



10-2006

MAX: A Maximal Transmission Concurrency MAC for Wireless Networks with Regular Structure

Rahul Mangharam

University of Pennsylvania, rahulm@seas.upenn.edu

Raj Rajkumar

Carnegie Mellon University

Follow this and additional works at: http://repository.upenn.edu/mlab_papers

 Part of the [Electrical and Computer Engineering Commons](#)

Recommended Citation

Rahul Mangharam and Raj Rajkumar, "MAX: A Maximal Transmission Concurrency MAC for Wireless Networks with Regular Structure", . October 2006.

Suggested Citation:

R. Mangharam and R. Rajkumar, MAX: A Maximal Transmission Concurrency MAC for Wireless Networks with Regular Structure" IEEE Third International Conference on Broadband Communications, Networks and Systems (IEEE BROADNETS), San Jose, CA, October 2006.

©2006 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

This paper is posted at Scholarly Commons. http://repository.upenn.edu/mlab_papers/5
For more information, please contact libraryrepository@pobox.upenn.edu.

MAX: A Maximal Transmission Concurrency MAC for Wireless Networks with Regular Structure

Abstract

Multi-hop wireless networks facilitate applications in metropolitan area broadband, home multimedia, surveillance and industrial control networks. Most applications require high end-to-end throughput and/or bounded delay. Current single-hop networks primarily employ random access link-layer protocols such as Carrier Sense Multiple Access (CSMA). These perform poorly in the multi-hop regime and provide no end-to-end QoS guarantees. The primary causes are uncoordinated interference and unfairness in exclusive access of the shared wireless medium. Furthermore, random access schemes do not leverage spatial reuse effectively and require routes to be link-aware. MAX is a time division multiplexed resource allocation framework for multi-hop networks with practical architectures for node scheduling algorithms. MAX tiling delivers optimal end-to-end throughput across arbitrarily large regularly structured networks while maintaining bounded delay. It outperforms CSMA-based random access protocols by a factor of 5-to-8. The MAX approach provides network services including: flexible uplink and downlink bandwidth management, deterministic route admission control, and optimal gateway placement. MAX has been implemented on IEEE 802.15.3 embedded nodes and a test-bed of 50 nodes has been deployed both indoors and outdoors.

Keywords

Scheduling Algorithms for Embedded Wireless Networks

Disciplines

Electrical and Computer Engineering | Engineering

Comments

Suggested Citation:

R. Mangharam and R. Rajkumar, "MAX: A Maximal Transmission Concurrency MAC for Wireless Networks with Regular Structure" IEEE Third International Conference on Broadband Communications, Networks and Systems (IEEE BROADNETS), San Jose, CA, October 2006.

©2006 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

MAX: A Maximal Transmission Concurrency MAC for Wireless Networks with Regular Structure

Rahul Mangharam & Raj Rajkumar

Department of Electrical & Computer Engineering, Carnegie Mellon University, Pittsburgh, PA. U.S.A.
{rahulm, raj}@ece.cmu.edu

Abstract–

Multi-hop wireless networks facilitate applications in metropolitan area broadband, home multimedia, surveillance and industrial control networks. Most applications require high end-to-end throughput and/or bounded delay. Current single-hop networks primarily employ random access link-layer protocols such as Carrier Sense Multiple Access (CSMA). These perform poorly in the multi-hop regime and provide no end-to-end QoS guarantees. The primary causes are uncoordinated interference and unfairness in exclusive access of the shared wireless medium. Furthermore, random access schemes do not leverage spatial reuse effectively and require routes to be link-aware. MAX is a time division multiplexed resource allocation framework for multi-hop networks with practical architectures for node scheduling algorithms. MAX tiling delivers optimal end-to-end throughput across arbitrarily large regularly structured networks while maintaining bounded delay. It outperforms CSMA-based random access protocols by a factor of 5-to-8. The MAX approach provides network services including: flexible uplink and downlink bandwidth management, deterministic route admission control, and optimal gateway placement. MAX has been implemented on IEEE 802.15.3 embedded nodes and a test-bed of 50 nodes has been deployed both indoors and outdoors.

Keywords: Multi-hop wireless networks, sensor networks, medium access controller, scheduling algorithms, topology

I INTRODUCTION

Multi-hop wireless mesh networks provide a distributed network organization where a service provider may place routers (or nodes) in an arbitrary topology as all nodes are interconnected by wireless links. Unlike traditional single-hop point-to-multipoint networks based on cellular architectures, multi-hop mesh networks require no infrastructure and facilitate flexible deployment as demand increases. In addition, as the density of routers is increased, the distance between routers is reduced to potentially provide higher link data rates. The network structure of interest here is of multiple wireless router nodes communicating across one or more hops to at least one gateway. This structure may be applied to metropolitan area broadband (IEEE 802.16 [1]), home multimedia (IEEE 802.11e, 802.15.3a [2, 3]), surveillance and industrial control (IEEE 802.15.4 [4]) networks. The goal is to deliver high end-

to-end throughput with bounded latency.

The two central problems with multi-hop wireless mesh networks are (a) granting users exclusive access to the shared wireless channel as all nodes operate within the broadcast medium and (b) effectively leveraging spatial channel reuse due to each node’s limited transmission range. It is therefore necessary to determine the *duration* a node should transmit (resource allocation) and *when* it should transmit for that duration (node scheduling). We define a node *schedule* as a sequence of fixed-length time slots where transmissions assigned to the same time slot do not collide. Determining a resource allocation with a minimum length schedule is NP-complete for multi-hop wireless networks with arbitrary topology and hence does not scale [5, 6].

The focus of this paper is on providing a theoretical resource allocation framework based on node scheduling algorithms in fixed multi-hop wireless networks with regular structure. We emphasize the key properties of our approach by simulating networks across a large dynamic range of demands and demonstrate the feasibility through protocol implementation and deployments. We first provide an overview of the problem of uncoordinated wireless link contention and then formally state the goal of resource allocation and scheduling.

A. Uncoordinated Link Contention

Single-hop random access protocols such as CSMA attempt to transmit a packet as soon as it is enqueued [7]. For example in Fig 1(a), each node along the chain is only able to communicate with its immediate neighbors. If, for example, a 10MB file is to be sent from a source node A to a destination node E across multiple hops, every intermediate forwarding node will contend in an uncoordinated manner with the previous two hop and next two-hop forwarding nodes. For example, once node A successfully sends one packet to node B using a single-hop MAC, it attempts to send the next packet without waiting for B to forward the first packet to C. By trying to send the next packet, node A thwarts the continued transmission of the previous packets it sent to B. An opportunistic local optimization to maximize the per-hop throughput is detrimental to the overall end-to-end throughput. Multi-hop CSMA performance studies [8] show




Figure 1. (a) Uncoordinated contention between packets of the same flow caused by single-hop random access MAC protocols. (b) Maximal concurrent transmission with a 3-slot transmission schedule

the maximum end-to-end throughput is $1/8$ and $1/24$ of the link throughput for a line of nodes and a 2-D grid of nodes respectively. The CSMA binary exponential back-off policy results in severe unfairness and complete starvation of flows over the same or neighboring links, and is unable to provide any end-to-end throughput or delay guarantee [9]. Furthermore, nodes with smaller degree (e.g. node A) experience lesser contention and tend to transmit more aggressively thus wasting a larger fraction of time backing-off as nodes with higher degree form the bottleneck.

