SmartNode: Achieving 802.11 MAC Interoperability in Power-efficient Ad Hoc Networks with Dynamic Range Adjustments

Edmond Poon, Baochun Li*

Abstract

The standard CSMA/CA based IEEE 802.11 protocol assumes that each node uses a certain fixed (or maximum) transmission power for the transmission of each packet. However, a MAC protocol with power adjustments can have significant benefits towards better power conservation and higher system throughput through better spatial reuse of spectrum. In this work, we propose a power efficient MAC layer algorithm, SmartNode*, that is compatible with the basic RTS-CTS-DATA-ACK MAC protocol defined in IEEE 802.11. Our algorithm uses the minimum required power level for the transmission of data packets, which requires special handling for the exchange of control packets. Compared with previously proposed power-controlled MAC protocols, our algorithm does not require multiple data channels at the physical layer, so that it is able to inter-operate with regular nodes running the existing IEEE 802.11 MAC protocol. Through extensive performance evaluations, we have demonstrated that our proposed algorithm is effective in a power-controlled ad hoc network — it is able to increase system throughput while conserving power with dynamic power adjustments.*

1 Introduction

In multi-hop wireless ad hoc networks, nodes relay packets for other nodes if the destination is out of the transmission range of the source. Since each node is energy-constrained and packet transmission consumes a certain amount of power, energy conservation may be achieved by *dynamically adjusting transmission ranges* on the fly at each node. The benefits of such dynamic power adjustments in ad hoc networks are three-fold: (1) It can significantly reduce power consumption rates; (2) total system throughput for an end-to-end flow may increase due to the spatial reuse of spectrum; and (3) contention among flows sharing the same channel may be minimized, since the total number of neighbor nodes involved is reduced. By using such a strategy, each node can make local decisions to adjust its transmission power to cover the minimum area, with a target of minimizing the number of nodes that it can reach.

However, the reduction of transmission power on each node may introduce serious problems in traditional RTS-CTS-DATA-ACK (CSMA/CA) MAC layer protocols, such as the IEEE 802.11 MAC standard. We present an example to illustrate one of the problems. At the MAC layer, proper handshaking between the sender and the receiver to acquire the floor before initiating the data packet transmission is important to avoid packet collisions caused by the well-known hidden station problem [1]. One solution proposed by the IEEE 802.11 MAC protocol [2] is to use RTS/CTS to reserve the channel, and to use data acknowledgments to guarantee successful transmission of data packets. The general assumption is that the transmission power levels of all nodes are uniform.

With a relaxed assumption that transmission power levels may be dynamically adjusted, nodes are no longer commutative. A successful RTS/CTS handshake may not be able to successfully reserve the channel, since some of the nodes are not able to overhear part of the ongoing packet transmissions, including the RTS/CTS exchange. In other words, even with successful RTS/CTS exchanges to establish a flow, the strategy of dynamic power adjustments tends to increase the possibility of introducing *hidden stations*, one of the problems that the original RTS/CTS proposal seeks to address. Such a problem yields the original RTS/CTS solution ineffective. With this ineffective RTS/CTS exchange to reserve the floor, packets may fail to be delivered in a flow with a small power range, suffering from the high levels of interference from a competing flow with a higher range.

To address these issues, the focus of this paper is to propose a novel algorithm, referred to as **SmartNode**, to extend and improve the existing IEEE 802.11 Medium Access Control protocol, so that it performs correctly and effectively when dynamic power adjustments are used to (1) conserve power; (2) enable better spatial reuse of spectrum; and (3) reduce contention levels within the local area. The core of SmartNode includes a scheme to determine the appro-

*Edmond Poon and Baochun Li are affiliated with the Department of Electrical and Computer Engineering, University of Toronto. Their email addresses are {*epoon,bli*}*@eecg.toronto.edu*.

priate power level to initiate transmission of a certain flow, as well as an improved RTS/CTS handshaking algorithm for appropriate channel reservation. The gist of the idea is that, nodes running the SmartNode algorithm attempt to derive the minimum transmission power to reach another node from the power strengths of received packets, and record the calculated information in a local table. Subsequently, data transmissions operate at the minimum power level, while the RTS/CTS exchanges may use a slightly higher transmission power than the minimum level, in order to compensate a disadvantaged flow for the sake of fairness, if higher levels of interference from other flows are encountered.

