Structuring Contention-based Channel Access in Wireless Sensor Networks*

Shane B. Eisenman
Electrical Engineering, Columbia University
New York, New York, USA
shane@ee.columbia.edu

Andrew T. Campbell
Computer Science, Dartmouth College
Hanover, New Hampshire, USA
campbell@cs.dartmouth.edu

ABSTRACT

In a wireless network using CSMA for MAC, packet collisions can result either because of the vulnerability stemming from the idle channel detection delay of the radio device or because of the hidden terminal problem. Both of these collision scenarios can be addressed by applying some additional structure to carrier sense-based channel access, sacrificing throughput to improve collision protection. Specifically, we assign a unique time slice to each contending transmitter that is designed to allow collision-free access to each data slot while still using carrier sense as the mechanism to determine channel state. The width of each time slice is equal to the idle channel detection delay of the radio. We prepend a fixed-length time interval to each transmitted packet whose length is proportional to the idle channel detection delay of the radio device and the node density. We discuss this method of channel access including tradeoffs and requirements, and analyze the performance in comparison with existing MAC strategies.

Categories and Subject Descriptors: C.2.1 [Computer-Communications Networks]: Network Protocols, Wireless Communications

General Terms: Algorithms, Design.

Keywords: Medium Access.

1. INTRODUCTION

In a wireless communication system, two or more transmitters may wish to access a shared channel concurrently. If these multiple transmissions are coincident on a radio receiver, interference occurs and the transmissions are corrupted at that receiver. In packet radio networks, this situation is referred to as a "packet collision". MAC (Medium or Multiple Access Control) protocols are typically employed in these systems to arbitrate wireless channel access and reduce the occurence of such corruption.

Wireless sensor networks are packet radio systems that often use low-cost, low-power, single channel transceivers [8] for communication. To date, the overwhelming majority of wireless sensor network system deployments, hereafter *sensor networks*, have employed variants of the Carrier Sense Multiple Access (CSMA) protocol [12] to reduce collision. This is primarily due to its low implementation complexity, which is well matched both to the archetypically memory-limited and computation-limited hardware [3] comprising such networks, and to the relatively low collection/transmission duty cycle required by a popular class of applications (e.g. [13] [14]).

With CSMA, a node that wishes to transmit listens to the wireless channel prior to sending its own transmission. If it hears an ongoing transmission it defers, or "backs off", until a later time; otherwise it transmits. In a time slotted channel, using carrier sensing provides for reasonably good collision avoidance. Assuming an independent Poisson packet generation rate at each transmitter, the throughput in a time slotted channel can reach about 54% of the offered load when employing carrier sense, as compared to only 37% without it [12]. Though packet generation in wireless sensor networks is not spatially uniform or independent, or temporally Poisson, due to inherent spatial and temporal correlations in the sensed events, using carrier sense is still considered good practice [15]. However, CSMA networks are still susceptible to collision due the limitations of carrier sensing. Among others [15], carrier sensing cannot help detect transmissions that are out of range of the transmitter but may still interfere at an intended receiver, and also fails when a transmission arrives at the transmitter after it has sensed the channel, determined it to be clear, and switched from receive mode to transmit mode. We will refer to the first as the *hidden transmitter problem* and the latter as the *idle channel detection delay problem.* Figure 1 shows scenarios where collisions result due to both of these problems. Using an RTS/CTS exchange [11] in combination with CSMA can effectively mitigate the hidden transmitter problem for the transmission of data packets, but can incur high overhead when used in wireless sensor networks where data packets are typically short [16] [2]. When offered load is high, this overhead can limit the goodput (data throughput) of the network. Further transmission of RTS/CTS packets themselves are subject to the *hidden transmitter problem.* The use of out-of-band busy tones has been proposed to pro-

*This work is supported by the Army Research Office (ARO) under Award W911NF-04-1-0311 on resilient sensor networks.

IPSN'06, April 19–21, 2006, Nashville, Tennessee, USA.

tect the RTS/CTS exchange (e.g., [24] and references), but requires the use of a second radio transceiver.

An alternative approach, called Time Division Multiple Access (TDMA), that solves both the hidden transmitter problem and the idle channel detection delay problem is to assign unique data packet time slots to each transmitter in a channel contention region (the area within which simultaneous transmission from two transmitters can interfere at a common receiver). However, adoption of TDMA protocols in wireless sensor networks has been slow for many reasons, as pointed out by the authors of S-MAC [4]. A main concern is the fact that when the traffic generation rate is not spatially uniform, data slots for some transmitters can go unused while other transmitters are rate limited due to their slot allocation. One way to address this problem is to assign more slots to transmitters that need to send more packets, but this can be difficult in networks where the traffic dynamics are unpredictable, e.g., in wireless sensor networks.

Recently, a class of hybrid approaches have been proposed that attempt to combine the best attributes of both CSMA and TDMA while making design choices that are appropriate for the wireless sensor network regime. Among these are S-MAC, T-MAC [5], TRAMA [6], Sift [18] and Z-MAC [17]. Our proposal, dubbed qMAC, joins this class of hybrid protocols, adding additional structure to the baseline contention-based channel access method in use in today's wireless sensor networks. Without the overhead of RTS/CTS exchange or the potential for wasted data transmission slots, we solve the hidden transmitter problem and the idle channel detection delay problem described previously. Assuming a slotted channel, we prepend a fixed-length time interval to each transmitted packet. The length of this time interval is proportional to the idle channel detection delay of the radio device and the node density. More specifically, this fixed-length time interval is composed of D time slices of length τ, where D is proportional to the number of nodes in the contention region. Figure 2 shows the resulting time structure of the transmit channel. Each one hop neighbor of a node i is assigned a unique τ slice (with additional constraints) in which to start transmissions to node i for each data packet slot. Further, proper assignment of τ slices to nodes in the two hop neighborhood of i eliminates collisions resulting from hidden transmitters, when coupled with additional signaling. By prepending a number of τ-length slots to each transmitted packet, we sacrifice bandwidth, but can eliminate packet collision (wasted energy) under all network conditions by doing so. We believe this is the right tradeoff to make in designing a MAC protocol for wireless sensor networks. In the following sections we discuss this method of channel access arbitration including tradeoffs and requirements, and analyze the performance in comparison with existing MAC strategies. Discussion of related work is integrated where appropriate.