B. Maximal Concurrent Transmission

In Fig. 1(b), we observe that when node A sends a message to B, if either node B or node C is transmitting, then node A's transmission will be unsuccessful. Therefore, for a chain of nodes with a 1-hop transmission and interference range, successful concurrent transmissions must be spaced by 3 hops [11]. To deliver high network utilization, it is essential to exploit the spatial reuse so that the *maximal set of concurrently transmitting nodes* is determined. We define a *k-order concurrent transmission set* to be a set of nodes that are mutually k or more hops away from each other. A *maximal k-order concurrent transmission set* is a k -order concurrent transmission set to which no other node of the network can be added. Nodes may transmit concurrently if they are mutually at least a distance of k -hops from each other such that k is greater than twice the communication range.

C. An Overview – The MAX Approach and Results

For a network with arbitrary topology, the cardinality of the maximal transmission sets can be different. A flow's transmission opportunity at a particular node depends on the duration each maximal transmission set is active, the number of such sets the node belongs to and the node degree with active flows. The inherent difficulty in arbitrating fair resource allocation and spatial reuse to nodes with different degree motivates us to solve the problem for networks with regular structure first and then generalize to networks with less regularity. As determining the maximum independent set of an arbitrary graph is NP-complete [10], we focus on network topologies that lend themselves naturally to minimal schedules.

In a network with regular structure, nodes are placed at regular intervals, each with uniform node degree as in Fig. 2. By assuming a regular topology, any locally optimal resource allocation solution with a feasible schedule is valid with the same properties for the entire network. The uniform node degree ensures the cardinality of all maximal transmission sets is the same and hence the optimal slot schedule is also fair.

DEFINITION: A *MAX Tile* is a periodic and symmetric network structure consisting of a group of nodes such that at most one




Figure 2. MAX Tiles and their respective tessellations overlaid over network topologies with (a) 3, (b) 4 and (c) 6 neighbors. Only one node may transmit at a time in each tile

node at a given position in each tile may transmit concurrently.

An illustration of MAX Tiles and their tessellations are given in Fig. 2. The nodes in the networks presented have a transmission and interference range of one-hop. Here a node in each tile is least 3-hops away from a corresponding node in the same position in all neighboring tiles. By assigning a synchronized slot schedule to nodes within a tile, we are able to schedule the entire tessellation of tiles and render the network interference-free. We summarize below the attractive properties of MAX Tiles of size n nodes, with regular structure:

Link Layer Properties:

- (a) A network with a tessellation of nodes scheduled as MAX Tiles is a tiling of maximal concurrent transmission sets of nodes. A MAX tiling results in an optimal schedule in terms of minimal length of the slot schedule for networks with transmission and interference range limited to one-hop.
- (b) Generalized capabilities in (a) for grid networks with any transmission and interference range.

Network Layer Properties:

- (a) Routes are interference-free from neighboring and non-overlapping flows with deterministic admission control.
- (b) Optimal gateway placements for shortest path routing are derived from MAX tiling.

Service Layer Properties:

- (a) Flexible uplink/downlink bandwidth asymmetry control.
- (b) Support for multiple path fine-granularity flows for enhanced end-to-end throughput.

MAX tiling outperforms CSMA-based random access protocols by a factor of 5-to-8 in end-to-end throughput while providing bounded delay. For flows with random source-destination pairs, the average network utilization exceeds 95%.

The rest of the paper is organized as follows. Section II presents related work followed by formal description of MAX Tile-based scheduling in Section III. In Section IV we generalize our results for interference-dominated networks and describe MAX bandwidth management mechanisms. Section V discusses gateway placement followed by performance analysis of routing in Section VI. Finally section VII presents our implementation followed by the conclusion.

II RELATED WORK

While the theoretical maximum throughput of random access MAC protocols such as p -persistent CSMA is 87% of

the offered link rate for single-hop communication [7], the upper bound for a multi-hop networks in contrast is 48.5% for CSMA and 35% for slotted-ALOHA [11]. In practice, the IEEE 802.11 standardized single-hop protocol achieves about 14% of the link rate for a one-dimensional chain of nodes [8]. The performance of 802.11 degrades further to 8% for a grid of nodes with horizontal flows. Our approach to maximizing transmission opportunity has a similar basis as [11, 12] but applies it to time-synchronized regular structures which do not require knowledge of relative node positions. In the multi-hop regime, as it is necessary to arbitrate transmission among all nodes within communication range and their neighbors, 802.11's opportunistic operation has been shown [9] to be unfair and starve TCP flows.

Node and link scheduling are the two primary approaches for resource allocation and scheduling in multi-hop wireless networks. In [13], a max-min fair resource allocation is proposed. The node connectivity graph of the network is resolved into a flow contention graph connecting all interfering links. From this, the network is decomposed into cliques of conflicting links and transmission durations are assigned to links in the descending order of the clique degree. While a fair resource allocation (i.e. transmission duration) is achieved, finding the slot schedule assignment still remains an NP-complete problem. It is therefore practical and desirable to decouple flow routing and link scheduling to jointly solve the resource allocation and node scheduling problems.

In [6], Ramanathan and Lloyd propose node and edge scheduling algorithms for tree, planar and arbitrary graphs with a distance-2 matching constraint for wireless networks. Their results provide an 8-10% improvement over greedy algorithms. For networks with gateways, they show node scheduling of tree networks is superior to link scheduling.

Silvester and Kleinrock [14] provide a comprehensive study of multi-hop scheduling with slotted-ALOHA for networks with regular structure. The maximum throughput for one-dimensional line networks is proportional to $1/e$ as with single-hop slotted-ALOHA. For grid networks, spatial reuse allows a capacity proportional to the square root of the number of nodes in the network. They also show that networks with smaller node degrees deliver higher average end-to-end throughput.

Simulation studies for on-demand routing protocols over 802.11 [15] observe that for a moderate-density connected network with 100 randomly placed nodes, the largest concurrent transmission set was of 7 nodes. This resulted in the available per-node throughput to be 50 times smaller than the apparent capacity. The offered loads in other routing studies [16] are limited to about 60Kbps despite using 2Mbps radios. Furthermore, traditional ad hoc routing protocols such as AODV and DSR [16] do not provide any guarantees of the quality of the route and the interference it experiences from neighboring and non-overlapping flows.