One of the most favorable properties of SmartNode is that, it assumes that all nodes use a single wireless channel at the physical layer for both data and control packet transmissions. This property is not available in many previous works in the area [3], which require an additional busy tone channel to perform functions such as collision detection. It is our belief that, for the purpose of arbitrating collisions in an ad hoc network with dynamic power adjustments and different ranges, the benefits of opening (or assuming) a new control channel — near-optimal arbitration between flows — may not outweigh the disadvantages of reducing the (already limited) channel capacity in wireless networks. With a single data channel, SmartNode is able to mix and interoperate well with nodes running the existing IEEE 802.11 MAC protocol, while still providing improved performance than an all-802.11 network.

The remainder of the paper is organized as follows. Section 2 describes the problems when dynamic power adjustments are used in a regular ad hoc network with 802.11-based nodes. Section 3 discusses previous work and compares them with SmartNode. In Section 4, we present SmartNode in details. Section 5 evaluates the performance of SmartNode in the ns-2 network simulator, in order to demonstrate the effectiveness of our algorithm compared to IEEE 802.11. Section 6 concludes the paper and discusses future work.

2 Problems

In the IEEE 802.11 MAC protocol, when a network node attempts to transmit a packet, the sender and the receiver initiate the RTS-CTS handshake to notify their neighbor nodes of their upcoming packet transmission. The effectiveness of the RTS-CTS handshake to acquire the floor depends on the general assumption that all nodes transmit using the same transmission range, and thus operate with identical power levels. Only with a coherent transmission range, neighbors are able to hear the reservation requests from the sender or the receiver, so that they may backoff for the specified period of time.

However, as we have briefly illustrated, if each sender dynamically adjusts the transmission power level to the minimum required to reach the receiver, we may not only conserve additional power, but also achieve better spatial reuse by minimizing the interference range as well. In other words, since more network nodes will not overhear other nodes' transmission, they can transmit packets at the same time to utilize the channel more efficiently. The benefits are more significant if the sender and the receiver are close to each other. To illustrate such benefits using an example, we assume $P_{\max}$ is the maximum transmission power and $P_{\min}$ is the minimum transmission power of a node. We denote P_{ij} as the minimum required transmission power for node i to reach node j and assume $P_{ij} = P_{ji}$ for $i \neq j$. Figure 1 has shown that, if a MAC protocol allows each node to use only the minimum required power for data transmission, the channel utilization will be tripled compared with the situation where nodes use the maximum transmission power.

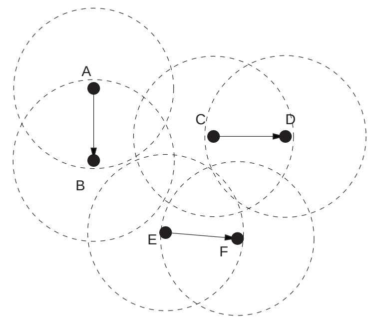

Figure 1. Benefits of spatial reuse of spectrum

With two examples, we present the problems of using unmodified 802.11 in a network with dynamic range adjustments. First, the sender may not be able to correctly receive the ACK packet from the receiver. As illustrated in Figure 2, after node A has initiated a data flow to node B, node C attempts to start another flow to node D. The flow will also be initiated with a successful RTS/CTS handshake, since C is unable to hear the RTS-CTS exchange between A and B. Because of the interference from node C, node A is unable to receive the ACK packet from node B correctly in its ongoing data flow.

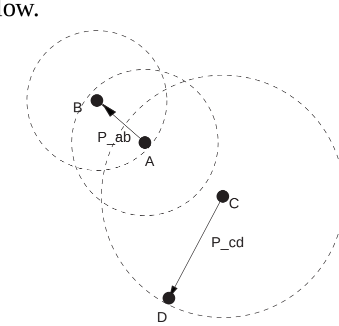

Figure 2. ACK packet fails to reach sender A

Second, there also exists a higher probability of collision among nearby flows due to the incomplete knowledge of the status of existing packet transmissions in the network. Figure 3 illustrates this problem. Nodes C and D initiates a flow first. Node A is unable to hear the RTS-CTS handshake between node C and node D if they are only transmitted at the minimum required power, P_{CD}. If node A attempts to transmit a packet to node B, it will generate a RTS packet, that collides with the ongoing data packets between node C and node D.

