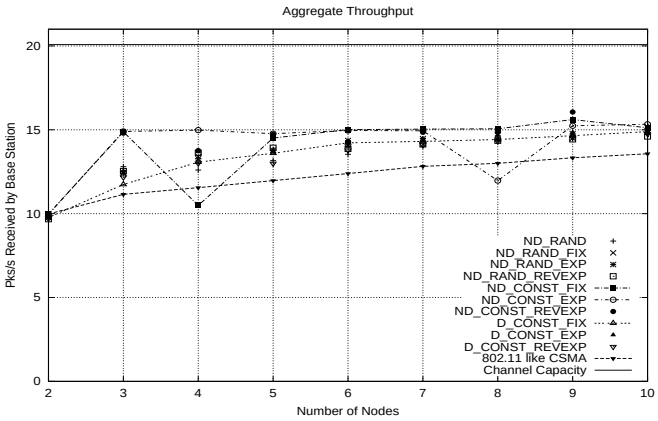

Figure 4: Aggregate delivered bandwidth for CSMA schemes with typical start conditions and offered load of 5 packets per second per node.

periodically at 5 packet/s with slightly different start times. Since the channel capacity is 20.8 packet/s, the traffic load will exceed capacity when more than 4 nodes are sending. All the variants of our simple CSMA schemes achieve greater bandwidth than the 802.11 scheme with its explicit ACKs. The three schemes with constant listen period and no random delay achieve highest bandwidth, especially where the load is just below the channel capacity (3 to 4 nodes). However, their aggregate bandwidth is not very robust, as indicated by the two dips in the figure. The dips are caused by repeated collisions, which these schemes are incapable of eliminating. The remaining schemes, with random delay or random listening intervals achieve slightly less bandwidth, but are more robust. As network load exceeds the channel capacity, all schemes except 802.11 utilize about 75% channel capacity, or 15packet/s of aggregate bandwidth.

The randomness introduced by the backoff mechanism may seem to be sufficient to avoid repeated collisions, however, without collision detection hardware greater attention must be paid to the listen phase. To focus on the robustness of the CSMA schemes, Figure 5 shows a worst case scenario where all nodes are synchronized to start transmitting at the same time. The three schemes with constant listening window and no random delay delivers zero bandwidth. With synchronized listen periods, the nodes fail to detect the collision, so the backoff mechanism is never triggered. The 802.11 scheme starts at 5 packet/s and delivers less bandwidth than in Figure 4 throughout. Although 802.11 has no random delay and constant listen period, the ACKs provide collision detection and trigger the backoff mechanism, desynchronizing the nodes. The rest of the curves correspond to the other seven schemes that have randomness either in delay or listening period. They achieve similar performance in both figures.

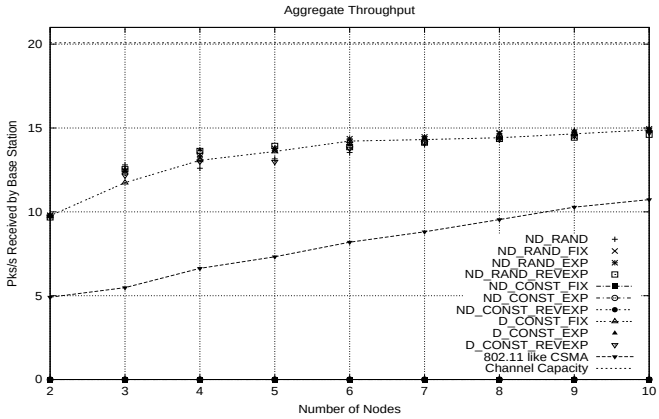

Figure 5: Aggregate Bandwidth for the different CSMA schemes with nodes that perform no random delay to begin transmission at different time.

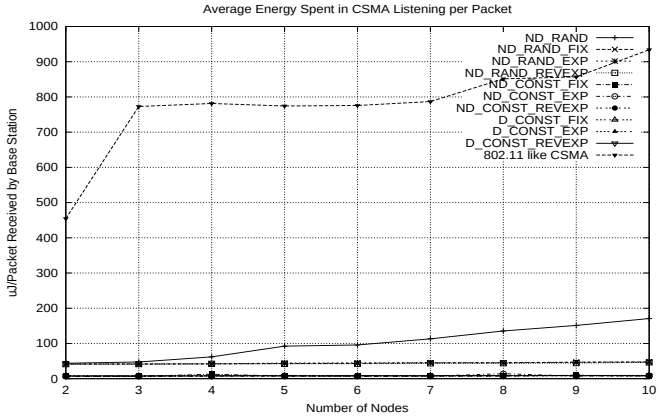

Figure 6: Average energy spent in listening the channel per unit of successful packet received by the base station.

In summary, all our simple CSMA schemes achieve good channel utilization and withstand high offered load. The performance is almost insensitive to backoff mechanism. Even the ND_RAND scheme, which has no backoff mechanism, performs as well as the other six schemes. However, randomness in the pre-collision phase is essential for robustness.

5.3 Energy Usage

In examining the energy consumed in communication, we separate the portion spent in actually transmitting and receiving packets from that spent listening. The former is determined primarily by the traffic load; differences result

from the happenstance of packets being dropped. The latter is primarily determined by the CSMA protocol. Figure 6 shows the average energy spent per packet in listening on the channel for the different CSMA schemes, assuming the energy model described in Section 2.2.

802.11 has the worst energy efficiency. Although 802.11 uses a constant listen period (DIFS and SIFS), listening on the channel throughout the backoff period makes it energy inefficient. The efficiency of 802.11 is even worse if we account for the energy spent in ACK transmission and reception.

Our CSMA schemes with constant listen period are the most energy efficient, at approximately $10uJ/packet$ independent of network size. Schemes with random listen period are more costly, at $40uJ/packet$. Closer examination indicates that the average number of backoffs per successful transmission is roughly constant, so this difference reflects an increase in the average listen time. The random listen period is drawn from a 64-bit window, although it drops out at the first collision bit. The energy cost of ND_RAND increases with network size. Under low traffic load, it performs the same as the other random listening period schemes. However, as the traffic exceeds channel capacity, with no backoff, the number of listens per packet increases.