III MAX RESOURCE ALLOCATION AND SCHEDULING FRAMEWORK

In this paper, we restrict network topologies to regular structures to provide optimal and fair spatial reuse. By exploiting the periodic node distribution and uniform node degree, we provide one solution of the distance-2 graph

coloring problem. This enables us to decouple flow routing and link scheduling and jointly solve the link resource allocation and scheduling problems.

While regular topologies may not always be achievable in practice, they provide an upper bound of end-to-end throughput and lend insight to the arbitration of fairness and spatial reuse. Irregular networks with node clusters of high density suffer from increased interference while low density clusters waste potential spatial diversity. All topologies presented in this paper describe logical, rather than physical, network connectivity graphs.

A. Preliminaries

The algorithms presented are applicable to regularly structured networks of any degree. For the sake of notational convenience, we focus on a rectangular grid of nodes, G , specified by a rectangular coordinate system. Each node has a uniform number of neighbors, N , at a logical communication distance, C and interference distance I , $I \geq C$. A tile is defined as a periodic and symmetric group of M nodes that form a tessellation across the grid. Each node within a tile may transmit for a specific duration at specified intervals.

B. Network Assumptions

To obtain analytical results about performance in a scheduled multi-hop wireless network, certain assumptions must be made about that network.

- (A1) Every node's transmission range is equal to its interference range of one-hop Euclidean distance. i.e. $C = I = 1$.
- (A2) Transmission range is limited to a fixed number of neighbors along rectangular coordinates and transmissions along a diagonal are not permitted. This could be achieved by using directional antennas.
- (A3) The topology of the network is known.
- (A4) All nodes are time synchronized and given fixed time slots that repeat at a fixed interval of M slots.
- (A5) All nodes transmit at a fixed link data rate over a single shared channel.

Not all of these assumptions are absolutely necessary, and the effects of relaxing them for practical network architectures will be discussed in Section VII.

C. MAX Tiling – Temporal Representation

In order to facilitate concurrent transmissions, it is necessary to maintain the 3-hop rule between transmitters. For 1-dimension networks as in Fig 1(b), we observe that nodes A and D may transmit concurrently. In order to maximize network capacity, consequently all nodes that are a multiple of 3-hops away from node A may transmit concurrently. With global time synchronization, a 3-slot Time Division Multiple Access (TDMA) cycle enables data to be pipelined in both directions within the network. The effective end-to-end data rate is $1/3^{\text{rd}}$ the available link rate with delay bounded to M^*h , where h is the number of hops along the path.

To extend the application of the 3-hop rule to networks in 2-dimensions, a *MAX Tile* structure is defined as in Fig 2. Each MAX Tile consists of a single node and its nearest neighbors resulting in a tile size of M nodes, where $M = N + 1$. More specifically, for a grid network each tile consists of $M=5$

nodes as shown in Fig. 3(a). Each node is assigned a fixed time slot to transmit in a TDMA cycle with M slots. Fig. 3(a) shows an example time slot allocation from 0 to $M-1$ within a MAX Tile. Each slot permits the transmission of one or more frames consisting of the payload and frame acknowledgements (ACK) to all neighbors for all successfully received frames from the prior cycle. The concatenated cumulative ACK ensures the protocol is not bidirectional and the hidden terminal problem needs to be resolved only at the frame receiver. While the time slot assignment within a tile may be arbitrary, it is necessary that the sequence of transmissions within a tile is consistent across all tiles in the network.

The key property of a tessellation of MAX tiles is that it ensures each transmitter is exactly three hops away from the closest concurrent transmitting node. Schedule assignment for a grid network based on rectangular coordinates may be described, for instance, by slot numbers in the x and y directions. If the time slot assigned to the top-left node $(0, 0)$ is 0 (Fig. 3(a)), then the assignment, s , to any node may be described by:

$$s = [x + (2C + 1)y] \bmod M \quad (1)$$

This ensures that nodes with the same slot are separated by a distance greater than twice the communication range. The end-to-end throughput of the network is $1/M$ or $1/5$ the available link rate. We now establish the correctness and optimality of MAX Tiles.

Theorem 1: The slot assignment, s , for a multi-hop wireless grid network described by $s = [x + (2C + 1)y] \bmod M$ is collision-free.

Proof: Assume two nodes, i and j , located at (x_i, y_i) and (x_j, y_j) respectively, are assigned the same time slots:

$$\begin{aligned} s_i &= x_i + (2C + 1)y_i \\ s_j &= x_j + (2C + 1)y_j \end{aligned}$$

A collision occurs if $s_i = s_j$:

$$(x_i - x_j) + (2C + 1)(y_i - y_j) = 0 \quad (2)$$

and the following conditions hold:

$$|x_i - x_j| + |y_i - y_j| \leq 2C + 1 \quad (3)$$

$$|x_i - x_j| + |y_i - y_j| \geq 1 \quad (4)$$

In (2), both nodes transmit during the same time slot and $(x_i - x_j)$ is a multiple of $2C + 1$. The first condition ensures the nodes are spaced by at most $2C + 1$. Together with (2), the first condition requires either $|x_i - x_j| = 0$ or $|y_i - y_j| = 0$. If $|x_i - x_j| = 0$, then by (2), $|y_i - y_j| = 0$ thus violating the second condition. The slot assignment rule described is collision-free. \square

In order for the resource allocation scheme to be optimal (i.e. resulting in the shortest schedule while being fair), the slot scheduling must determine the maximal concurrent transmission sets and ensure every node has the same duty cycle with no residual idle time. Such a network delivers the maximum attainable throughput and is fair. Let T_i be the set of nodes that are i -hops away from a . If all nodes of a maximal $(2C+1)$ -order independent set transmit, and these are the only transmitting nodes in the network, all of their transmissions will be successfully received collision free.

Theorem 2: For any maximal $(2C+1)$ -order concurrent transmission set S , no node not in the set can transmit without causing interference with a reception of at least one node in S .

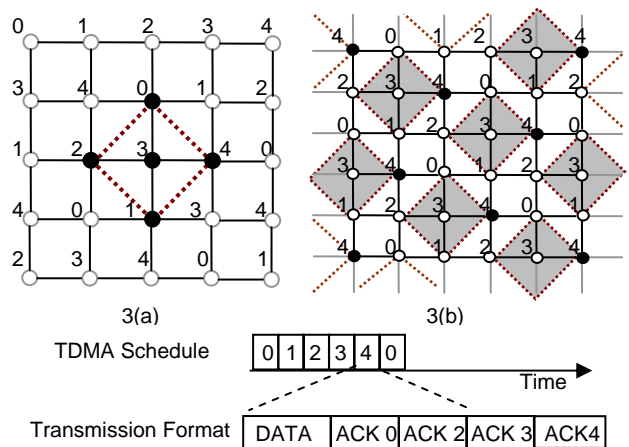


Figure 3. (a) A MAX Tile marked over a grid of nodes with assigned time slots. (b) A tessellation of MAX Tiles

Proof: Suppose a was such a node. It is clear that $T_1 \cap S = \emptyset$ (null set) otherwise a would be transmitting and receiving at the same time. If $T_{2c} \cap S \neq \emptyset$ then this implies there is a node that simultaneously receives a signal from a and from a node in S . Thus $T_{2c} \cap S = \emptyset$, but this implies that S is not maximal since a is at least $2C+1$ hops away from every node in S . \square

Fig. 3(b) illustrates a tiling where the right-most node of each tile is active in slot 4. For broadcast scheduling based on node scheduling, this is the tightest packing of tiles such that all nodes transmitting in a particular slot are at a Manhattan distance of $2C+1$ from the closest concurrent transmitters.