2. PROTOCOL DESCRIPTION

qMAC comprises three components: Neighbor Discover and Access Quantum Assignment that are executed mandatorily at startup and as needed thereafter, and Medium Access that is used during every data packet transmission. These three components are described below, after introducing some notation used throughout the rest of the paper and stating our design assumptions.

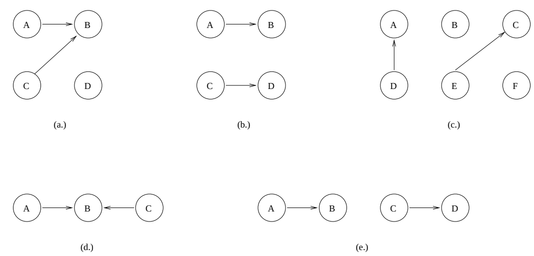

Figure 1: CSMA collision scenarios. Node separation is d and transmission radius is $d\sqrt{2}$. Arrows indicate the intended receiver. (a.) through (c.) represent collisions resulting from the idle channel detection delay, while (d.) and (e.) are hidden transmitter collision scenarios.

2.1 Preliminaries

2.1.1 Notation

Symbol	Description
τ	Idle channel detection delay. Equal to the sum of the RSSI sampling time and the receive-to-transmit switch time of the radio circuitry.
$\mathcal{N}_i^1$	Set of nodes in i's one hop neighborhood.
$\mathcal{N}_i^2$	Set of nodes in i's two hop neighborhood.
q_i	Channel access quantum assigned to node i.
t_{skew}	Worst-case clock skew allowed by the synchronization protocol.

2.1.2 Assumptions

We assume a slotted channel and as such require a protocol to provide synchronization between nodes in a contention region. qMAC performs best with perfect synchronization, but our analysis accounts for possible deviation from synchronicity by means of the parameter t_{skew}. qMAC's performance degrades gracefully in the face of such deviation, as long as it is known in advance. That is, the throughput decreases and average delay increases, but no packet collisions occur due to the hidden transmitter or idle channel detection delay problem. There are several synchronization proposals in the wireless sensor network literature that are sufficient for our purposes, e.g., [20] [19].

qMAC is based on the observation that smallest time scale at which a node can interact with the network is limited by τ. Hence, we term τ the *time quantum* of the radio communication system. In particular, a node that is sensing the channel to detect ongoing transmissions takes τ time to capture and analyze a sample (an RSSI value) and switch to transmit mode to begin sending a packet. The radio state can thus be modeled as a three state machine (receive state, q-state, and transmit state), where transitions between receive and transmit states must pass through q-state. Implicit in this view is the assumption that the receive-to-transmit delay is equal to the transmit-to-receive delay. We further assume that perfect hearing (channel fading is not considered except where otherwise noted), though this is simply for ease of exposition, and also propagation delay is assumed to be negligible, though it could easily be incorporated in the value of τ.

Finally, we assume the availability of an efficient distributed

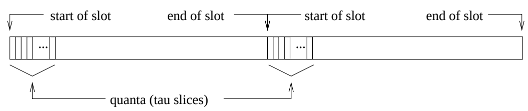

Figure 2: Slotting of the transmit channel. If there are k quanta (τ slices), then the slot length is $k\tau + pktTime$.

coloring algorithm that can be implemented on the relatively resource-poor hardware of today's wireless sensor networks. DRAND [21] may be used to fulfill this requirement.

2.2 Neighbor Discovery

In qMAC, each node requires knowledge of one hop, and optionally its two hop neighbors. As described later in Section 2.4, knowledge of a node's one hop neighbors is required to assign quanta such that the idle channel detection delay problem is solved, while knowledge of two hop neighbors is required to solve the hidden transmitter problem.

A simple lightweight beaconing scheme can be used by nodes to learn their two hop neighborhood. As nodes boot up, each broadcasts a packet containing its MAC address, along with any one hop neighbors it has discovered. The neighborhood information gleaned from the beacon exchange is incorporated into the channel access quanta assignment (see Section 2.3) so it is important that the start-up beaconing period be long enough to allow for most nodes to get an accurate neighborhood map. So called "late joiners" to the network by definition cannot be learned during the startup beaconing phase, but can still participate in the network as described in Section 2.4.

In practice, many wireless sensor networks will utilize a routing protocol (e.g., [7] [1]) to facilitate data delivery, rather than flooding every packet. If so, it is likely that qMAC's neighbor discovery requirement can be met by snooping route control packets, or reading from a cross-layer state cache (e.g., [22]), at least for one hop neighbors.

2.3 Access Quantum Assignment

For the purposes of qMAC's channel access scheme, described in Section 2.4, we assign a unique τ-length time slice, or access quantum, to each node in a contention region, taking care to avoid assignments that lead to packet collision. Thinking of the access quantum assignment problem as a graph edge coloring problem, a successful coloring will have the property that no edges that enter or leave a node have the same color (the "transmit edge" from a node is considered as only one color regardless of how many receivers are connected). However, because of the broadcast nature of the wireless channel and the implications of the idle channel detection delay, there are additional constraints on the coloring.

Let q_{ij} denote the quantum integer assigned for node i to transmit to node j in the coloring problem. Since transmissions on the wireless channel are intrinsically broadcast transmissions, $\forall j \in \mathcal{N}_i^1$, q_{ij} must be the same. We therefore simplify notation by using q_i to denote the quantum integer assigned for node i to transmit. Aside from the requirement that access quantum assignments must be unique, further restrictions are imposed in order for qMAC's medium access flow to guarantee collision-free channel access. Namely, to solve the idle channel detection delay problem, $\forall j \in \mathcal{N}_i^1$, $|q_i - q_j| > 1$ must hold. That is, one hop neighbors can not have adjacent quantum assignments. Further, to solve the hidden transmitter problem, $\forall k \in \mathcal{N}_i^2 - \mathcal{N}_i^1$ (nodes exactly two hops from i) $|q_i - q_k| > 3$ must hold. Section 2.4 discusses why these constraints must be met; in this section we comment on the method of access quantum assignment itself.