Delay uses essnetially no energy, since the radio is off during that period. The most energy efficient schemes are those with constant listen period and a random delay provides robustness. Thus, for the rest of the analysis we focus on the three CSMA schemes that use random delay and constant listen period.

5.4 Fairness

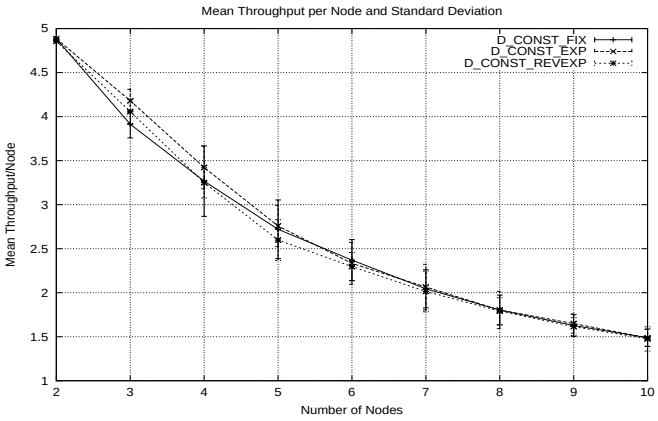

Figure 7: Mean throughput per node delivered to the base station and standard deviation as a measure of fairness for the different CSMA schemes.

Figure 7 and Figure 8 show the deviation of mean throughput per node among the three CSMA schemes and 802.11 as an indication of fairness. The three CSMA schemes are very similar. Their standard deviations are approximately 0.25 packet/s and tend to decrease as traffic increases. Thus, we conclude that the difference in backoff mechanism is insignificant in terms of fairness at uniform load.

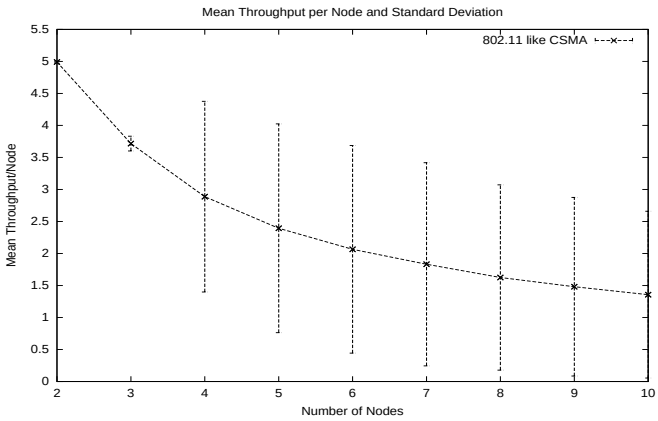

Figure 8: Mean throughput per node delivered to the base station and the corresponding standard deviation as a measure of fairness for the different CSMA schemes.

On the contrary, 802.11, shown in Figure 8 gives an unfair allocation of bandwidth among the nodes, with a standard deviation of more than 1 packet/s when the traffic exceeds the channel capacity. From the 802.11 data that are not shown in the figures, we found that nodes, which have an earlier transmission start time, end up capturing the channel and result in this unfairness. We will discuss this capturing effect later in this section.

Proportional fairness is another interesting metric we should consider. In a multihop network, where thru-route traffic is likely to exceed a node's own traffic, disproportional send rate is common. We would like to observe that given the existence of thru-route traffic, is proportional fairness in a single cell likely to be maintained by these CSMA schemes.

We assume the single cell in Figure 3 is part of a multihop network with node 0 routing all packets that it receives to the next hop. Nodes are set to send at different rates to create an uneven amount of traffic. Table 3 lists the send rate of each node and the resulting bandwidth allocation among different nodes for each scheme in Table 1. All data is normalized to the bandwidth of node 10 in order to observe the relative proportion of bandwidth allocation. The ideal case is for node 1,2,and 3 to send 500% more traffic than node 10, node 4,5,and 6 to send 250% more, and node 8 and 9 to send the same amount traffic as node 10. The result suggests that backoff mechanism has an effect on proportional fairness, with binary exponential increasing backoff being the worst. The data also shows that 802.11, as compared to conventional wireless network, performs worst in proportional fairness in this regime.

5.5 Sensor Phase Shifting

The half duplex nature of the networking stack makes CSMA vulnerable to the capturing effect. That is, during reception of neighboring node's transmission, the network stack will not start any transmissions issued by the application. Instead, it will fail the transmission back to the application. If the two nodes remain in synchrony with one starting its

Node	Send Rate (packet/s)	802.11	D_CONST _FIX	D_CONST _EXP	D_CONST _REVEXP
1	10	215%	698%	981%	692%
2	10	199%	598%	933%	676%
3	10	205%	620%	1028%	627%
4	5	258%	198%	297%	227%
5	5	248%	200%	294%	197%
6	5	246%	188%	325%	192%
7	2	14%	92%	164%	116%
8	2	91%	100%	111%	103%
9	2	90%	114%	128%	97%
10	2	100%	100%	100%	100%
		bad	good	bad	good

Table 3: This table illustrates the proportional fairness for each node sending at different rate shown in column 2 for each CSMA scheme listed. The data is normalized as an percentage to the send rate of node 10.