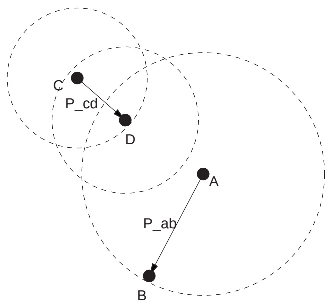

Figure 3. Fairness issue with power consideration

From these two examples, we can clearly derive the problem that, *flows that use lower power levels generally have a higher probability of transmission failure*, since a successful reservation for the channel of a low-power link may not be heard by other nodes that are potentially close enough to disrupt its data exchange (with an interference range that covers either the sender or the receiver of the low-power flow). This is not *fair* to the low-power flows, since it penalizes the good will of power-conserving nodes. In the worst case, some low-power flows may be indefinitely unable to access the channel. We refer to these low-power flows as *disadvantaged* flows. For the sake of fairness, our algorithm is required to determine when a node should start a flow or accept any transmission requests, so that the power consumption can be kept to the minimum, and any newly granted flows only have a low probability of affecting ongoing flows.

To summarize, in order to achieve a balanced trade-off between fairness and the benefits of dynamic power adjustments, any proposed solutions need to achieve the following design objectives: (1) *Fully distributed to each node.* Only local information is used in the computation. Such an algorithm guarantees *scalability* with respect to the network size, and *zero communication overhead* of additional protocol-specific message exchanges across the network. (2) *Fairness.* The reduction of transmission power levels should not introduce disadvantaged flows as previously illustrated. (3) *Backward interoperability.* The algorithm should maintain a close resemblance to the existing IEEE 802.11 standard to ensure compatibility and interoperability in the same network. Modifications and extensions are made only when absolutely necessary.

3 Related Work

Power-aware MAC protocols in 802.11-based wireless ad hoc networks have received much recent attention in the research literature, due to better spatial reuse of spectrum and lower power consumption. Recent work focuses on integrating dynamic power adjustments into the 802.11 RTS-CTS-DATA-ACK handshaking protocol. As one example, Monks *et al.* has proposed a power controlled MAC protocol (PCMA) [3] that utilizes a data channel and a busy tone channel. Their main contribution is to achieve power controlled transmission while still preserving the collision avoidance property of standard multiple access protocols. The drawback is that these papers do not consider the fairness among flows and the compatibility issues with IEEE 802.11.

Unlike previous work [3, 4, 5], SmartNode aims to provide a compatible solution that improves IEEE 802.11. To achieve this goal, it is critical to consider a single wireless channel for both data and control messages. Although the use of a busy tone channel as an indicator for channel access can minimize packet collisions, such approach limits available capacity and does not guarantee interoperability with 802.11 nodes. In this aspect, SmartNode attempts to offer a balanced tradeoff between performance and compatibility. Such property allows easy migration from an all-802.11 network to a SmartNode network.

In addition, our paper studies the fairness issue in using the dynamic transmission power adjustment. SmartNode considers the distribution of the bandwidth to each individual flow and offers certain dynamic tunable parameters to achieve best-effort fairness. In [3], the authors only attempt to propose new MAC protocols to optimize the system throughput without investigating the issue of fairness. Some flows may then be unable to obtain a fair share of the channel due to the multi-hop nature. Through simulation, SmartNode can achieve better fairness than IEEE 802.11.

4 Proposed Algorithm: SmartNode

In this section, we present SmartNode, an algorithm to balance the trade-off between fairness and the benefits of dynamic power adjustments. Within the scope of this work, we assume that each node is equipped with an omnidirectional antenna, and all wireless links are bi-directional. We further assume that an error-free channel is perceived, which may be achieved by appropriate channel coding. We do not consider node mobility in our analysis of the algorithm.

The main idea of the algorithm is that, each node should use local information to determine the minimum required transmission power before initiating a flow to another node. The power levels used by RTS and CTS packets should be determined based on the congestion levels of the environment for the optimal performance.

4.1 Determining the minimum required transmission power

At the physical layer, we assume that nodes are able to adjust the transmission power for outgoing packets in a particular flow, and to determine the power strength of received signals on receiving a packet. If node i transmits a packet to node j, let P_{tr} be the transmission power of node i, P_{re} be the received power at node j, and $\text{SIR}_{\text{thresh}}$ be the minimum signal to interference ratio threshold such that the intended receiver can successfully receive the packet. The following equations must hold for every successful packet transmission:

$$P_{\min} \leq P_{tr} \leq P_{\max}, \text{ and}$$

$$\text{SIR}_{\text{thresh}} N_0 \leq P_{re} \leq P_{\max}$$

where N_0 is the background noise power (including interference). To estimate the transmission power level to reach a particular node, in this chapter we only consider large-scale path loss characteristics in the fading channel model, where the channel gain is given by $1/(|X_i - X_j|^\alpha)$, where X_k is the location of a node k, and α is a parameter greater than 2 (usually between 2 and 5 depending on the physical environment). If we assume the minimum required transmission power levels $P_{ji} = P_{ij}$ for any pair of nodes i and j, we may then estimate the required minimum transmission power to transmit packets *in subsequent flows* from j to i, by observing received packets *in the current flow* from i to j in the channel. For example, assuming the two-ray ground propagation model, the following equation holds for a flow from i to j using transmission power P_{tr}:

$$P_{re} = \frac{P_{tr} G_t G_r h_t^2 h_r^2}{|X_i - X_j|^\alpha} \tag{1}$$

where G_t and G_r are the antenna gain at the sender and the receiver, respectively, and h_t and h_r are the height of the antenna at the sender and receiver, respectively. Assuming, instead, that i transmits to j using the minimum transmission power P_{ij}, the equation becomes:

$$\text{SIR}_{\text{thresh}} N_0 = \frac{P_{ij} G_t G_r h_t^2 h_r^2}{|X_i - X_j|^\alpha} \tag{2}$$

By combining Equation (1) and (2), we can obtain

$$P_{ji} = P_{ij} = \frac{\text{SIR}_{\text{thresh}} N_0}{P_{re}} P_{tr} \tag{3}$$

Again, P_{re} is the received power level at node j, and P_{tr} is the transmission power level at node i.

Assuming a known $\text{SIR}_{\text{thresh}} N_0$, if we can include the transmission power P_{tr} in a *RTS packet* from node i, and measure the received signal strength in a packet received by node j as well, the minimum required transmission power to send a packet from node j *back to* node i can be estimated using Equation (3). To be conservative, in the proposed algorithm, we multiply the calculated P_{ji} by a preset coefficient, μ, that is greater than 1, to ensure that j can reach i even with an inaccurate estimate of the channel fading characteristics. Finally, we believe that it is not appropriate to use the Global Positioning System (GPS) to estimate the minimum transmission power level, since the estimates based on geographical distances are usually far from accurate.

4.2 Modifications to 802.11 RTS-CTS Exchanges

In SmartNode, transmission of each data packet follows the existing IEEE 802.11 MAC protocol, which adopts the RTS-CTS handshaking protocol to acquire the floor. However, the appropriate power levels to be used for either RTS or CTS packets are still open to be determined. One trivial solution is to use the maximum power level for both RTS and CTS packets, and then to use the minimum power level for data packets. However, since each single data packet activates one or more rounds of RTS-CTS handshaking, using the maximum power for both RTS and CTS will significantly reduce the benefits of power conservation. On the other hand, using the minimum power levels for RTS and CTS is not suitable either, since this will introduce severe fairness problems introduced in Section 2. in SmartNode, we attempt to balance such trade-offs. We present our algorithm as follows.

Initially, when a node i has packets in its packet queue, it waits until its backoff counter reaches zero, and then initiates the RTS packet with the maximum transmission power level. As we have noted, for the purpose of determining the minimum transmission power levels, the RTS packet from node i includes the transmission power i uses. On overhearing the RTS packet, any of the neighboring nodes is able to estimate the minimum transmission power based on Equation (3) to reach node i. Each node, say node j, maintains a local table to record the minimum transmission power required to reach i, i.e. $\langle \text{node id}, P_{ji} \rangle$, that it has calculated to reach its neighboring nodes. After an initial converging phase, in a network without node mobility, all nodes will converge to a local table of the minimum power levels to reach neighboring nodes. In this case, for subsequent initiations of new flows, a node will first perform a table lookup. If the lookup fails, it will transmit the RTS packet with the maximum power; otherwise, it uses the results of the lookup as the minimum power level to transmit the RTS packet.

After the RTS packet is successfully received, the receiver will reply with a CTS packet to the sender to begin the packet transmission, as in 802.11. However, so far, it is still an open issue with respect to what is the transmission power level that the receiver should use to transmit the CTS. One option is obviously to transmit at the maximum power;

however, such a scheme works against the goal of power conservation. Even when we assume that such an option is feasible (that the penalty of additional power consumption is acceptable), we still need to include the transmission power level in the CTS packet, so that the sender may use the embedded information to compute the proper minimum transmission power to deliver the subsequent data packets.