There are a number of distributed graph coloring algorithms proposed in the literature that can be applied to provide contention-free MAC scheduling. These can be implemented with various degrees of efficiency in terms of the computation, communication and memory required. Among these DRAND seems well matched to the likely operating environment of qMAC, wireless sensor networks. DRAND uses knowledge of the two hop neighborhood to provide a broadcast schedule where no two nodes within a two hop communication neighborhood are assigned to the same slot. While in qMAC we do not assign packet *slots* to nodes, the assignment of access order implied by the *quantum* integer is a similar problem that can be solved in the same way. The chromatic performance and message complexity of the method scale as $\mathcal{O}(\delta)$, where δ is the density of the contention region [21]. However, DRAND is neither designed nor demonstrated with the additional node coloring constraints qMAC requires.

While it is likely that one of these techniques (e.g., DRAND) can be adapted to provide the coloring necessary for qMAC, for the present we describe an alternative centralized approach used for proof of concept. The method is to start with a node (in practice, the information sink of the sensor network would be a good place to initiate the coloring), and pass a token across the graph in a breadth-first fashion along the wireless communication links. The token enables the holder to select and announce its chosen quantum q. Each node k in the network maintains a mapping of node ID to quanta q_i, $\forall i \in \mathcal{N}_k^2$. This mapping is built from information contained in received quantum assignment broadcast packets. Upon receiving the token, a node k scans its list of mappings and selects the lowest quantum integer not present, subject to the previously defined constraints w.r.t. quantum separation. This integer is set as q_k and broadcasted to all one hop neighbors of k, along with all the mappings $q_i, \forall i \in \mathcal{N}_k^1$ known by k. For correctness, it is required that ties in the breadth-first token propagation are broken in favor of the node with the highest *connectivity*. That is, for nodes at the same "level" in the breadth-first search, the connectivity tiebreaker gives the token to the node that has the most neighbors that have already possessed the token. The intuition for this is that the node with the most connections to previous token holders also knows the most about existing access quanta assignment. For determinism, a secondary tie-breaker can be added that gives the token to the node with the lower ID. We numerically simulate the performance of this centralized breadth-first token propagation method in Section 3.1.

2.4 Medium Access

Recall that the channel is divided into slots as shown in Figure 2. From a transmitter's perspective channel access comprises two phases, reservation and transmission. When a node i wishes to transmit to node j in a given slot, it defers until q_i τ slices have elapsed since the start of the slot and then transmits a τ-length *reservation burst*. Node i waits 3τ

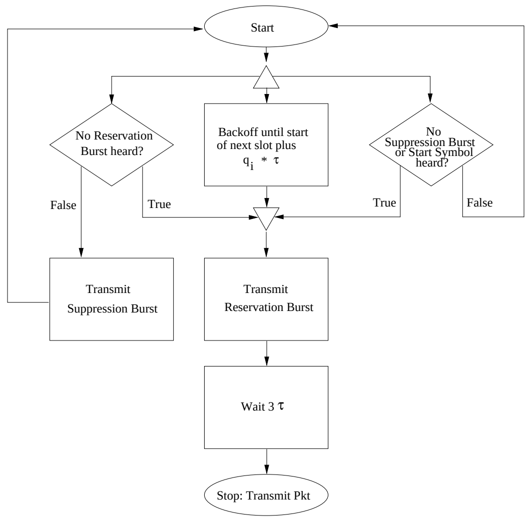

Figure 3: Channel access flow diagram for transmitters. Not shown: a receiver that hears a *reservation burst* transmits a *suppress burst*, waits τ and then begins receiving the data packet. A receiver that hears a suppress burst does nothing.

and then commences transmission of its data packet. Upon receiving a reservation burst from i, nodes $j \in \mathcal{N}_i^1$ transmit a τ-length *suppress burst* in slice $q_i + 2$. The suppress burst acts as a busy tone, indicating to transmitters hidden from i that they should not transmit in the current slot. Figure 3 gives a graphical representation of the channel access procedure for a node wishing to transmit a packet. If a perfect coloring (i.e., assignment of access quanta) has been done, as described in Section 2.3, the packet of node i is guaranteed not to collide with any other at node $j \in \mathcal{N}_i^1$. *Specifically, the minimum two quantum distance between the assigned quanta of one hop neighbors guarantees that all transceivers have enough time to sample and react to channel activity, thereby eliminating collisions due to the idle channel detection delay of the radio device. Further, the minimum four quantum distance between the assigned quanta of nodes exactly two hops away (hidden transmitters) allows for one hop neighbors of a transmitter i to warn the transmitters hidden from i that a transmission is already reserved for the current slot in time to avoid hidden transmitter collisions.*

Figure 4 shows an simple illustrative example of the access timing, including the need for the aforementioned constraints on access quantum assignment. With slight abuse of notation, we will use τ here also to mean just the radio mode switch time. Nodes A, B and C are placed such that A can communicate with B, and B with C, but A and C are out of range of each other. In the figure, all nodes wish to transmit, and $q_A = 0$, $q_B = 2$ and $q_C = 4$. A sends a reservation burst in its assigned quantum to indicate its intention to transmit, and then goes back to receive mode to listen for the suppress burst. Although B also wishes to transmit and has the ability to reserve the channel in the third quantum, it listens before that time to see if any other node with a lower quantum assignment is reserving the channel. Thus it receives the reservation burst sent by A and sends a suppress burst as soon as it can. Since it takes τ to switch from receive mode to transmit mode, B sends the suppress burst in the third quantum interval. B then takes τ to switch back into receive mode to receive the data packet transmission. A and C hear the suppress burst in the third quantum and therefore C refrains from sending its reservation burst in the fifth quantum, instead deferring to the start of the next slot.

It may happen that reservation bursts from different nodes in a contention region are sent in the same quantum. This can happen if the access quantum assignment is faulty, if a node is moved post-assignment to a different contention region, if a "late joiner" randomly selects a quantum that collides with the valid assignment of another transmitter, or due to time-varying wireless transmission range. Whatever the cause, in essence we can think of the situation as a number of transmitters M in a contention region that is split into D groups, $M > D$, each group contending for one valid quantum. Luckily, this situation can often be detected if a suppress burst is not heard in response to a transmitted reservation burst. In such a case, nodes involved in the quantum collision (i.e., *q-colliders*) use a p-persistent approach and, despite not hearing a suppress burst in response to their reservation burst, transmit their data packet with probability p in the current slot. Clearly we must have $0 < p < 1$ to avoid deadlock. Ideally p should be adaptive to $\frac{1}{E[\text{num of } q\!-\!colliders]}$ to be fair, but since this is not known in general p should be set according to the knowledge of the density of nodes in the contention region (which is known from the neighbor discovery phase, and maintained by periodic beaconing) since those are the potential quantum colliders. Experience with p-persistent CSMA shows that, under Poisson packet generation assumptions, a smaller p provides a smaller collision probability and higher throughput for a given offered load in the contention region [12]. It should be noted that this way of mitigating quantum collisions is only effective when M is modestly more than D. When $M \gg D$ then qMAC devolves to medium access that behaves like p-persistent CSMA, where instead of carrier sense to determine an idle channel the idle channel is implicit in the τ-slicing of the slot. In this case, a reassignment of access quanta in the contention region should be done to give each transmitter its own quantum.