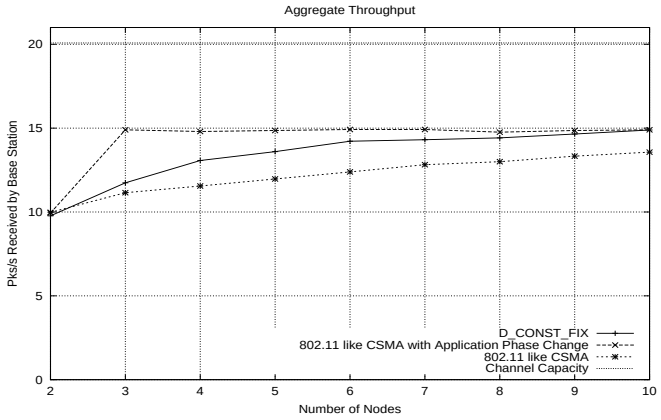

Figure 9: Aggregate throughput comparison between 802.11 CSMA and 802.11 CSMA with application phase change.

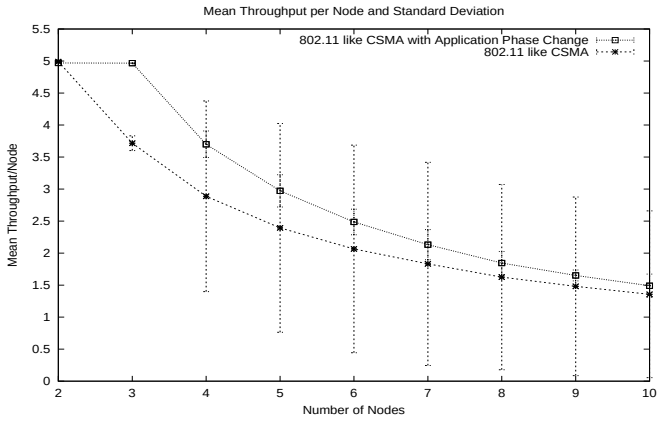

Figure 10: Fairness comparison in terms of mean throughput standard deviation between 802.11 CSMA and 802.11 CSMA with application phase change.

transmission before the other, the channel may remain captured even though traffic load is extremely low. It appears that this capturing effect causes the behavior in Figure 8 for the 802.11 case.

Our CSMA mechanisms include an application level adaptation, where the phase of the sensor sampling interval is shifted by a random amount in response to transmission failure. This provides a way to break away from unfortunate synchrony, which both CSMA listening and backoff mechanisms fail to detect. The amount of phase change is totally application dependent. In our simulation, a phase change corresponds to a random delay bounded by the window of the application's transmission period.

When this phase-shift is incorporated into the 802.11 scheme, bandwidth and fairness improve substantially, as shown by Figure 9 and Figure 10. In Figure 9, the augmented 802.11's aggregate bandwidth reaches 75% channel capacity under traffic load and it reaches this level quickly. Figure 10 shows that the deviation in mean throughput per node is reduce to 0.25 packet/s. Proportional fairness is also comparable to the best CSMA scheme, however, energy efficiency shows little change from the base 802.11 case.

5.6 Empirical Results

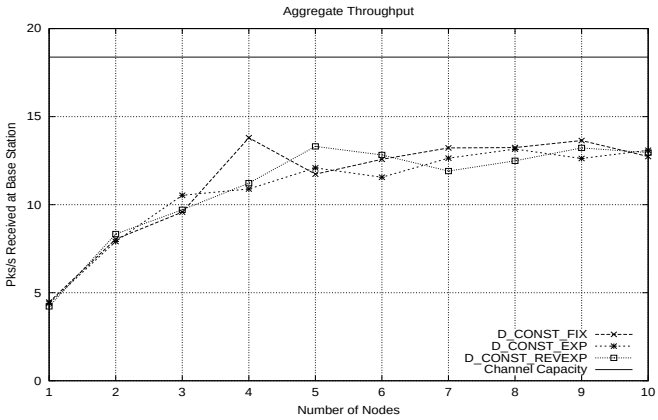

Figure 11: Actual aggregate bandwidth obtained over our network sensor prototypes for the different CSMA schemes. (Empirical)

We have implemented the CSMA schemes on our networked sensor platform, and we compare here the three CSMA schemes with random delay to the simulation result. The nodes are placed to provide the the topology illustrated in Figure 3.

With each node sending at a rate of 5 packet/s, Figure 11 shows the resulting aggregate throughput. The empirical measurement closely matches the simulation prediction, with aggregate bandwidth reaching 70% of channel capacity. Figure 12 shows the average energy spent per packet in CSMA listening phase, and it also agrees with the prediction of around $10uJ/packet$. Finally, Figure 13 shows the fairness comparison. The deviation among the three schemes vary from 0.3 packet/s to 0.5 packet/s, and the deviation decreases as the network traffic increases. Given the uncertainties in the actual measurement, these results match very

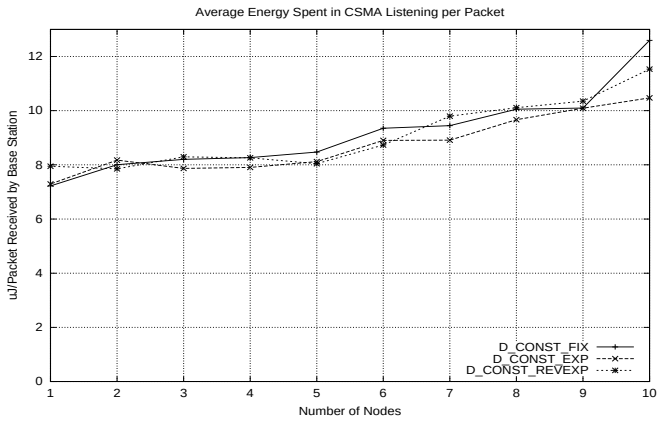

Figure 12: Actual energy spent in CSMA listening over our network sensor prototypes for the different CSMA schemes.(Empirical)

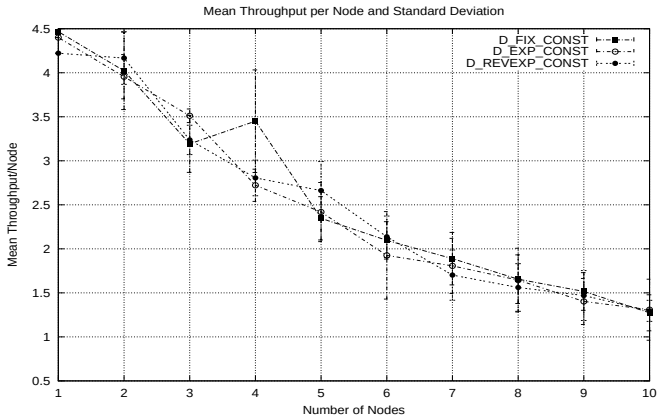

Figure 13: Fairness comparison in terms of deviation of mean throughput per node. (Empirical)

well with the prediction from the simulator.