Based the results from our performance evaluations of such a scheme, we have found that modifying the CTS packet to include the transmission power of sending CTS will render *interoperability with 802.11 nodes impossible*. This is due to the fact that, since each 802.11 node backs off for the duration of the CTS, DATA and ACK packets, if the CTS packet has a longer header, 802.11 nodes will need further modifications to identify appropriate backoff times.

In SmartNode, we propose that, to transmit the CTS, the receiver should use the *same transmission power* as that is indicated in the RTS packet. The sender is then able to determine the minimum required transmission power to reach the receiver, using the previously proposed method, without any modifications in the CTS packet format. With such an approach, we allow nodes to determine the minimum required transmission power to reach their neighbors using the RTS/CTS/DATA/ACK handshaking, which is compatible with the IEEE 802.11 MAC protocol.

In order to solve the fairness problem illustrated in Figure 3, a higher transmission power level is preferred to enforce all neighbors of both the sender and the receiver to backoff to avoid possible collisions. On the other end, being overly conservative with high power levels discourages spatial reuse of spectrum, since some of the new flows may not be established even though they will not disrupt any ongoing packet transmissions. In SmartNode, we need to have a mechanism to allow packet transmissions to complete successfully, while still able to improve the system throughput. In other words, we need to determine appropriate transmission power levels for RTS/CTS to balance between being conservative and being unfair, since the two goals are incompatible in nature.

The adaptive solution to this problem in SmartNode is to use the minimum power level (which is calculated previously and stored in local tables) to transmit both RTS and CTS packet transmission on the first t attempts, where t is a *dynamically tunable* parameter. If *all* t attempts fail, the power level of both RTS and CTS will be incremented by $\omega(P_{\max} - P_{ij})$ for every additional failure, where ω is a tunable coefficient. In the worst case, every time the sender will use the maximum transmission power to deliver the RTS/CTS packets. With such a strategy, neighboring nodes are notified by the ongoing data transmission and therefore packet collision can be avoided. In addition, even when the propagation environment slightly changes, a source node is still able to send packets to reach its destination with such compensation. Figure 4 illustrates how the SmartNode RTS/CTS handshaking works. Initially, node C and node D use the minimum required transmission power, P_{cd}, as represented by the dotted line. After t unsuccessful packet transmission attempts, node C and node D will increase the transmission power level to increase the success probability of competing with other flows and establishing the connection, since more nodes can hear their RTS/CTS handshaking. Note that such approach can also compensate for those flows which have difficulties to deliver data packets to the destination due to high multipath fading, which is very common in the wireless environment.

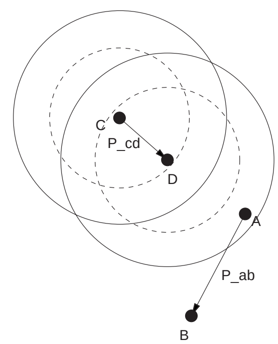

Figure 4. SmartNode: RTS/CTS handshaking

Initially, a small ω and a large t is preferred to enable a better spatial reuse of spectrum for higher system capacity, since such setting will minimize the effective interference region of a packet transmission. On the other hand, such a scenario may introduce higher collision probabilities and thus affect the overall system throughput. In our performance evaluations, we will evaluate different settings of ω and t to determine the effects on fairness among flows and the system performance.

To be specific, the highlights of SmartNode is the ability to adjust transmission power levels to compensate flows that may be unable to access the channel. It integrates the following original contributions, based on standard 802.11: (1) The RTS packet is transmitted with minimized power levels, which is adaptive based on the congestion levels of the environment; (2) the CTS packet is transmitted using the same power as the RTS packet; and (3) the ACK and DATA packets are transmitted with the pre-calculated and pre-stored minimum transmission power levels. In comparison, we believe that previous work using the IEEE 802.11 MAC protocol were not able to use such strategies to achieve better fairness. Instead, the fairness among flows in a 802.11 network is topology-dependent, where flows may starve if they are located in unfavorable locations. We believe that SmartNode offers a balanced tradeoff between the fairness and throughput/power concerns. As follows, we summarize the three steps that we have proposed in the SmartNode algorithm, assuming the sender i attempts to establish a data flow with the receiver j:

Step 1: In order to access the channel, after the backoff counter has expired, i uses the calculated minimum transmission power level to transmit the RTS packet to j, if a lookup is successful from the local table. Otherwise, it uses $P_{\max}$ instead to perform the initial discovery of the appropriate power level to reach j. Notice that the RTS packet includes the transmission power level used by i, such that each neighboring node is able to overhear and calculate the required power to reach i. The results of the calculation will be stored in the local table.