The slot assignment scheme based on (1) requires nodes to be aware of their direction relative to their neighbors. We employ that approach for ease of mathematical representation in a rectangular coordinate system. In practice, however, a distributed tile replication algorithm which assigns schedules based on a seed MAX Tile and its tiling is employed. The tile replication algorithm is detailed in Section VII.

D. Theoretical and Practical Significance of Regular Topologies

We focus on using time-synchronized networks with regular structures to derive the upper-bound of link layer throughput in multi-hop mesh wireless networks. Regular structures provide a uniform node density so a locally optimal scheme is also globally optimal and may be deployed in a distributed manner. Regular structured networks have been employed in several theoretical network studies [14, 22, 23] and therefore form a basis for comparison of network capacity and protocol efficiency.

In [17] we describe a TDMA topology control method to prune an arbitrary physical connectivity graph to a logical topology with uniform degree. For a given physical topology with average degree $\geq d^*$, a connectivity graph is determined from which a spanning tree is extracted by the gateway. Following this, links between nodes are incrementally marked as active until the maximum degree of each node is at most d^* . MAX Tiling is then performed on the (partially) uniform network topology by allocating time slots only to active links. This scheme does not require knowledge of node positions. Another scheme for reducing arbitrary physical connectivity graphs to near-regular structured topologies is

presented in [25]. The authors first use the node positions and transmission range to determine the connected dominating set of the network and then assign slots to nodes based on a distance-2 node coloring heuristic.

From a practical deployment perspective, mesh networks with regular structures have been deployed in military experiments in [24, 25], involving over 1,000 nodes, and are feasible for large-scale networks in factory and warehouse ceilings, parking lots and cargo areas. While it may not be practical to control the network topology in most deployments, we demonstrate the feasibility in a test-bed deployment, described in Section VII, to highlight (a) the theoretical properties of MAX in regular topologies and (b) an economical and easy method to achieve global time synchronization.

IV GENERALIZED MAX

In this section we relax the assumptions presented for basic MAX tiling in Section III to more realistic regular network topologies. As the communication range increases, the node degree increases, resulting in an enlarged tile. The end-to-end data rate varies inversely with the tile size and is given by $1/M$. On the other hand, as the interference range increases, the tile size increases at a slower rate and the throughput decreases due to reduced spatial reuse. Finally, the basic MAX tiling approach may be extended from a “one-size-fits-all” uniform TDMA slot assignment to control link asymmetry and, for example, dilate the bandwidth of select (e.g. gateway) nodes.

A. Communication Range-Dominated Networks

We relax assumption **A1** to cater to dense networks with limited power control capability. If the communication range, C , spans a distance greater than 1-hop, the tile size is $(\sum_{i=1}^C i * N) + 1$, where N is the uniform number of neighbors, and the slot assignment is then described by:

$$s = (x_i + (2C + 1)y_i) \cdot \text{mod}((\sum_{i=1}^C i * N) + 1) \quad (5)$$

Fig. 4(a) illustrates a MAX Tile with a 2-hop communication and interference range (i.e. $C = I = 2$). A tiling continues to ensure all concurrent transmitters are at least a distance of $2C+1$ hops apart. As nodes at a $C+1$ Manhattan distance do not lie within the $2C$ Euclidean distance range, the shape of the MAX Tile is consistent.

B. Interference Range-Dominated Networks

We relax assumption **A2** to cater to networks whose interference range exceeds their transmission range. In Fig. 3(b) we assumed a node is unable to communicate to the nearest diagonally located node. We observe that each node is to receive from the nearest transmitter while being unable to receive from the nearest diagonally located concurrent transmitter. While this has been achieved in outdoor experiments with 802.15.4 nodes [17], it imposes a tight 6dB signal-to-noise (SNR) sensitivity budget to differentiate between a neighbor and a diagonal node’s transmission. For networks without such power control capability, the interference range, I , may be greater than the communication range. For the case when $I > C$, it is necessary to separate




Figure 4. (a) A MAX Tile for nodes with 2-hop communication and interference range. (b) Slot assignments for nodes with 1-hop communication range and 2-hop interference range.

concurrent receivers by a distance of at least $I+1$. Consequently the tile size is given by $(I+1)^2+1$ with a slot assignment described by:

$$s = (x_i + (I + 1)y_i) \cdot \text{mod}((I+1)^2 + 1) \quad (6)$$

In Fig. 4(b) we illustrate the slot assignment for $C=1$ and $I=C+1$. Concurrent transmitters are separated by 5 hops (i.e. $2I + 1$) and the tile size is 10 (i.e. $M = (I+1)^2$). Using this approach, a larger tile size (i.e. $M = (I+1)^2$) is required and results in an exponential reduction of per-node network capacity. Furthermore, the application of (6) requires knowledge of the relative position of nodes in the x and y directions.

For networks with $I > C$, we may use multiple communication channels and retain the $M = N+1$ MAX Tile size. Fig. 5 illustrates the use of two communication channels to tessellate MAX Tiles with $C=1$ and $I=C+1$. In a given time-slot, the transmitting node and its neighbors tune to one of the two pre-assigned frequencies. During the node X 's time slot (Fig. 5(a)), all its four neighbors tune to one of the two channels pre-assigned to X . During the next time slot (Fig. 5(b)), when node Y transmits, all Y 's neighbors tune to its pre-assigned channel.