Once access quanta assignments have been made, nodes with lower quanta always have channel access priority over nodes with higher quanta. Despite the spatial correlation of sensed data in sensor networks, especially within a contention area, over time and depending on traffic patterns this can result in a skewed picture of the sensor field as constructed by data packets received at the information sink of the sensor network. Further, such prolonged disparity in the access priority between transmitters can contribute to unequal energy consumption. To address this problem in qMAC, we require a distributed, computationally lightweight function that performs an injective mapping of a valid two hop neighborhood assignment of quanta to a new assignment, while maintaining the properties of a valid assignment. Such a function could be run periodically (and synchronously) on each node, perhaps every slot if desired. Quantum reassignment would have the additional benefit of evenly distributing the effect of quantum colliders, discussed in the previous paragraph, over all the nodes in the contention region. Design of such a function is an open problem. We note that inclusion of a hash function (which is injective) in the overall remapping function may have po-

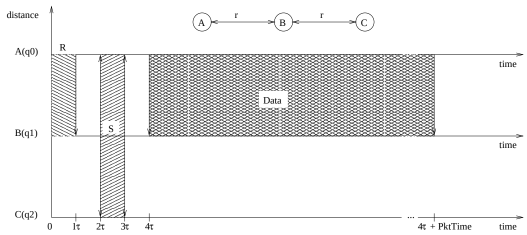

Figure 4: Channel access example for three nodes, labeled A, B and C, that lie on a line separated by the transmission radius r. $q_A = 0$, $q_B = 2$ and $q_C = 4$. The sequence of events for A to transmit to B is shown. R designates a reservation burst and S designates a suppress burst.

tential.

3. EVALUATION

With qMAC we resolve collisions due to both the hidden transmitter problem and the idle channel detection delay problem, with the introduction of some bandwidth overhead. In this section we analyze qMAC in terms of this overhead and also provide a comparison between qMAC and three related MAC schemes, CSMA, TDMA and Z-MAC. The overhead analysis focuses on the reduction in maximum theoretical throughput caused by prepending a number of channel access quanta to each data packet transmission, and by the use of reservation and suppression bursts introduced in Section 2.4. The effect of imperfect synchronization is also discussed. The protocols used to provide synchronization and to provide channel access quanta assignment also add measurable overhead to the qMAC solution. However, as these protocols run orthogonal to the core qMAC channel access proposal, an explicit overhead/complexity analysis is not included here.

3.1 Access Quantum Overhead

The penalty of solving the idle channel detection delay problem and the hidden transmitter problem with qMAC is the sacrifice of bandwidth from adding access quanta to each data slot. The number of slots required is intuitively proportional to the number of nodes in a contention region, but is also strongly impacted by the coloring method used to assign quanta while respecting the necessary constraints on quantum separation.

It can be shown (omitted) that an ideal centralized scheme can achieve a coloring such that the cardinality of the palette scales linearly with the node density but is constant with respect to network size, when the only requirement is uniqueness within a two hop neighborhood. In fact, the distributed solution DRAND provides offers this level of chromatic efficiency as well [21]. However, the same has not been shown (to the best of our knowledge) when the additional constraints of chromatic spacing (discussed in terms of quantum assignment in Section 2.3) are imposed. While we think DRAND can be modified to satisfy our more stringent requirements on color assignment, here we analyze the performance of the breadth-first token passing approach to color (quantum) assignment described in Section 2.3 as a proof of concept.

Figure 5 shows the required number of access quanta as the network scales in size. Performance for two topologies, two dimensional grid and randomly scattered according to a two dimensional uniform distribution, are shown. Note that the quantum assignment algorithm is deterministic so that for a fixed topology the coloring will always be the same. As such, the curves based on the grid topology are from a single trial, whereas the curves based on the uniformly distributed random topology represent the average of five different randomly generated topologies. From the figure we see that the number of required quanta is roughly constant with increasing network size, that is, as the number of nodes in the network increases but the node density and the average number of nodes in the contention region stays constant. However, the less regular node placement of the uniformly distributed random topology requires somewhat more access quantum slots. Figure 5 also shows the relative costs, in terms of required number of quanta, of solving the hidden transmitter problem ("2-hop constraints" curves) and the idle channel detection delay problem ("1-hop constraints" curves). From this numerical simulation we see that roughly twice as many quanta are needed to satisfy the additional coloring constraint necessary for qMAC to solve the hidden transmitter problem. This result is inline with the geometry of the problem. Approximating the one hop coverage area of node i with transmission radius r as the square circumscribed by a disc of radius r, we calculate the number of one hop neighbors of i as $2r^2\delta$ for a uniformly distributed random topology, where δ is the node density, and $(\frac{2r}{\sqrt{2}d}+1)^2$ for a grid topology, where d is the grid spacing. Similarly, we calculate the number of nodes exactly two hops from i as $6r^2\delta$ for a uniformly distributed random topology and $(\frac{2\sqrt{3}r}{\sqrt{2}d})^2 + \frac{4r}{\sqrt{2}d}$ for a grid topology. Since there are approximately thrice as many nodes that must conform to the quantum spacing requirement for nodes exactly two hops away as must conform to the quantum spacing requirement for nodes exactly one hop away, the observed jump in required quanta from the "1-hop constraints" curves to the "2-hop constraints" curves is reasonable.

Figure 6 shows the required number of access quanta as the number of nodes in the contention region increases. Again, performance for grid and uniformly distributed random topologies, and the relative costs of solving the hidden transmitter problem ("2-hop constraints" curves) and the idle channel detection delay problem ("1-hop constraints" curves) are shown. We see that the required number of access quanta scales linearly with increasing number of nodes in the contention region (increasing transmit radius) for both grid and uniformly distributed random topologies. We also again see the jump in quanta required to solve the hidden transmitter problem.