If i retries more than t times to send RTS, it increments the transmission power by $\omega(P_{\max} - P_{ij})$ in order to compete more favorably with other (possibly higher-power) ongoing flows. In the worst case, maximum transmission power will be used.

Step 2: When j receives RTS from i, it replies with a CTS packet if the channel is idle. j will use the *same power level* as the the transmission power of i. In such a way, node i can determine the minimum required transmission power to reach j without incurring any overhead on the CTS packet.

Step 3: When i receives CTS, it starts the transmission of data packet to j with the power level P_{ij}, which is previously calculated. j will reply an ACK packet at the minimum power required (P_{ji}) to reach node i when it receives the entire error-free data packet.

5 Performance Evaluation

We have implemented the proposed SmartNode algorithm as a module within the ns 2.1b8a network simulator. With the assumption that all flows are single-hop in our simulation, we extensively evaluate the performance of SmartNode using comparative studies in three different network scenarios: (1) All nodes run unmodified IEEE 802.11 MAC; (2) all nodes run SmartNode; and (3) mixed deployment of 50% 802.11 nodes and 50% SmartNode nodes. Note that with such settings, the maximum transmission range of a node is approximately 250 meters and an interference range using the maximum transmission power is about 550 meters. The performance metrics we are interested to evaluate are: (1) The *effective throughput* of data flows in terms of bytes per second actually delivered; and (2) the *normalized energy consumption*, which is defined as the amount of energy consumed per byte that is successfully delivered. The throughput represents the Quality-of-Service provided by the network under dynamic power adjustments, while the normalized energy consumption represents the overall power efficiency.

Comparisons under varying offered load

In the first experiment, we randomly deploy 64 nodes in the rectangular region, and initiate 16 Constant Bit Rate (CBR) flows as *offered load* with different bit rates. Our objective is to examine the effects of different offered load levels on the performance of SmartNode vs. 802.11 networks, with respect to both the *effective throughput* and *normalized energy consumption*. The simulation results, respectively, are shown in Figures 5 and 6.

In Figure 5, the throughput of a SmartNode network is considerably higher than a 802.11 network when the offered load is heavy (i.e., a heavily congested network), and roughly identical when the offered load is low. SmartNode may not offer a much higher throughput than 802.11 due to the following practical reasons. (1) One of the design choices of SmartNode — *interoperability* — leads to the design choice of using a single channel, which means that SmartNode may not completely avoid packet collisions. (2) We note that, on average, each node covers a large physical area since we deploy only 64 nodes in a region of 1000 x 1000 m^2, and there is one initiated CBR flow for every four nodes. If node density is higher, a better throughput may be obtained for using SmartNode over 802.11, since each packet transmission of SmartNode covers a smaller area.

In Figure 6, it is clearly demonstrated that SmartNode consumes less power for each successfully delivered packet. When the offered load is high, SmartNode is approximately 50% better than 802.11. In the mixed network that contains both 802.11 and SmartNode nodes, the normalized energy consumption is also slightly lower, with the same effective throughput. Note that there are three main factors that will affect the result of the normalized energy consumption: (1) *throughput*: a higher throughput generally lead to lower normalized energy consumption (Joules per packet); (2) *node density*: with higher node density, fewer nodes will overhear a packet transmission in SmartNode compared to 802.11; and (3) the difference between E_{re} (energy consumed when receiving) and E_{idle} (energy consumed when listening). Since SmartNode reduces the number of nodes overhearing a packet for each packet transmission, significant power saving can be achieved especially when the power consumption of receiving a packet is much higher than the power consumption of listening to the channel.

The clustered case

Compared with random deployment of a certain number of nodes, we believe that nodes are usually clustered in more realistic scenarios and they only communicate within its cluster (e.g. conference room). To evaluate the SmartNode performance in such situations, we consider a scenario where an ad hoc network contains two separate clusters. Each cluster is 50 meters in radius, and located 150 meters apart. Nodes are randomly placed within each of the clusters, and there are a total of 16 initiated CBR flows. With respect to the effective throughput and normalized energy consumption, the simulation results in the clustered case are illustrated in Figure 7 and 8, respectively. Figure