Our simulation result in Figure 5 shows that in the worst case when correlated nodes transmit at the same time, no successful transmission is possible if no randomness exists in either delay or listening. However, in practice, one may expect that the many source of noise would break the synchronization among nodes. However, we find that if we start the nodes synchronized down to one bit time, the receivers still receive corrupted packets due to repeated collisions if no explicit randomness is added.

All in all, given that energy efficiency and fairness are our main metrics, the *CSMA mechanism should incorporate random delay and a constant listen period with radio powered down during backoff period.* If the software platform uses a half-duplex networking stack, phase-shifting at the application level is important. Furthermore, the backoff mechanism has an affect on maintaining proportional fairness, but plays no significant role in terms of aggregate bandwidth and fairness. Our results suggest that backoff with a fixed window size or binary exponential decrease in window size are effective in maintaining proportional fairness.

6. ANALYSIS OF MULTIHOP SCENARIO

This section extends our analysis to multihop networks where two essential challenges are present. First, if nodes near the base station originate too much traffic, little will bandwidth will be available for more distant nodes. Second, if distant nodes collectively originate more traffic than is available as the flows approach the base station, packets will be dropped and the effort in routing them will be wasted. The CSMA scheme developed for media access control is augmented with a transmission control protocol so that nodes adapt their data origination rate to give a fair share to downstream nodes and to match available upstream bandwidth. We consider both a traditional RTS/CTS contention control scheme and an adaptive rate control (ARC) scheme that avoids explicit control packets. Like TCP, it adjusts its rate based on observed packet loss. However, in this multihop scenario, each node is both source and router, so rate control can be applied to either flow. For comparison, the base CSMA and 802.11 schemes are carried forward to multihop networks, as well. As with the single hop CSMA, we begin with simulations and study real implementations for the most viable options.

6.1 Reference Topology

We focus our multihop study on a single topology consisting of 11 nodes and one base station, as shown in Figure 14. Bidirectional connectivity is represented by edges. Nodes are hidden from each other if they are not linked by an edge. This topology is sufficient to reveal many of the general challenges, including variations in depth, flow rates, and loading. Constructing this topology by physical placement of the nodes is challenging, especially with automatic route discovery, given the variability in cell shapes. Thus, to better control the implementation experiments, we fix the routing to match the simulated topology and place the nodes to give a reasonable approximation of the desired cell coverage. Still there is significant connectivity and interference not present in the simulation.

The ideal available bandwidth in such a multihop network is much lower than the capacity of a single cell, since each packet occupies the parent for three packet times: it is sent to the parent, sent by the parent, and sent by the grandparent. The ideal uniform data origination rate for each node, X, is given by the cell bandwidth divided by the amount of traffic in the busiest cell. For the multihop scenario shown in Figure 14, the channel capacity of node 2 will determine the ideal transmission rate of the overall network. If each node sends at a rate of X packet/s, the total traffic in the cell of node 2 is as shown in Table 4. Thus, under simulation, the maximum uniform origination rate is $20/24 = 0.83$ $packet/s$. In the implementation, this limit is $15.7/24 = 0.66$ $packet/s$ because we have used an encoding scheme that permits SECDED (Single bit Error Correction and Double bit Error Detection), rather than using Mancester encoding with CRC, as in the CSMA studies. In either case, the available bandwidth through the bottleneck is a small fraction of the 4 $packet/s$ load that each node offers. The transmission control protocol tries to match the offered load to available bandwidth.

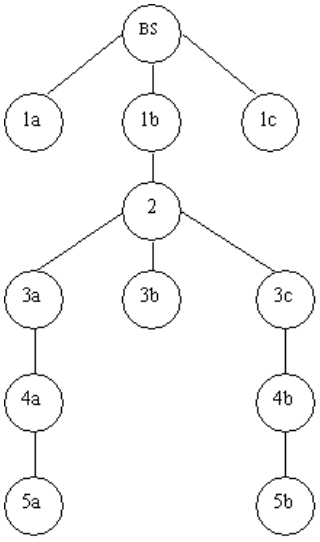

Figure 14: Five level deep reference multihop network with 11 nodes. Edges represent bidirectional connectivity.

Node	Number of Children	Transmission Rate
1b	9	9X
2	8	8X
3a	3	3X
3b	1	X
3c	3	3X
Total		24X

Table 4: Transmission rate breakdown of each node in the cell of node 2 in the multihop scenario with each node sending at the ideal transmission rate of X packet/s

6.2 Simulation Measurements

The simulation runs with each node sending packets to the base station at rate of 4 packet/s with the same start time. The base station will echo each packet it receives in all schemes in order to make a fair comparison.

Figure 15 shows the bandwidth delivered to the base station from each node in the network for four different schemes. D_CONST_FIX and 802.11 CSMA with ACK include only the MAC mechanism, as in Section 5. RTS/CTS uses an explicit contention control scheme like 802.11 in ad hoc mode, but over our energy efficient D_CONST_FIX CSMA scheme. ARC uses this same MAC, but with the TCP-like adaptive rate control.