Figures 5 and 6 demonstrate that our centralized breadth-first token passing method of quantum slot assignment (node coloring) achieves the same *order* of chromatic efficiency, while meeting the quantum spacing requirements of qMAC, as DRAND achieves without considering these constraints. Since it is known that breadth first search can be done in a distributed fashion [23], this result gives us confidence that an existing distributed protocol can be modified to take the quantum spacing requirements into consideration and give similar or better performance than what we show here.

To get an idea about how the required number of access quanta impacts the maximum throughput in a realistic sensor network, we consider values of τ for the CC1000 radio [8]

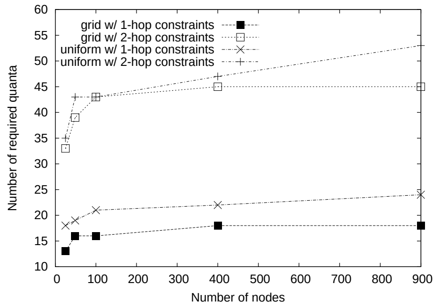

Figure 5: The required number of access quanta is roughly constant with increasing network size (fixed transmission radius, fixed density, increasing area). The transmission radius is $\sqrt{2} * d$, where d is the grid spacing.

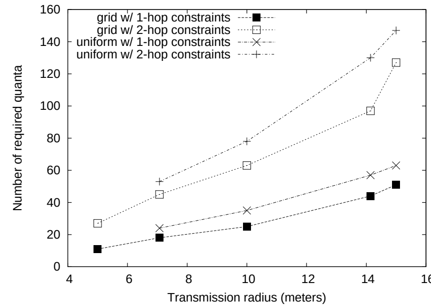

Figure 6: The required number of access quanta increases linearly with increasing number of nodes in a contention region. The number of nodes is fixed (900) as is the area (29dx29d, where d is the grid spacing).

and the TR1000 radio [9] when used with TinyOS [10]. On a Mica2 mote the worst-case ADC sample time is 686μs and the radio mode switch time is 200μs, giving an upper bound on τ of 886μs. The byte transmission time, with a transmission rate of C=19.2kbps is 417μs. Ignoring the difference between 0-based and 1-based numbering of the quanta, we denote $\max(q)$ to be the number of quanta required for qMAC. With a default TinyOS packet length of L=36 bytes (ignoring preamble and start symbol), we can quantify the overhead O and maximum link goodput G_{max} as follows.

$$O = \frac{\max(q)(2t_{skew} + \tau)}{\max(q)(2t_{skew} + \tau) + L/C},$$

$$G_{max} = \frac{L}{\max(q)(2t_{skew} + \tau) + L/C},$$

where t_{skew} is the worst-case bound on the pairwise clock skew in a contention region provided by the synchronization protocol. Note that t_{skew} acts a multiplier to spread the effective width of each quantum. For example, from Figure 5 for a 900 node grid network we require 45 quanta. Conservatively assuming a t_{skew} of 10μs within the two hop contention region (FTSP [19] reports a per-hop sync error of about 1μs) we have $O \approx 73\%$ and $G_{max} \approx 5.2$kbps when using the CC1000 radio on Mica2. The Mica mote [3] uses the TR1000 radio which has a τ of only 250μs and C = 10kbps, yielding $O \approx 26\%$ and $G_{max} \approx 7.4$kbps. Clearly, the overhead of qMAC is highly dependent on the radio's τ.

We observe that in qMAC there is a tradeoff between throughput and packet loss due to collision. qMAC sacrifices bandwidth to eliminate collisions due to the limitations of τ and the hidden transmitter problem even in the face of imperfect synchronization. We believe this approach is justified since in wireless sensor networks wasting energy from packet collisions is a more important concern than maximizing throughput.

3.2 Comparison with MAC Alternatives

qMAC uses the carrier sensing aspect of CSMA to determine the state of the wireless channel. Recognizing the limitations of this technique due to τ, and the inability of carrier sense to identify the presence of hidden transmitters, qMAC adds synchronization and the notion of unique transmission opportunity for each node in a contention region, similar in spirit to TDMA. The access quantum that qMAC assigns to each node in essence constitutes a prioritized backoff similar to "owner" and "non-owner" groups in Z-MAC [17]. In this section we compare qMAC in turn to each of these, its closest relatives.

3.2.1 Comparison with CSMA

In Section 1, we identified the idle channel detection delay problem and the hidden transmitter problem that plague simple CSMA. The addition of an RTS/CTS exchange [11] provides more robust channel access by reducing the occurrence of the hidden transmitter problem and this exchange inspires the reservation burst/suppress burst mechanism that is part of qMAC. In the comparison of these two approaches we consider the likelihood that each will be successful at eliminating data packets collisions, their relative bandwidth overhead in terms of wasted transmit opportunities, and implementation complexity.

First we observe that although a successful RTS/CTS exchange is effective at mitigating the hidden transmitter problem, the transmissions of RTS and CTS packets themselves are protected only by plain CSMA and thus subject to collision due to the fundamental τ problem and hidden transmitter problem. Because of the channel synchronization and access quantum assignment in qMAC, collisions involving the reservation burst or suppress burst should be less common and can be addressed as discussed in Section 2.4. Further, RTS and CTS packets contain at least the transmitter and receiver addresses and therefore are typically several bytes longer than the τ-length bursts used in qMAC. For example, even with the relatively long τ of the CC1000 radio, still the combined length of the reservation and suppress bursts is 1772μs while the combined length of two TinyOS packets whose payloads contain only the source address is 18 byte times (5 byte header, 2 byte payload, 2 byte CRC for each) or 7506μs. Not only does this increased length imply more energy spent, but also a higher probability of failure due to channel fading and collision. A more rigorous quantitative analysis of collision probability is omitted due to space constraints.