By alternating the channel used for tiling, tiles operating on the same channel are always separated by distance of $2(I+1)$. We observe that nodes switch between the two channels on a time-slot basis to ensure that all neighbors of the current transmitter are tuned to the same channel. By employing the 2-channel tiling, the end-to-end network throughput is still $1/M$, where $M=N+1$ as opposed to $1/(I+1)^2$. From a practical

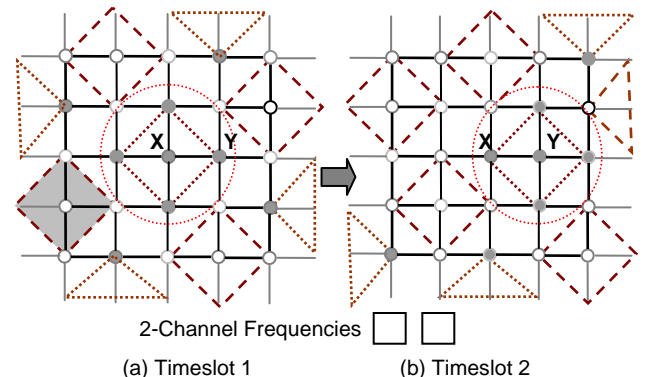


Figure 5. A 2-channel TDMA scheme for nodes with 1-hop communication range and interference with diagonal neighbors.

perspective, the multi-channel MAX tiling approach is more efficient than employing larger tiles specified by (6) as most wireless standards such as 802.11 and 802.15.4 support multiple channels. Our implementation of a network of 802.15.4 nodes is capable of supporting 15 channels in the worldwide ISM band [4]. In Section VII we show that this scheme is practical and robust.

C. Bandwidth Management with Link Asymmetry

MAX Tiles, with $C=I=1$, offer an uplink transmission opportunity of $1/M$ and downlink reception from all neighbors of N/M during each TDMA cycle. Thus, for a grid network, each node has an uplink bandwidth of $1/5^{\text{th}}$ and downlink capability of $4/5^{\text{th}}$ the link rate. As all users are not alike and have different uplink/downlink demands, it is useful to control the bandwidth asymmetry. For a network with few gateways and several end users, the $1/M$ bandwidth reduction is very limiting and restricts the gateways' maximum outgoing throughput to the entire network. Likewise, for end users, the maximum downlink throughput is more critical than uplink. To control the bandwidth asymmetry, we apply a simple transform by artificially setting a gateway's neighbor count to be L rather than M , where $L > M$. The gateway may use the additional $(L-M)$ slots for transmission. The additional $(L-M)$ slots are said to be accounted for by *virtual nodes*.

For example in Fig. 6, a gateway with initial $1/5$ uplink and $4/5$ downlink capability, may set its tile size to be 8 rather than 5. By using the additional 3 time slots due to virtual nodes in addition to its assigned time slot, the gateway's transmission duty cycle is increased to $3+1$ time slots every cycle. As $M=8$, the gateway now has $4/8$ slots for transmission and $4/8$ slots for reception, while end user nodes have $1/8$ uplink and $4/8$ downlink bandwidth asymmetry. Due to the tiling, all nodes that occupy the same position within a tile as the gateway have the same uplink/downlink ratio. For example, if the gateway is a center node in its tile, all center nodes in the network will have the $1/2$ uplink and $1/2$ downlink ratio. Thus by varying the number of virtual nodes, we may conveniently adjust the bandwidth asymmetry to suit the network's requirements.

V MAX ROUTING AND GATEWAY PLACEMENT

A. Interference-free Routing with Bounded Delay

Once the multi-hop MAC slot assignments are resolved by the tile replication procedure (described in Section VII), the network has been initialized and the routing procedure may be executed. Unlike CSMA, as all node transmissions are collision-free, there is no interference from overlapping and




Figure 6. MAX Bandwidth Asymmetry control increases gateways uplink ratio from 20% to 50%

non-overlapping flows. For example, consider flows 1 and 2 in Fig. 7. If flow 1 is started first under the CSMA regime, upon starting flow 2, which requires a higher rate, flow 1's throughput is reduced and may starve due to interference [9]. Under the MAX TDMA regime, non-overlapping flows do not interfere and enjoy the maximum end-to-end throughput.

For overlapping flows, assuming equal distribution of bandwidth, a node's throughput is given by $1/(M \cdot F)$ where F is the number of flows traversing through the node. The routing objective is significantly simpler and is to minimize the maximum overlap of flows across all nodes along the path. This problem is similar to VLSI global routing with k -overlaps in a rectilinear grid [18], where k is the number of metal layers. In Fig. 7, we observe TCP flows, 3 and 4, overlap and therefore enjoy only half the offered throughput.

B. Routing Enhancements: Multiple Path and SuperNodes

Routing schemes based on a shortest-path criterion result in congestion at the center of the grid due to a large number of route overlaps. When the total required data rate of all flows passing through a node is $1/M$ of the link data rate, the node is said to be *fully-utilized*. As the offered throughput with the MAX approach is deterministic, flows requesting routes across fully-utilized nodes are *blocked* by the route admission control policy. In order to maximize the network capacity, the blocking probability must be minimized. We employ two mechanisms to achieve this.

1) Multiple Path Routing: In order to reduce the probability a flow is blocked along a path, the source splits a single high-rate flow into multiple flows with a lower data rate requirement. Intuitively, we can relate this to the analogy of filling a container with stones. If the stones are big, there will be gaps in the container. To fill those gaps better, the approach should be to break the stones into smaller pieces. In Fig. 8(a) we observe a flow which requires full (i.e. $1/M$) link capacity and fully utilizes all nodes along its path. This creates a partition between the two halves of the grid such that no routes can be created across this partition. On the other hand, in Fig. 8(b) the same flow is split into two half-rate flows and will not partition the grid. The net effect of employing finer granularity flows is a reduction in the overall blocking rate and an increase in the offered network capacity. As in the case of any multiple-path routing scheme, the end-points will incur an overhead for segmentation and reassembly.




Figure 7. Routing of overlapping and non-overlapping flows




Figure 8 (a) Overlapping flows cause route blocking
(b) Multi-path flows prevent route blocking

2) SuperNodes: An alternative to reducing the flow blocking rate is by employing SuperNodes. SuperNodes are nodes with one or more additional time slots than other nodes in each MAX Tile. As described in Section IV.C, a node in every tile may acquire additional time slots at the cost of increasing the overall TDMA cycle duration. The effect of SuperNodes is to provision additional capacity required for flow overlaps which may otherwise be blocked due to a node being fully utilized. SuperNodes are analogous to vias in VLSI routing.

Consider for example, the flow with full bandwidth (i.e. 1/6 the link rate when SuperNodes are granted an additional time slot) requirement running through the centre of the grid as in Fig. 9(a). It partitions the grid as all nodes along its path are fully utilized. Hence flows that originate from the top of the grid and destined to the other side will not be able to find a route to cross over. The idea of the SuperNode mechanism is to let some nodes have one extra time slot so that flows are allowed to intersect at such nodes. As shown in Fig. 9(b), the center nodes of all MAX Tiles are converted to SuperNodes with a transmission opportunity of 2/6 as opposed to 1/6 for the remainder of nodes in the tile. Although a flow with full bandwidth requirement cuts through the grid, flows originate from one side of the grid are still able to overlap at a SuperNode and are not blocked due to the additional time slot.