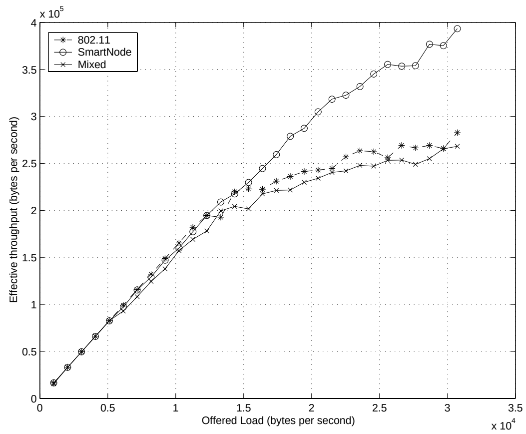

Figure 5. Effective throughput under varying offered load

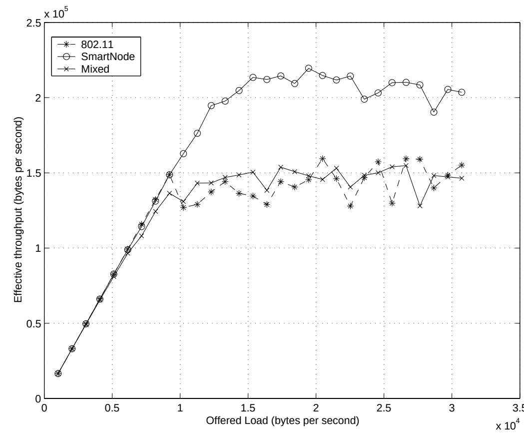

Figure 7. Effective throughput: the clustered case

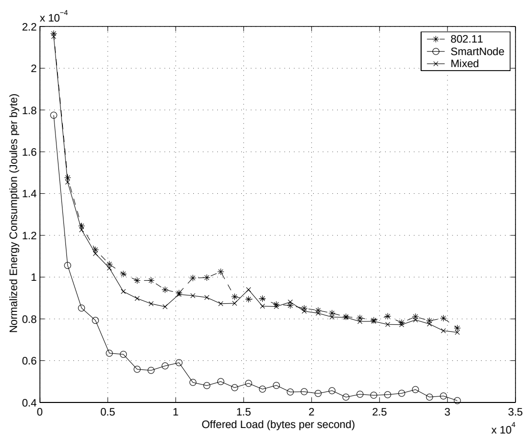

Figure 6. Normalized energy consumption under varying offered load

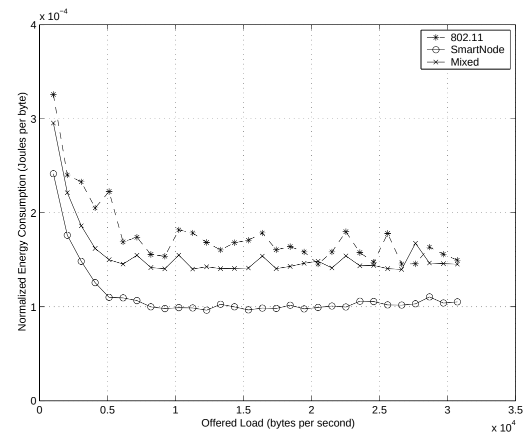

Figure 8. Normalized energy consumption: the clustered case

7 shows that SmartNode introduces an throughput improvement of approximately 60% compared with 802.11. This is due to the fact that SmartNode allows simultaneous flows in each of the clusters, while 802.11 only allows one at a time. The normalized energy consumption is also shown (Fig. 8) to be reduced significantly, due to lower power consumption rates at each of the nodes. In addition, the mixed network also performs slightly better than the 802.11 network, since it enjoys better spatial reuse of spectrum when those sender-receiver pairs that actually run SmartNode transmit at the same time. Based on these results, we conclude that, SmartNode enjoys significantly higher throughput and lower power consumption in a clustered network, which is more realistic.

Evaluation of fairness

In order to evaluate SmartNode with respect to fairness, we demonstrate how the bandwidth is distributed among 16 flows, with each of them offering a fixed load of 25.60 KB/s. Table 1 illustrates the total number of packets successfully transmitted for each flow under IEEE 802.11 and SmartNode with different setting of ω. In this table, we can see that IEEE 802.11 starves flows F5, F6, and F7. But for nodes running SmartNode algorithm, F6 and F7 are able to enjoy a better share of the bandwidth, due to lower interference from other flows when those flows only use the minimum required transmission power level for packet delivery. Overall, SmartNode can enjoy a higher throughput and a lower standard deviation (STDEV in Table 1) as well. A small standard deviation is preferred since it signals that there exists a small variance of throughput among flows. In other words, the system is able to distribute the bandwidth more fairly to each flow. Since maintaining strict fairness guar-