Additionally, there is the overhead of pre-RTS backoff to consider. In general, it is difficult to quantify this overhead since many CSMA backoff schemes have been proposed. As a reasonable benchmark for sensor networks we consider the backoff strategy employed by the default MAC proto-

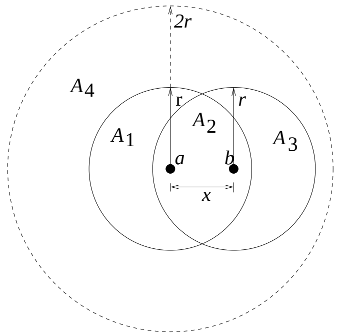

Figure 7: Areas of interest w.r.t. a packet transmission from a to b.

col (B-MAC [16]) currently released with TinyOS. B-MAC implements two backoffs, an "initial backoff" that is applied before each new packet transmission attempt and a "congestion backoff" that is applied if the channel is deemed busy by the carrier sense mechanism. Though both backoffs are uniformly chosen at random the average "initial backoff" is 16 byte times, and the average "congestion backoff is 8 byte times" for the Mica2 radio stack. In qMAC there is no analog to "initial backoff" - it is not needed because the channel is slotted. "Congestion backoff" is analogous to the time a node must wait to transmit in the next slot. However, in qMAC a such a node always backs off the perfect amount before attempting to send in the next slot, whereas CSMA backs off some multiple of "congestion backoff" with an average overwaiting of half the congestion backoff, or 4 byte times in the default B-MAC implementation. In summary, taking advantage of the slotted channel, qMAC implements an ideal backoff scheme, while the default implementation of B-MAC spends up to 48 byte transmission times on backoff for each transmitted packet, an average case sub-optimality of 24 byte times.

Further, the synchronized nature of channel access in qMAC allows a node to sleep after it has determined that it is not the intended recipient of a packet in a given slot. Thus, we can approximate the sleep duty cycle for a given node in this case by $(\frac{D\tau}{2} + \frac{L}{C})/(D\tau + \frac{L}{C})$, where D is the number of required τ slices and $\frac{L}{C}$ is the data packet transmission time. The actual percentage of time a given node can sleep is thus determined by the density and the length of the radio's idle channel detection delay. This ability to sleep directly translates to energy saved, and hence longer network lifetime. On the other hand, since CSMA with RTS/CTS is still totally random access, nodes can never sleep without risk of missing a packet transmitted to them.

The area where the RTS/CTS exchange outperforms qMAC's burst exchange is in reducing the number of *exposed nodes* in the network, thereby increasing the number of packet transmission opportunities in the network. In the remainder of this section we quantify this disadvantage of qMAC. A node k is said to be exposed w.r.t. a transmission from node i to node j if it is prohibited from transmitting by an RTS/CTS (reservation burst/suppress burst) exchange, though its transmission would not interfere with the reception of i's transmission at j. For the RTS/CTS exchange, a node $k \in \mathcal{N}_i^1 - \mathcal{N}_j^1$ is potentially exposed whereas for the qMAC burst exchange a node $k \in \mathcal{N}_i^2 - \mathcal{N}_j^1$ is potentially exposed. We calculate the "exposed area" with reference to Figure 7, which shows the areas of interest when node a transmits to node b. In the figure, A_1 contains the nodes in $\mathcal{N}_a^1 - \mathcal{N}_b^1$, A_2 contains the nodes in $\mathcal{N}_a^1 \cap \mathcal{N}_b^1$, A_3 contains the nodes in $\mathcal{N}_b^1 - \mathcal{N}_a^1$ and A_4 contains the nodes in $\mathcal{N}_a^2 - (\mathcal{N}_a^1 \cup \mathcal{N}_b^1)$. The nodes in A_3 are hidden from a and suppressing interfering transmissions from these nodes is the goal of both RTS/CTS and qMAC burst exchanges. The nodes in A_1 are potentially exposed by the RTS/CTS exchange, and along with nodes in A_4 are potentially exposed by the qMAC burst exchange. To quantify the penalty of qMAC in terms of potentially exposed nodes we label as *Suppression Factor* the ratio of exposed nodes with qMAC to exposed nodes with RTS/CTS. Assuming a spatially uniform node density, the Suppression Factor is equivalent to the ratio of the respective exposed areas A_E^{qMAC} and $A_E^{RTS/CTS}$. Since $r_a = r_b = r$ for RTS/CTS the exposed area is equal to the hidden area regardless of the distance x separating a and b, and after some trigonometry and algebra can be expressed as

$$A_E^{RTS/CTS} = r^2[\pi - 2cos^{-1}(\frac{x}{2r}) + sin(2cos^{-1}(\frac{x}{2r}))].$$

For the qMAC case, the exposed area is simply

$$A_E^{qMAC} = 3\pi r^2.$$

We plot Suppression Factor versus node transmitter/receiver separation x in Figure 8.

From the figure we see that the RTS/CTS exchange provides a more targeted suppression of potentially interfering transmitters, especially when the distance x between transmitter and intended receiver is small. For example when x/r is 0.1, qMAC creates an exposed area nearly 50 times larger than RTS/CTS, implying a huge network throughput penalty. However, this is not as serious a problem as it first appears since any decent routing algorithm should attempt maximize its forward progress toward the destination (within link quality constraints) implying that qMAC should be operating in the part of the curve where x is closer to r.

Finally, we consider the implementation complexity of RTS/CTS and qMAC burst exchange. RTS and CTS signals are actual packets that contain structured data which must be interpreted by the MAC. On the other hand, qMAC bursts need only be differentiated from each other - they can't be mistaken for data since they can only occur during the access quanta portion of the data slot which is known at all nodes in the contention region because of the synchronization. Thus, they need only be distinguishable bit patterns that could even be fixed in hardware. An alternative is to use a Quaternary FSK radio, with two frequencies used for the qMAC reservation and suppress bursts and the other two used to modulate data bits on the carrier.

3.2.2 Comparison with TDMA

TDMA and qMAC share several common requirements and characteristics. Each requires node synchronization to create a slotted channel and a node coloring algorithm to designate a unique opportunity for channel access to each node in a contention region, so both bear the overhead burden of these mechanisms. While both offer contention-free channel access, with TDMA each data slot is reserved for a particular transmitter in the contention region, while with

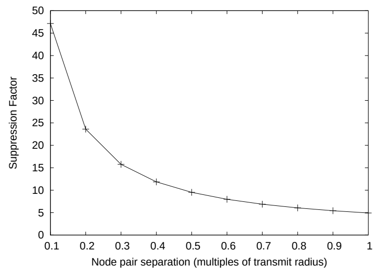

Figure 8: The RTS/CTS exchange provides a more targeted suppression of potentially interfering transmitters, especially when the distance between transmitter and intended receiver is small. qMAC creates a fixed size area of suppression with respect to this separation distance, while RTS/CTS is adaptive.