The two basic CSMA schemes fail to deliver any packets from nodes which are more than two level deep. This is expected since these mechanisms are not designed to cope with the hidden node problems that occur between levels of the topology. Note that node 2 is able of send more than all nodes in the first level. Since node 2 can hear 1b's traffic and traffic from 1a and 1c is synchronized with 1b, by avoiding collision with 1b, node 2 can also avoid colliding with node 1b and 1c, and get it packets through. This is an example of the unfairness resulting from schemes that fail to accommodate the collective behavior.

The RTS/CTS contention scheme performs better than the two CSMA mechanisms. Nodes deep in the network are capable of delivering packets to the base station. However, the resulting bandwidth allocation is very unfair. "Gateway" nodes, which are close to the base station, dominate the channel and use up most of the channel capacity for delivering their own packets. Thru-route traffic obtains a small fraction of the channel around these "gateway" nodes and therefore suffers high loss rates.

The ARC scheme provides the most fair delivered bandwidth. "Gateway" nodes successfully decrease their rate of originating data and open up the channel for thru-route traffic to make it to the base station. For this particular result, we set $\alpha = 0.08$ and $\beta = 0.5$. Figure 15 shows that nodes below first level of the tree achieve about 0.2 packet/s of delivered bandwidth, or about 25% of the ideal rate that saturates the bottleneck.

Figure 16 shows how variance, a measure of fairness, changes for different α and β. For all values of α and β tested, our adaptive scheme always achieves a substantially lower variance than the other schemes. Furthermore, as α increases, β is irrelevant since the linear increase is so high that it overcomes the penalty exercised by β. The interesting regime is $\alpha < 0.2$, where β is more important.

Figure 16, 17, and 18 show that β plays an important role in controlling the tradeoff among fairness, energy efficiency, and aggregate bandwidth. A small β will lead to a conservative scheme that is energy efficient, but with low aggregated bandwidth (Figure 17) and high variance (Figure 16). A large β will impose a smaller penalty, especially for route-thru traffic, such that the aggregate bandwidth is higher and is more fair. However, higher contention of the channel at every hop in the network will degrade energy efficiency because a large fraction of the packets are dropped en route. Although both the RTS/CTS and 802.11 CSMA/ACK scheme achieve higher aggregate bandwidth than the adaptive scheme, they do so by heavily favoring nodes near the base station.

Finally, Figure 19 shows the *yield*, which measures the ratio of total packets received by the base station to total packets sent in the entire network. The CSMA schemes, as expected, perform the worst. The RTS/CTS scheme performs best. The reason is that RTS/CTS handshake avoids the hidden node problem in many cases. For our adaptive rate control scheme, α is the only parameter that affects yield. A large α leads to an aggressive scheme which increases the transmission rate rapidly and results in more contention. In a sense, packets are used as a contention unit like RTS, so a higher α naturally leads to a lower yield. This is also the reason that as α or β increases, the scheme becomes more aggressive and less energy efficient. Table 5 summarizes the effect in changing β based on our simulation with our hypothetical topology.

Our ARC scheme can be further improved by inferring po-

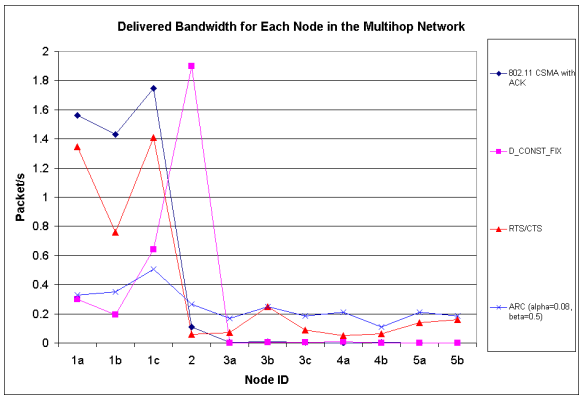

Figure 15: Delivered bandwidth from each node in the network to the base station for the different transmission control schemes. All nodes attempt to originate traffic at 4 packet/s.

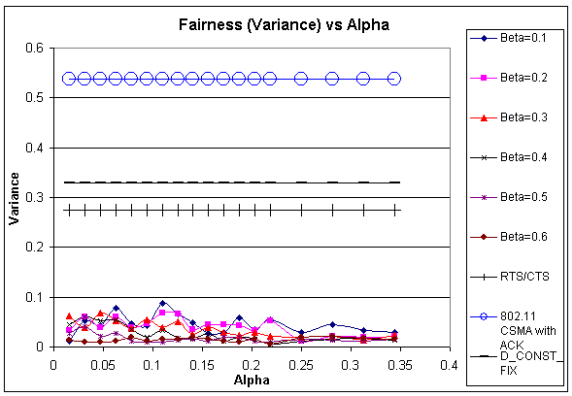

Figure 16: Variance, a measure of fairness, of the bandwidth received by the base station from each node.

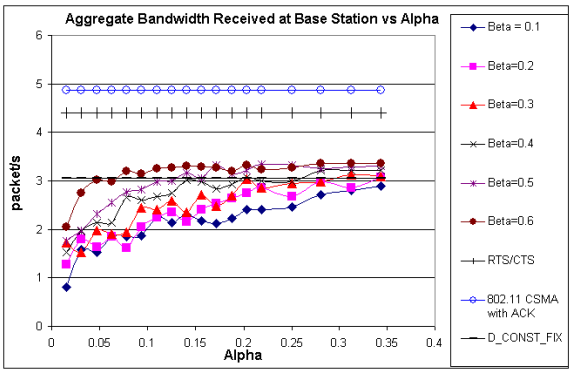

Figure 17: Aggregate bandwidth received by base station for the different schemes.