Table 1. Successful packet delivery comparisons with different settings of ω

Flow	802.11	SN(0%)	SN(4%)	SN(8%)	SN(12%)
0	3462	4782	4314	4782	3246
1	2428	5006	5050	5050	5045
2	2195	4984	5059	5052	5043
3	5025	5053	5049	5050	5052
4	3681	5060	5053	5056	5054
5	468	337	302	274	343
6	566	3832	3740	4340	3700
7	159	2872	4192	2233	2763
8	1913	4523	4695	4786	4646
9	4282	2975	2495	3883	3372
10	4221	4305	4210	4040	4853
11	5047	5053	5057	4613	5057
12	5047	5048	5050	5050	5050
13	5050	5050	5050	5051	5051
14	1427	5048	5049	5049	5026
15	5050	4611	5050	5050	5051
Total	50021	68539	69415	69359	68352
STDEV	1816	1276	1286	1311	1311

Table 2. Successful packet delivery comparisons with different settings of t

Flow	SN(0)	SN(2)	SN(4)	SN(6)
0	4576	4314	3179	4782
1	5010	5050	5044	5006
2	4492	5059	4589	4984
3	5053	5049	5050	5053
4	4142	5053	5056	5060
5	391	302	409	337
6	3454	3740	3892	3832
7	1174	4192	4077	2872
8	4224	4695	4771	4523
9	2694	2495	2859	2975
10	3405	4210	4831	4305
11	5051	5057	5054	5053
12	5045	5050	5050	5048
13	4602	5050	5050	5050
14	5050	5049	5048	5048
15	4610	5050	5050	4611
Total	62973	69415	69009	68539
STDEV	1417	1286	1259	1277

antees based on a certain fairness definition is beyond the scope of this paper, we do not claim that SmartNode may achieve perfect fairness in all scenarios.

In this scenario, we also attempt to investigate the effect of ω when it is equal to 0%, 4%, 8%, and 12%. In Table 1, the results have shown that as ω increases, the system throughput tends to decrease since it discourages spatial reuse of spectrum in most cases. With a small ω, some flows may have difficulty accessing the channel due to incomplete knowledge at each node (such as F5). A successful low power transmission of the RTS/CTS handshake may fail to reserve the channel. These flows may also have higher collision rates, which leads to a lower throughput as well. The standard deviation is very close for all cases using SmartNode, but the system throughput is higher when $\omega = 4\%$ and $\omega = 8\%$.

In addition, Table 2 shows the results of running the SmartNode with different settings of t. Clearly, the system throughput is low if $t = 0$, since it will use a higher power for packet transmission even at the first attempt. Such approach will discourage spatial reuse of spectrum, and may be even more unfair to some flows due to higher interference. On the other hand, if SmartNode adopts a large t value, higher packet collision may lead to lower throughput. From this scenario, we notice that there does not exist an optimal value of ω or t for the best performance, since the actual result depends heavily on the network topology and different situations of flow contention.

6 Concluding Remarks

We have proposed SmartNode, an extension of IEEE 802.11 MAC protocol to better support dynamic range adjustments in multi-hop ad hoc networks. The objectives are to be simple, scalable (low overhead), and offers backward compatibility with 802.11 (with a single wireless channel). With SmartNode, the benefits of dynamic range adjustments — power conservation and spatial reuse of spectrum — may be achieved compared with a standard 802.11-based network. Extensive results of performance evaluation in ns-2 demonstrate that a SmartNode-enabled network generally conserves more energy, enables better spatial bandwidth utilization, and provides better bandwidth allocation among flows, which verifies our intuitions and achieves our goals. These results also lead us to believe that a smooth transition is possible toward a power-aware MAC protocol standard.

References

[1] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, "MACAW: A media access protocol for wireless lan's," in *Proceedings of ACM SIGCOMM*, 1994.

[2] "IEEE 802.11 Standard," *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, 1999.

[3] J. P. Monks, V. Bharghavan, and W. W. Hwu, "A Power Controlled Multiple Access Protocol for Wireless Packet Networks," in *Proceedings of IEEE INFOCOM*, 2001.

[4] S. Singh and C.S. Raghavendra, "PAMAS - Power Aware Multi-Acces protocol with Signalling for Ad Hoc Networks," *ACM Computer Communication Review*, 1999.

[5] S. L. Wu, Y. C. Tseng, and J. P. Sheu, "Intelligent Medium Access for Mobile Ad Hoc Networks with Busy Tones and Power Control," *IEEE Journal on Selected Areas in Communications*, 2000.