qMAC each node in the contention region has a reservation to transmit in any particular data slot. The design choice here for qMAC trades off maximum throughput, by adding access quanta at the beginning of each data slot for increased adaptability to the variable traffic patterns prevalent in sensor networks. As events are detected in disjoint parts of the sensor field, traffic bursts are likely to arrive in a contention region via different nodes. A static allocation of data slots to nodes can not adjust to this situation and may result in some nodes being backlogged while other nodes have nothing to send in their slot. This can have severe implications on the delay performance of a particular traffic stream entering a contention region. In particular, comparing with qMAC we see that the inter-transmit opportunity time for a node in TDMA is $(M-1)\frac{L}{C}$, where M is the number of nodes in the contention region and $\frac{L}{C}$ is the packet transmit time, while for qMAC it is $\max(q)\tau+\frac{L}{C}$, where $\max(q)$ is the largest required integer to do access quantum assignment as described in Section 2.3. In general, for qMAC to show a delay advantage over TDMA in such an unbalanced traffic scenario we need

$$\max(q)\tau + \frac{L}{C} < (M-1)\frac{L}{C}. \quad (1)$$

As a numerical example, substituting values from Section 3.1 of $\tau = 886\mu s$ (CC1000 radio) and $\max(q) = 45$ for the grid topology in Figure 5 ($M = 25$), the inequality in Equation 1 holds $39870 < 360288$, giving a delay reduction of $(360288 - 39870)/360288 \approx 89\%$ for qMAC. Further, we observe from Figure 6 that $\max(q)$ is linear in M, over the tested range. Therefore we can write

$$\ell M\tau + \frac{L}{C} < (M-1)\frac{L}{C}$$

$$\ell < \frac{L}{C}\frac{(M-2)}{M}\frac{1}{\tau} \approx \frac{L/C}{\tau}.$$

The large delay reduction evident in the previous numerical example shows that this requirement on ℓ is not difficult to meet in the state of the art wireless sensor networks.

3.2.3 Comparison with Z-MAC

As as a representative of the existing class of hybrid CSMA/ TDMA MAC proposals, Z-MAC also seems to be closest in spirit to qMAC. In Z-MAC, nodes in a two hop neighborhood are uniquely assigned to a data slot to facilitate collision-free channel access, just as in TDMA. To achieve this end, Z-MAC employs DRAND [21] for slot assignment and TPSN [20] and a local synchronization scheme borrowed from RTP/RTCP to ensure nodes in a contention region have a consistent notion of slot boundaries [17].

The node assigned to slot k is termed the "owner" of slot k, while all other nodes in the contention region are "non-owners" of slot k. During periods of "high" contention, the region enters HCL mode and each node is only allowed to transmit in slots it owns. When channel contention falls below some threshold the region enters LCL mode and non-owners are allowed to contend for data slots. Slot owners still get priority access by backing off a random time in the interval $[0, T_o]$, while non-owners back off randomly in $[T_o, T_{no}]$. T_o and T_{no} are selected based on a stochastic analysis to maximize effective throughput [17]. It must be emphasized that channel access by non-owners during LCL mode is still arbitrated by simple CSMA.

Thus, in HCL mode the hidden transmitter problem and the τ problem are eliminated by use of slot synchronization and smart slot assignment. Additionally, by allowing non-owners to transmit in each slot in LCL mode Z-MAC provides additional flexibility compared to TDMA in handling spatially non-uniform traffic generation. However, in Z-MAC's LCL mode the $M-1$ non-owners contend for the channel using simple CSMA and are thus susceptible to the hidden transmitter problem and the vulnerability due to τ .

We calculate the probability of collision within a contention region for a given data slot when in LCL mode as

$$P_c = 1 - P_{\bar{c}(\tau)\cap\bar{c}(H.T.)},$$

where $P_{\bar{c}(\tau)\cap\bar{c}(H.T.)}$ is the probability that neither a τ collision or a hidden transmitter collision occurs in the slot. To show that Z-MAC's LCL mode is non-negligibly affected by the hidden transmitter and τ problems, it is sufficient to show that probability of collision is bounded away from zero at packet generation rates consistent with low contention. Thus, we provide a lower (and upper) bounds on P_c as follows:

$$1 - min(P_{\bar{c}(\tau)}, P_{\bar{c}(H.T.)}) \le P_c \le 1 - P_{\bar{c}(\tau)}P_{\bar{c}(H.T.)},$$

Due to space constraints, in lieu of the probability development, Figure 9 numerically summarizes the results for the case of $M = 3$ (two non-owners in a two hop neighborhood contending for the given available data slot). In general, the collision probability is sensitive to the ratio $\frac{\tau}{T_{no}-T_o}$ and to the channel load. For the results shown in Figure 9, we have assumed independent Poisson packet generation at the two contenders, and based on constants for the backoff windows given in [17] we use $\frac{\tau}{T_{no}-T_o} = 0.092$.

Figure 9 clearly shows a non-negligible probability of collision for Z-MAC LCL mode as the packet generation rate rises above zero. In contrast, qMAC provides collision-free channel access for all nodes in all data slots through the proper assignment of access quanta to each node in a contention region.

4. CONCLUSION

In this paper we have described qMAC, a channel access scheme to eliminate packet collisions in wireless sensor network due to the idle channel detection delay problem and the hidden transmitter problem. We show that the strategy of using τ-length access quanta to arbitrate wireless channel

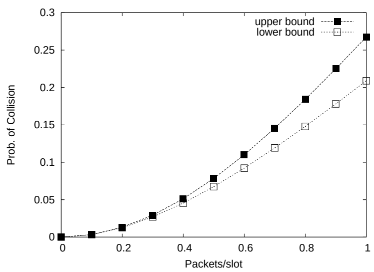

Figure 9: Analytical bounds on the collision probability of Z-MAC in LCL mode for the case where M=3

access has the effect of reducing the maximum achievable throughput compared to a saturated TDMA network (up to approximately 25-75% depending in the radios we evaluated). Further, the exposed node area compared to CSMA with RTS/CTS is increased by a factor of at least 5 with qMAC, and this penalty increases as the distance between the transmitter and intended receiver decreases.

However, we believe that in wireless sensor networks maximum achievable throughput is less important than minimizing packet loss (reducing wasted energy) for a large class of applications, and the design choices made in qMAC reflect this view. Perhaps the closest relative of qMAC among the hybrid sensor network MAC protocols is Z-MAC. When confronted with this throughput/loss tradeoff, Z-MAC chooses to optimize throughput by its choice of backoff timers for slot owners and non-owners, allowing collision between non-owners when in LCL mode. qMAC makes the opposite choice, introducing bandwidth overhead in the form of access quanta in each data slot to eliminate collision under all network conditions. Yet, this overhead is minimized by choosing τ, the minimum time scale at which the radio hardware can interact with its neighbors, as the quantum length.