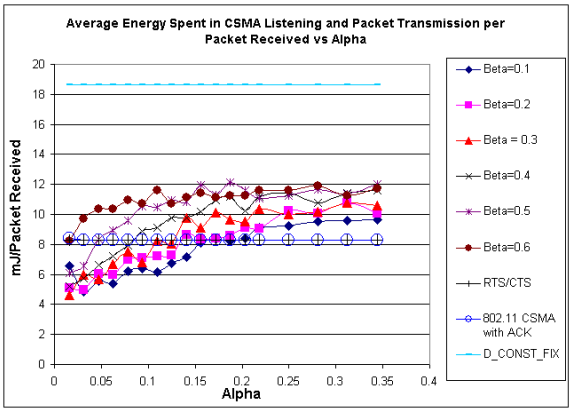

Figure 18: Average energy cost per packet received by the base station. The energy cost constitutes of energy spent in CSMA listening and packet routing and transmission.

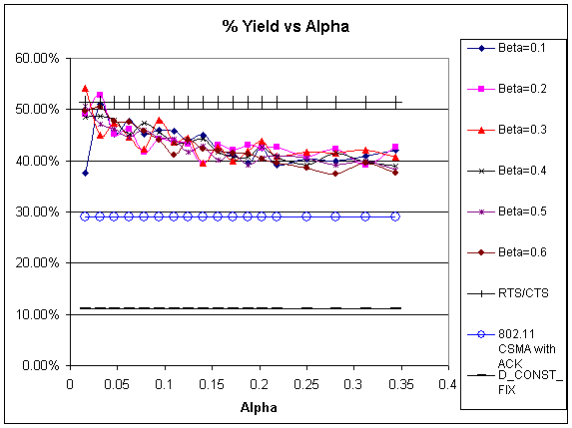

Figure 19: Percent Yield of the total packet sent as received by the base station for the different schemes.

tential hidden node problems by listening for parent transmissions, as explained in Section 4. Table 6 shows the comparison between the case with and without the inference mechanism on our adaptive scheme. The result clearly shows that such an inference scheme is effective in improving the fairness of the adaptive transmission scheme with no significant changes in the other metrics.

6.3 Empirical Measurements

As discussed in Section 1, an ad hoc multihop network topology is dynamically determined by how node placement and physical environment influence radio propagation, and by the choice of route discovery algorithms. Cell boundaries are not sharp, so interference effects may overreach useful communication cells and complicate the problem. Increasing the distance between node avoids cell interference, but causes an increase in loss rate. All in all, a defined topology with sharp cell boundaries is difficult to achieve. In our empirical study, we try to follow the topology shown in Fig-

Metric	Increase β	Decrease β
Fairness	Increase	Decrease
J/Packet Delivered	Increase	Decrease
Aggregate Bandwidth	Increase	Decrease
Yield	-	-

Table 5: Summary of the effects in changing β in the simulation.

Node ID	Bandwidth Delivered w/o Inference Scheme (packet/s)	Bandwidth Delivered with Inference Scheme (packet/s)
1a	0.33	0.31
1b	0.35	0.29
1c	0.51	0.38
2	0.27	0.30
3a	0.17	0.20
3b	0.25	0.21
3c	0.19	0.22
4a	0.21	0.18
4b	0.11	0.16
5a	0.21	0.20
5b	0.19	0.34
Aggregate Bandwidth	2.79	2.79
Variance	0.012	0.05
Yield	46.05%	44.03%
Delivered Bandwidth/J	0.32	0.29

Table 6: Measuring the effectiveness of the hidden node inference scheme over ARC with $\alpha = 0.08$ and $\beta = 0.5$.

ure 14. Cell overlapping and interference is inevitable, but this is expected in any real implementation in this regime. Finally, the impact of loss rate due to interference and decrease in signal strength over distance arises where they were absent from simulator. While a 5% loss rate in a single hop scenario is not significant, a 5% loss rate per link over a multihop network will significantly hinder data propagation. It is true that retransmission will eliminate some of the loss rate. For the sake of empirical study, we did not use this option. Since loss rate naturally acts as a damping factor of traffic, β_{route} may no longer be necessary. In the following studies, we set $\beta_{route} = 1$ so it asserts no penalty for thru-route traffic.

In the real implementation, we set $\alpha = 0.05$ and vary β from 0.125 to 0.25 and 0.5. ARC runs with the D_CONST_FIX CSMA scheme. We place the nodes on an open field and fix the routes to follow the topology in Figure 14. Measurements are done to compare the effect of different ARC settings and how they perform with respect to D_CONST_FIX alone. Following the same scenario as the simulation, each node sends at a rate of 4 packet/s. Furthermore, base station sends an explicit ACK for each packet it receives, rather than the entire packet.

Figure 20 shows the delivered bandwidth from each node. The results are quite close to the simulations. In the D_CONST_FIX case, "gateway" nodes clearly dominate the

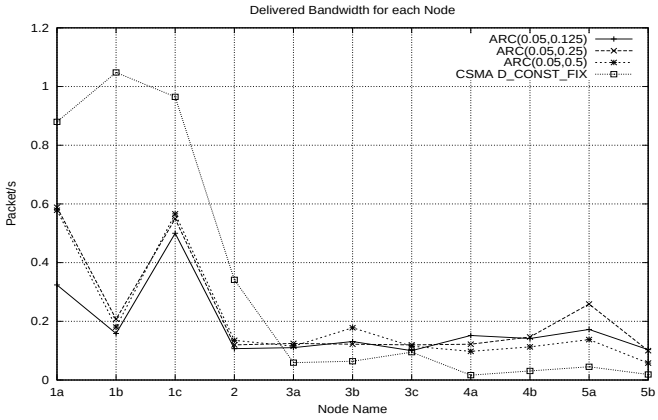

Figure 20: Delivered bandwidth received by the base station from each node. (Empirical)