C. Optimal Gateway Placement

For most network applications such as metropolitan area broadband, home multimedia, surveillance and industrial control there is a need to communicate with an external entity via a gateway portal. Within the context of a multi-hop wireless network, the gateway is a node that interfaces between nodes in the network and the external entity. As the maximum link throughput is determined by the resource allocation methods mentioned earlier, it is desirable to minimize the number of overlapping flows to avoid congestion. This is achieved by employing shortest path routing while load balancing the flows across the available gateways.

For a four-node network in Fig 10(a), flows from each node may be distributed evenly across all links as shown in the flow matrix in Fig. 10(b). Using this approach recursively, it is




Figure 9 (a) Overlapping flows cause route blocking
(b) SuperNodes prevent route blocking




Figure 10 (a) A basic 4-node network. (b) A load balanced flow matrix

possible to determine the load-balanced flow matrix of a regular network with n nodes. By placing the gateways evenly as shown in Fig. 11, we are able to minimize the maximum path lengths to the nearest gateway in a load balanced network. As each gateway services roughly a equal number of nodes, the network capacity is the highest for such a gateway distribution. The gateway placement in Fig. 11(a) may be recursively used in Fig. 11(b), Fig. 11(c) and likewise in larger networks with additional gateways. By applying this hierarchal tiling of gateways followed by MAX tiling, we evenly distribute the load across the network and provide optimal network capacity due to min-max route lengths.

VI PERFORMANCE ANALYSIS

In this section we present numerical results for simulations comparing the multi-hop network performance with CSMA and TDMA-based MAX Tiling. The results corresponding to CSMA were determined over a 1 Mbps 802.11 link using the ns-2 network simulator. The inter-frame spacing and frame structure is consistent across both the 802.11 and MAX MACs. All other results used a MAX TDMA-based simulator that we built.

A. End-to-End Throughput

In order to compare the end-to-end throughput across multiple hops with CSMA and TDMA-based MAX Tiling, we simulated a 1-dimension chain of nodes with a single flow. In Fig. 12(a) we observe the throughput for 64 byte packets as the length of the chain is increased. The throughput for the 802.11 MAC decreases rapidly to 1/20 of the offered link rate while the MAX MAC offers a steady throughput that is 1/4 the link rate. The throughput of the MAX MAC is lower than the theoretical 1/3 of the link rate due to header and guard time overheads.

In Fig. 12(b) the end-to-end throughput for a 12x12 grid of nodes is presented. All flows are routed horizontally across the grid with no overlaps. We observe that the throughput of the MAX MAC is slightly less than the expected 1/5 of the link throughput. The MAX MAC outperforms the 802.11




Figure 11. Optimal gateway placement for Min-Max path lengths. Placements for 1, 4 and 5 gateways are marked in solid black. The shaded regions mark the nodes routed to the given gateway.

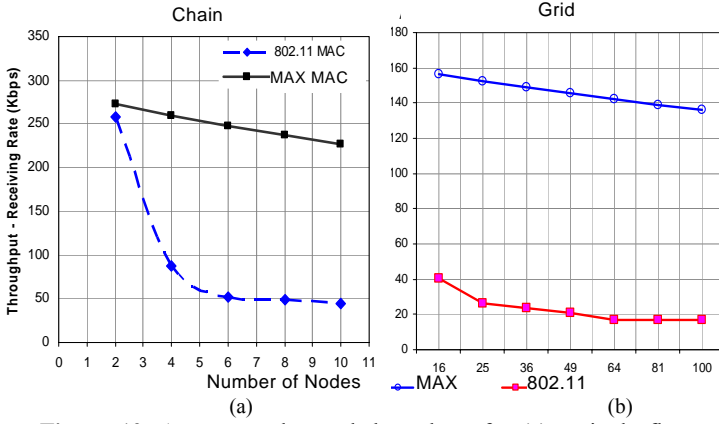


Figure 12. Average end-to-end throughput for (a) a single flow along a chain of nodes and (b) parallel horizontal flows in a grid

MAC by a factor 5-8x. Furthermore, it is important to note that the end-to-end delay offered by the MAX MAC is bounded.

B. Network Utilization

We define the network utilization as the ratio of the aggregate link capacity that is utilized by the routed flows. The average network utilization lends insight into the overall efficiency of the network and the number of flows that may be routed across it. In this test, flows are routed from a randomly chosen source and destination pair in a 7x7 grid. For the routing we employ an exhaustive search by which if a set of routes is feasible, then the shortest route is selected. Each flow requires a constant bit rate (CBR) equivalent to 50% of the maximum possible end-to-end throughput (i.e. 1/10 the link rate). The experiment was repeated 50 times. The results are consistent across various network sizes. We do not compare network utilization with CSMA as the throughput saturates for small networks and the per-link utilization is always below a small fraction of the available link capacity as observed in Fig 12(b).

In Fig. 13(a), we observe the average network utilization approaches 70%. As the number of flows routed across the network increases, the rate at which the utilization increases diminishes. This is due to the fact that as the network gets congested, certain under-utilized nodes are unreachable due to blocking.

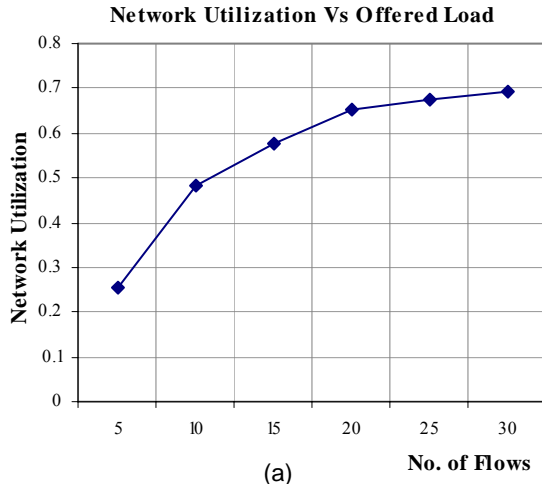


Figure 13. (a) Avg. network utilization for a grid network.

C. Effect of SuperNodes

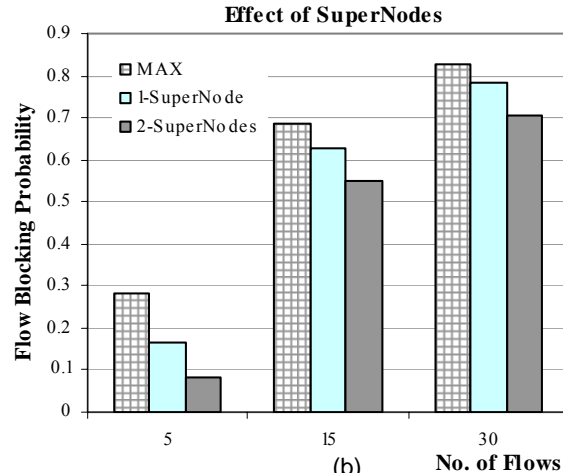
Fig. 13(b) presents the blocking rate as the number of flows is increased across the same network. We analyze the performance of the MAX MAC without SuperNodes, with one SuperNode per tile and with two SuperNodes per tile. Over the 7x7 node network, there were 8 and 17 SuperNodes for the latter two cases. We observe that as the number of flows increases, the blocking rate naturally increases due to congestion. It is interesting to note that the ability of SuperNodes to allow twice as many flows to overlap lowers the blocking rate significantly. This illustrates the usefulness of SuperNodes in increasing the network utilization.