For future work we plan to verify the collision performance benefit of qMAC with a testbed implementation. This effort will involve the modification of an existing distributed coloring protocol to add the additional constraints required for qMAC (see Section 2.3), and empirical analysis of the synchronization skew offered by existing protocol implementations. Such experimentation will also allow us to confirm that qMAC offers a performance benefit in the face of time-varying, irregular radio propagation patterns.

Acknowledgment

The authors thank B.-J. Ko and J. Tan for helpful discussions on graph coloring, and A. Constantinides and the anonymous reviewers for suggestions to improve the paper.

5. REFERENCES

[1] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks. In *Proc. of the 6th Annual Int'l Conf. on Mobile Computing and Networking*, pages 56–67, Aug 2000.

[2] A. Woo and D. Culler. A Transmission Control Scheme for Media Access in Sensor Networks. In *Proc. of the 7th Annual Int'l Conf. on Mobile Computing and Networking*, pages 221–235, Jul 2001.

[3] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. System Arthitecture Directions for Network Sensors. In *Proc. of the 9th Int'l Conf. on Arch. Support for Programming Languages and Operating Systems*, pages 93–104, Nov 2000.

[4] W. Ye, J. Heidemann, and D. Estrin. An Energy Efficient MAC Protocol for Wireless Sensor Networks. In *Proc. of the 21st Int'l Annual Joint Conf. of the IEEE Computer and Communications Societies (INFOCOM 2002)*, pages 1567–1576. New York, June 2002.

[5] T. V. Dam and K. Langendoen. An Adaptive Energy-efficient MAC Protocol for Wireless Sensor Networks. In *Proc. of 1st ACM Conf. on Embedded Networked Sensor Systems*, pages 171–180. Los Angeles, Nov 5-7 2003.

[6] V. Rajendran, K. Obraczka, and J.J. Garcia. Energy-efficient, collision-free medium access control for wireless sensor networks. In *Proc. of 1st ACM Conf. on Embedded Networked Sensor Systems*, pages 181–192. Los Angeles, Nov 5-7 2003.

[7] A. Woo and D. Culler. Taming the Underlying Challenges of Reliable Multihop Routing in Sensor Networks. In *Proc. of the 1st ACM Conf. on Embedded Networked Sensor Systems*, pages 14–27. Los Angeles, Nov 5-7 2003.

[8] Chipcon. CC1000 Single Chip Very Low Power RF Transceiver. http://www.chipcon.com.

[9] RFM. TR1000 916.50 MHz Hybrid Transceiver. http://www.rfm.com.

[10] Tinyos homepage. http://www.tinyos.net.

[11] V. Bharghavan, A. Demers, S. Shenker, L. Zhang. MACAW: a media access protocol for wireless LAN's. In *Proc. of the ACM Conf. on Communications Architectures, Protocols and Applications*, pages 212–225. London, 1994.

[12] A. S. Tanenbaum. Computer Networks. Prentice Hall PTR, Upper Saddle River, NJ, 1996.

[13] R. Szewczyk, E. Osterweil, J. Polastre, M. Hamilton, A. Mainwaring, and D. Estrin. Habitat Monitoring with Sensor Networks. In *Communications of the ACM*, Vol. 47, No. 6, pages 34–40, June 2004.

[14] G. Tolle, J. Polastre, R. Szewczyk, D. Culler, N. Turner, K. Tu, S. Burgess, T. Dawson, P. Buonadonna, D. Gay, and W. Hong. A Macroscope in the Redwoods. In *Proc. of the 3rd ACM Conf. on Embedded Networked Sensor Systems*, pages 51–63. San Diego, Nov 2-4 2005.

[15] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan. Understanding the Real-World Performance of Carrier Sense. In *Proc of ACM SIGCOMM Workshop on Experimental Approaches to Wireless Network Design and Analysis*, Philadelphia, Aug 22 2005.

[16] J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. In *Proc. of the 2nd ACM Conf. on Embedded Networked Sensor Systems*, pages 95–107. Baltimore, Nov 3-5 2004.

[17] I. Rhee, A. Warrier, M. Aia, and J. Min. Z-MAC: a Hybrid MAC for Wireless Sensor Networks. In *Proc. of the 3rd ACM Conf. on Embedded Networked Sensor Systems*, pages 90–101. San Diego, Nov 2-4 2005.

[18] K. Jamieson, H. Balakrishnan, and Y. Tay. Sift: A MAC protocol for event-driven wireless sensor networks. Technical Report MIT-LCS-TR-894, MIT Laboratory for Computer Science, 2003.

[19] M. Maroti, B. Kusy, G. Simon and A. Ledeczi. The Flooding Time Synchronization Protocol In *Proc. of the 2nd ACM Conf. on Embedded Networked Sensor Systems*, pages 39–49. Baltimore, Nov 3-5 2004.

[20] S. Ganeriwal, R. Kumar, and M. B. Srivastava. Timing-Sync Protocol for Sensor Networks. In *Proc. of the 1st ACM Conf. on Embedded Networked Sensor Systems*, pages 138–149, Nov 5-7 2003.

[21] I. Rhee, A. Warrier, and L. Xu. Randomized dining philosophers to TDMA scheduling in wireless sensor networks. Technical Report, Computer Science Department, North Carolina State University, Raleigh, NC, 2004.

[22] J. Polastre, J. Hui, P. Levis, J. Zhao, D. Culler, S. Shenker, and I. Stoica. A Unifying Link Abstraction for Wireless Sensor Networks. In *Proc. of the 3rd ACM Conf. on Embedded Networked Sensor Systems*, pages 76–89. San Diego, Nov 2-4 2005.

[23] B. Awerbuch. Randomized distributed shortest paths algorithms. In *Proc. of 21st ACM Symp. on Theory of Computing*, pages 490–500. Seattle, 1989.

[24] Z. J. Haas, and J. Deng. Dual Busy Tone Multiple Access (DBTMA) A Multiple Access Control Scheme for Ad Hoc Networks. IEEE Transactions on Communications, pages 975–985, Vol. 50, No. 6, Jun 2002.