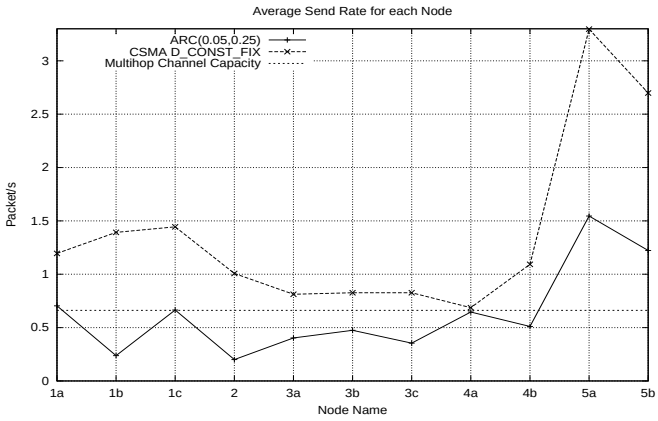

Figure 21: Average transmission rate of each node. (Empirical)

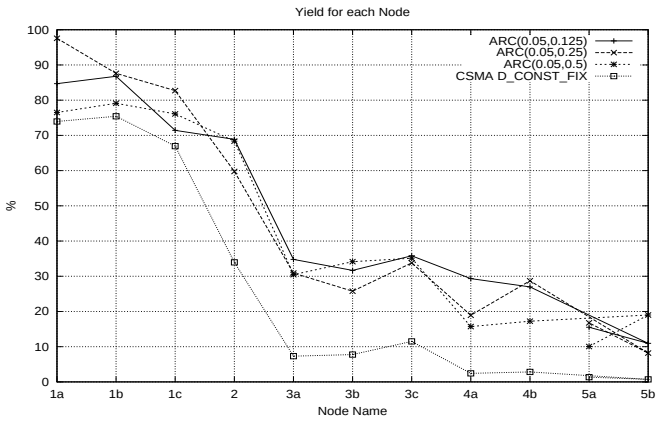

Figure 22: Percent yield (packet received by base station/packet transmitted) of each node. (Empirical)

channel around the base station and hinder traffic deep down in the network. However, node 2 does not hog the channel, as in simulation, because nodes at first level are no longer hidden. With CSMA alone, nodes deep in the network still achieve some level of delivered bandwidth rather than none as predicted by the simulator. It is expected that cell overlaps actually make CSMA perform better in practice.

In term of fairness, all three settings of ARC is more fair than D_CONST_FIX alone. It is clear that both traffic along a line, such as (1a,2,3a,4a,5a), and at a symmetric merge, such as 3a and 3b, achieve roughly equal delivered bandwidth of about 0.1 packet/s, or 15% to 20% of the multihop channel capacity. The "gateway" node, 1b, has successfully lowered its rate for thru-route traffic as its delivered bandwidth is lower than its neighboring nodes (1a and 1c). A lower β of 0.125 achieves a slightly fairer schedule than higher β, which is different from what the simulator predicts.

Figure 21 shows the actual transmission rate of each node. The data shows that ARC is much more effective in adapting the traffic to the multihop channel capacity without any global knowledge of the network. Almost all the nodes send below this capacity, rather than over-committing the channel. CSMA scheme alone performs much worse. Almost all nodes over-commit the channel, especially nodes deep in the network where the number of neighbors they can hear is small. However, in the cell of node 2, where most of the nodes are capable of hearing each other, CSMA scheme alone is quite effective as it lowers the send rate of 3a,3b,and 3c to a level close to the multihop channel capacity.

Figure 23 shows the actual dynamics of the rate control mechanism in lowering the transmission rate of nodes 1a, 1b, and 3c. Node 1a's send probability has greater oscillation than Node 1b. This is expected because node 1b's α is much smaller than 1a's α since 1b has to deliver traffic from its nine children. Node 1a has only its own traffic to handle. Node 3c's amplitude of oscillation lies between that of 1a and 1b. This is the correct behavior, since it has to deliver traffic from only two children.

Mechanism	Aggregate Bandwidth (packet/s)	Energy Efficiency (mJ/packet)
ARC(0.05, 0.125)	1.99	15.76
ARC(0.05, 0.25)	2.45	15.29
ARC(0.05, 0.5)	2.27	15.20
D_CONST_FIX	3.56	16.76

Table 7: Summary of aggregate bandwidth and energy efficiency measurement. The energy efficiency metric is the average amount of energy spent in CSMA listening and packet transmission per packet received by the base station.

Table 7 shows the aggregate bandwidth delivered by each scheme. The CSMA scheme achieves the highest aggregate bandwidth because it give preference to nodes near the base station. The simulator predicts that increase of β will increase aggregate bandwidth. Table 7 shows the opposite. As β increases, aggregate bandwidth actually decreases. In fact, Figure 22 shows the yield of each node, which suggests that a lower beta achieves higher yield for each packet sent,