While SuperNodes are assigned a relatively larger number of slots, it is at the cost of increasing the overall number of slots in the TDMA cycle. SuperNodes cause other nodes in the network to have a relatively lower throughput. Thus, assigning SuperNodes with a large number of additional slots will be detrimental to the overall network capacity. From our simulations of routes between randomly selected source-destination pairs, we find that 2 SuperNodes per tile offers the best results.

D. Effect of Multiple-Path Routing

We now look at the benefits of multiple-path routing with finer granularity flows. For this test, a 7x7 node network is used with an offered link rate of 2Mbps. As in the earlier case, the routing scheme employs an exhaustive search by which if a set of routes is feasible, then the shortest route is selected. Three types of flows are routed across random source-destination pairs. The first set of flows requires 100% of the maximum end-to-end throughput (i.e. 400kbps). The next test split the offered load into 200Kbps and likewise into 100Kbps flows. As the flows were split, the offered load was the same but the number of finer granularity flows increased. For example, for an offered load of 8Mbps, there were 40 flows each with a 200Kbps requirement or 50% of the maximum end-to-end throughput of 1/5 the 2Mbps link rate.

In Fig. 14, we observe that employing multiple path routing with fine granularity flows provides a significant benefit. This occurs due to a lower blocking rate achieved as a consequence of evenly spreading the load across the network.



(b) Max. blocking probability is reduced by use of SuperNodes.

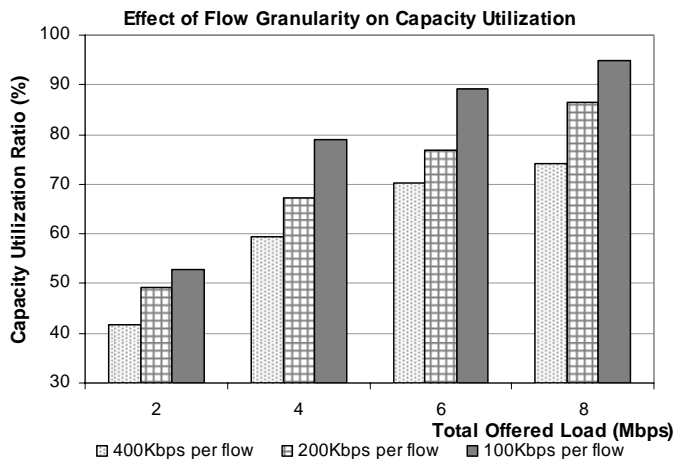


Figure 14. Effect of per-flow link data rate on the maximum network utilization

To summarize, our simulation study shows that the MAX MAC outperforms the CSMA-based 802.11 MAC by a factor of 5-8x on average for line and grid mesh networks. SuperNodes provide a consistent and significant reduction in flow blocking rate when compared to basic MAX. Finally, the use of multiple path routes for finer granularity flows enhances network utilization considerably.

VII IMPLEMENTATION

In this section we discuss some practical deployment issues of MAX and experiences from a 30-node network deployment.

A. Distributed Tile Replication

We propose a simple tile replication algorithm to schedule a network with a known regular topology. The tile size is determined either as a function of the maximum number of neighbors (most congested links) or after using a topology pruning scheme described in Section III.D.

Tile replication is started with a seed tile composed of nodes with pre-set identities and pre-set slot assignments for each tile member-node as in Fig. 15. Simple rules for information exchange and tile replication have been defined in [17] to assign each node its slot schedule. No node transmits until it knows its identity unambiguously. During each step of the neighbor discovery and tile replication phase, each node transmits its identity and the identity of its known neighbors and their slot assignments. For irregular structures, our simulation in Fig. 15 shows that over 95% of the node identities are determined when up to 10% of the nodes in the network were randomly selected and shut down.

The MAC does not assume any routing information and it is expected that the routing process is executed after the MAC resource assignments are complete. As each node is given an equal opportunity to transmit, flows routed in any direction across the network will receive similar throughput and delay.

B. Embedded Wireless Mesh Deployment Experiences

To verify the feasibility of MAX Tiling, a 30-node network of embedded nodes was deployed both in an open field and a 3-storey campus building. Each node (see Fig. 16), developed by us, consists of an IEEE 802.15.4 transceiver, an Atmel ATMEGA32 microcontroller [19, 20] and 6 sensors (i.e. PIR,

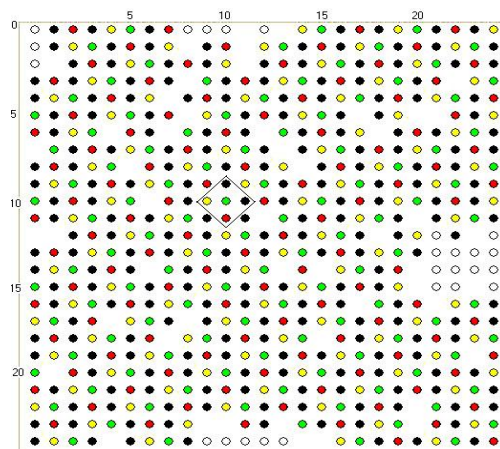


Figure 15. Simulator output of MAX Tile replication algorithm in a grid with holes

light, temperature, audio, acceleration and humidity). In addition, to provide global time synchronization, each node is equipped with an amplitude modulation (AM) receiver for indoors or an atomic clock receiver for outdoors. A carrier current-based AM transmitter [21] is plugged into an electrical outlet in the building and uses the power grid as an antenna to radiate a periodic (e.g. 5 sec) global time synchronization beacon within the building. Upon reception of the beacon, the AM circuit wakes activates the microcontroller and 802.15.4 transceiver. Each node transmits and receives in its allocated 4ms slots and returns to sleep mode when inactive. Protocol implementation and experimentation details are further described in [17].

We conducted three experiments in an open field. In the first experiment, we determined the minimum spacing between concurrent transmitters by placing three nodes: a receiver (RX), transmitter (TX) and jammer in a line. With the jammer off, we first measured the RX-TX distance for stable and successful reception. We notice that 100% of the 2000 transmitted packets are received up to a distance of 10m at power level 6. We repeated the experiment with the jammer at different distances from the receiver. We observed that the jammer has no effect beyond 20m and a concurrent transmitter can be placed at a minimum distance of 30m.

In the second experiment, we profiled the radiation pattern and packet reception behavior for the on-board antenna. A receiver was rotated on top of a servo motor in the middle of a field and the signal strength and packet reception success rate were logged. A transmitter placed 8m away transmitted 100 packets at every 0.5 degree turn of the receiver. In Fig. 17, we observe that the packet reception success rate on a radial axis is almost uniform. This indicates that equidistant nodes in a regular topology experience similar performance in all directions.




Figure 16. FireFly time-synchronized embedded nodes

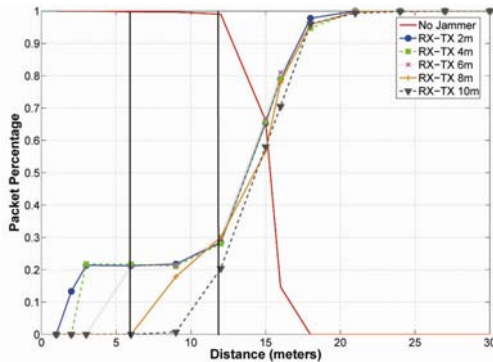


Fig 17. Packet success rate in the presence of a jammer

In the third test, we laid out a grid of nodes with fixed slot assignments as described in Fig. 3(b) with 10m spacing between neighboring nodes. We observed the average successful packet reception ratio for 2,000 transmitted packets to be only 68% due to interference from diagonal nodes located 14m away. We repeated the test with a grid as in Fig. 5 employing the 2-channel scheme described Section IV.B and observed an average successful packet reception ratio of 98%. Although, we do not expect such topology control to be available in most deployments, these experiments verify the feasibility and performance of regular structured networks. We are currently extending our work to irregular networks within an 8-storey campus building.

VIII CONCLUSION

This paper presents MAX, a time division multiplexed resource allocation framework for multi-hop networks with practical architectures for node scheduling algorithms. Unlike traditional random access protocols, the MAX MAC delivers optimal end-to-end throughput by identifying the maximal set of concurrent transmitters across arbitrarily large regularly structured networks while maintaining bounded delay.

The MAX MAC outperforms CSMA-based random access protocols by a factor of 5-to-8 in terms of end-to-end throughput. The MAX approach provides network services including: (a) flexible uplink and downlink bandwidth management, (b) deterministic route admission control, and (c) optimal gateway placement. MAX delivers an average network utilization of 94% and is scalable to irregular networks.

While regular topologies may not always be achievable in practice they provide an upper bound of end-to-end throughput and lend insight to the arbitration of fairness and spatial reuse. As future work we will extend MAX to less regular networks and explore energy efficient tiling schemes.




Fig 18. Packet reception for omni-directional antenna

REFERENCES

1. IEEE 802.16-2004 Standard for Local and Metropolitan area networks Part 16: Air Interface for Fixed Broadband Wireless Access Systems. Oct. 2004
2. IEEE 802.11e/D4.4, Draft Supplement to Part 11: MAC Enhancements for QoS, June 2003.
3. IEEE 802.15 High Rate Alternative PHY Task Group (TG3a) for Wireless Personal Area Networks
4. IEEE 802.15.4-2003 Specifications for Low Rate Wireless Personal Area Networks Part 15.4: Wireless MAC & PHY
5. Hari Balakrishnan, et al. "The Distance-2 Matching Problem and its Relationship to the MAC-layer Capacity of Ad hoc Wireless Networks" IEEE Jnl. on Selected Areas in Comm., Vol. 22, No. 6, pp. 1069–1079, Aug. 2004
6. S. Ramanathan, E. L. Lloyd. "Scheduling Algorithms for Multi-hop Radio Networks" ACM Trans. on Net., Vol. 1-2, 1993.
7. L. Kleinrock and F. A. Tobagi, "Packet Switching in Radio Channels: Part I - Carrier Sense Multiple Access Models and their Throughput-Delay Characteristics," IEEE Trans. Comm., vol. 23, no. 12, pp. 1400–1416, Dec. 1975.
8. Jinyang Li, et al. "Capacity of Ad Hoc Wireless Networks", MOBICOM 2001.
9. Xu S, Saadawi T. "Does the IEEE 802.11 MAC Protocol Work well in Multi-hop Wireless Ad hoc Networks?" Comm. Mag., IEEE, Vol. 39, No. 6., pp. 130-137. 2001.
10. Garey, M. R. and Johnson, D. S. Computers & Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman, '83.
11. Randolph D. Nelson, Leonard Kleinrock: Maximum Probability of Successful Transmission in a Random Planar Packet Radio Network. INFOCOM 1983: 365-370
12. S. Kulkarni and M. Arumugam. "Collision-free Communication in Sensor Networks", Symp. on Self-Stabilizing Sys, June 2003.
13. X. L. Huang and B. Bensaou, "On Max-min Fairness and Scheduling in Wireless Ad-hoc Networks: Analytical Framework and Implementation," MobiHoc, pp. 221-231, 2001.
14. J. Silvester and L. Kleinrock, "On the Capacity of Multi-hop Slotted-ALOHA Networks with Regular Structure," IEEE Trans. Comm., vol. 31, pp. 974–982, Aug. 1983.
15. S. Das, C. Perkins, and E. Royer. "Performance Comparison of Two On-demand Routing Protocols for Ad hoc Networks." In INFOCOM-00, pages 3-12, 2000.
16. J. Broch, et al. "A Performance Comparison of Multi-hop Wireless Ad hoc Network Routing Protocols", ACM MobiCom 1998.
17. A. Rowe, R. Mangharam, and R. Rajkumar, "RT-Link: A Time-Synchronized Link Protocol for Energy Constrained Multi-hop Wireless Networks" Carnegie Mellon Tech. Report, Nov 2005. http://andrew.cmu.edu/~rahulm/Research/Pubs/rtlink_tr1105.pdf
18. J. Hamkins and D. J. Brown. "Routing in a Rectangle with k-ary Overlap" Proc. of the 2nd Great Lakes Symp. on VLSI. 1992
19. Chipcon Inc., <http://www.chipcon.com/>. Chipcon CC2420
20. Atmel Corporation, <http://www.atmel.com/> ATMEGA128L
21. Radio Systems 30W TR-6000 AM Transmitter Data Sheet. <http://radiosystems.com>
22. G. Mergen and L. Tong, "Capacity of regular ad hoc networks with multipacket reception," Proc. Allerton Conf. USA Oct 2001.
23. V. Mhatre, et al., "Design of Surveillance Sensor Grids with a Lifetime Constraint" in EWSN, Berlin, January 2004.
24. A. Arora, et al., "ExScal: Elements of an Extreme Scale Wireless Sensor Network," in IEEE RTCSA. 2005.
25. V. Naik, A. Arora, P. Sinha, H. Zhang, "Sprinkler: A Reliable and Energy Efficient Data Dissemination Service for Wireless Embedded Devices," in RTSS, Miami, USA. 2005