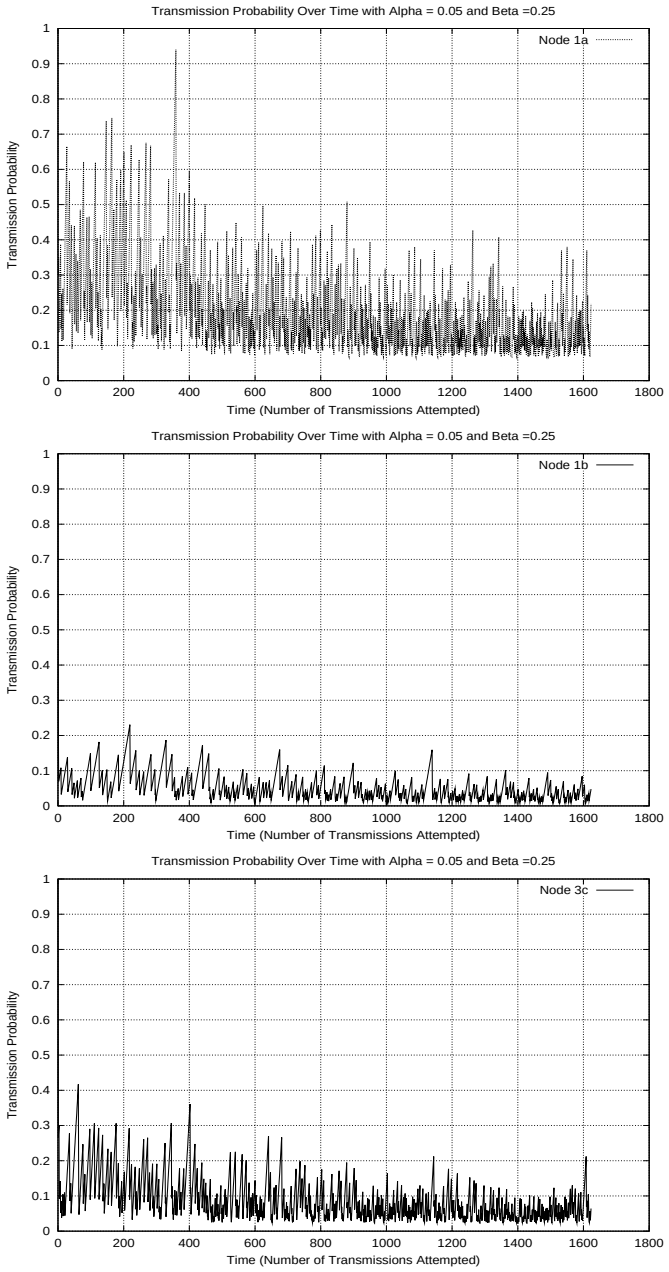

Figure 23: Dynamics of transmission probabilities of node 1a, 1b, and 3c over time. (Empirical)

and therefore, result in higher aggregate bandwidth. Finally, Table 7 shows the listening and transmission energy cost per packet received by the base station. The CSMA scheme bares the highest energy cost per packet delivered. In ARC, differences in β do not have a significant impact. The reality of cell overlaps and interference in the multihop scenario is probably the main cause of the discrepancy between simulation and empirical measurements.

7. CONCLUSION

The paper has shown how the application scenario, resource limitation, and network traffic characteristics in sensor networks differ from conventional computer networks and explained why existing MAC protocols are not suitable in this regime. We have defined multihop fairness and energy efficiency in bandwidth delivery to base station to be appropriate metrics in evaluating MAC protocols for sensor networks, and we based our design on these goals. A comprehensive study has been performed in understanding the appropriate carrier sensing mechanism. The conclusion is that random delay should be introduced prior to any transmission, with backoff acting as a phase shift for the periodicity of the application. A new, simple adaptive rate control scheme for achieving the desired metrics in a multihop network has been proposed and compared with conventional contention based schemes. The adaptive rate control scheme together with the new CSMA mechanism provides an effective media access control without explicit control packets. Simulation have shown that our proposed mechanism is effective in achieving fairness while maintaining good aggregate bandwidth with reasonable energy efficiency. Our adaptive scheme is extremely efficient in energy for low traffic situation which is the common case in sensor networks. These simulation results are further supported by real implementation on our tiny network sensor platform.

8. ACKNOWLEDGMENTS

This work is supported in part by the Defense Advanced Project Agency (grant DABT 63-98-C-0038) "Ninja" and (grant N66001-99-2-8913) "Endeavour", and by the National Science Foundation (grant RI EIA-9802069). Support is provided as well by the California MICRO program, Intel Corporation, IBM, Sun Microsystems and Philips. We would like to thank our TinyOS team, Jason Hill and Robert Szewczyk, for all the discussions related to this work, UCLA sensor networks team for their valuable feedback, and Wilson So for his assistance in collecting some of the measurements.

9. REFERENCES

[1] Bluetooth. http://www.bluetooth.com.

[2] ANSI/IEEE Std 802.11 *1999 Edition*.

[3] Atmel, Inc. At90s4434/ls4434/s8535/ls8535 *Preliminary (Complete) Datasheet*.

[4] V. Bharghavan, A. Demers, S.Shenker, and L. Zhang. MACAW: A media access protocol for wireless lans. In *Proceedings of SIGCOMM Conference*, pages 212–225, 1994.

[5] S. Floyd and K. Fall. Promoting the use of end-to-end congestion control in the internet. In *IEEE/ACM Transactions on Networking*, 1998.

[6] V. Jacobson. Congestion avoidance and control. In *Proceedings of ACM SIGCOMM Conference*, 1998.

[7] Jason Hill, Robert Szewczyk, Alec Woo, Seth Hollar, David Culler, Kristofer Pister. System architecture directions for networked sensors. In *ASPLOS*, 2000.

[8] L. Kleinrock and F. Tobagi. Packet switching in radio channels, part 1: Carrier sense multiple-access models and their trhoughput-delay characteristics. 23(5):1400–1416, 1975.

[9] S. Lam. A carrier sense multiple access protocol for local networks. In *Computer Networks*, volume 4, pages 21–32, 1980.

[10] K. S. J. Pister, J. M. Kahn, and B. E. Boser. Smart dust: Wireless networks of millimeter-scale sensor nodes. In *1999 UCB Electronics Research Laboratory Research Summary*, 1999.

[11] T. S. Rappaport. *Wireless Communications, Principles and Practice*. Prentice Hall PTR, New Jersey, 1996.

[12] T. J. Shepard. A channel access scheme for large dense packet radio networks. In *Proceedings of ACM SIGCOMM*, 1996.

[13] S. Singh and C. Raghavendra. Pamas - power aware multi-access protocol with signalling for ad hoc networks. In *ACM Computer Communication Review*, 1998.

[14] S. Singh, M. Woo, and C. Raghavendra. Power-aware routing in mobile ad hoc networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 181–190, 1998.

[15] A. Tanenbaum. *Computer Networks*. Prentice Hall Inc., 1981.

[16] Thyagarajan Nandagopal, Tae-Eun Kim, Xia Gao, Vaduvur Bharghavan. Achieving mac layer fairness in wireless packet networks. In *MOBICOM*, 2000.

[17] F. Tobagi and L. Kleinrock. Packet switching in radio channels, part ii: Hidden-terminal problem in carrier sense multiple access and the busy-tone solution. 23(5):1417–1433, 1975.