

SHUSH: Reactive Transmit Power Control For Wireless MAC Protocols

Anmol Sheth and Richard Han
Department of Computer Science
University of Colorado, Boulder

Email: anmol.sheth@colorado.edu, rhan@cs.colorado.edu

Abstract—Asymmetric transmission ranges caused due to transmit power control have the undesirable effect of increasing the number of hidden terminals in the network as well as increasing the unfairness in channel access. In this paper we present a new reactive power controlled MAC protocol, SHUSH, which tackles the above problems. We compare the performance of SHUSH with four other transmit power controlled MAC protocols and demonstrate that SHUSH achieves superior aggregate goodput, spatial reuse, fairness, and minimal energy consumption.

I. INTRODUCTION

Transmit power control is a highly effective technique for minimizing interference and energy consumption in wireless networks. If two communicating nodes can lower their transmission powers and corresponding transmission radii to the minimal values that maintain communication between the two nodes, then the transmission between the two nodes will minimize interference with other nodes, enabling higher aggregate throughput and greater spatial reuse than wireless networks that lack power control [19]. Transmit power control has further value in terms of minimizing energy consumption due to transmission. Portable wireless access devices such as laptops, PDAs, and cell phones have extremely limited battery lifetimes. Energy constraints are also severe in wireless sensor networks. Transmit power control offers an effective technique for prolonging the lifetime of these wireless access devices by communicating at the minimum power that still maintains connectivity.

An assumption made by existing MAC layer protocols is that communication ranges are symmetric. However, in presence of power control this assumption is no longer valid. A critical issue that is encountered by MAC layer design given asymmetric transmit powers is the *general hidden terminal problem*, which can result in increased collisions and degraded throughput. We define the general hidden terminal problem as any situation in which an interferer is unable to hear a hidden ongoing conversation and then interrupts that hidden conversation. The *specific hidden terminal problem* addressed by MACA [2] is a subset of the general problem and assumes symmetric transmit powers. In a symmetric transmit power environment, the only way that an interferer is unable to hear an ongoing conversation and then interrupts that conversation is if the interferer cannot hear the sender and then interrupts the receiver by transmitting. In this case, the sender is hidden from the interferer, hence the classic hidden terminal terminology.

Asymmetric transmit powers introduce new and more general hidden terminals. For example, in Figure 1, nodes A and B are unable to hear the low power transmission of data from C to D. Nodes A and B have set their transmit powers to higher power in order to communicate. If node A transmits data to B, then the A-B conversation will unwittingly interfere with the C-D conversation. In this case, nodes C and D constitute *general hidden terminals* that have a hidden ongoing conversation that is interfered with by nodes A and B. Note that the interfering terminals are no longer confined to be within the carrier sense range of receiver D, as in the case of symmetric transmit powers. Instead, these new interfering terminals can be located anywhere outside of the union of C and D's carrier sense ranges (shown) yet within the union of carrier sense ranges formed at maximum transmit power from C and D (not shown), e.g. where A and B are located. Let us term this region the expanded interfering terminal zone. The interfering nodes collide with the hidden conversations, forcing retransmissions and degrading throughput. The increase in collisions caused by transmit power controlled hidden terminals was also observed in [1].

SHUSH provides the first complete solution that is both fair and efficient to the problem of general hidden terminals introduced by asymmetric transmit powers in wireless 802.11 LANs. This general solution should preserve many of the original motivations for asymmetric transmit power, namely minimal energy consumption and minimal interference to enhance throughput and spatial reuse.

Existing solutions for power control require sending RTS/CTS at maximum power. This leads to poor spatial reuse and poor aggregate throughput and is at odds with our original motivations for transmit power control. Consider the BASIC protocol, which transmits the RTS/CTS at maximum power to address hidden terminals and transmits the DATA/ACK at optimum (minimal) power to conserve energy. This BASIC protocol was proposed as PARO in [6], and a similar protocol was proposed in [7]. PARO was subsequently renamed as BASIC [1]. The problem with BASIC is that the RTS/CTS reserves a large floor space for the subsequent DATA/ACK conversation that is carried on in a much smaller floor space. In BASIC's zeal to suppress potential interferers, RTS/CTS max also suppresses other concurrent conversations that would not interfere with the low power DATA/ACK conversation.

The problem with large floor space reservations is exacerbated in 802.11 by the fact that the RTS/CTS frames have

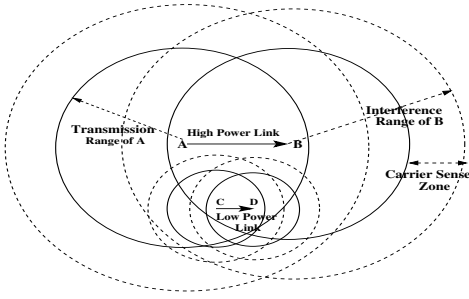


Fig. 1. Example of Low Power Link hidden from a High Power Link




Fig. 2. Overview of transmit power controlled MAC protocols

extended range compared to data packets transmitted at the same power, because RTS/CTS frames are sent at a lower “basic” data rate, which is 1Mbps for 802.11b.

An alternative approach taken by SHUSH is to be reactive rather than proactive in an asymmetric power control environment. In this approach, nodes such as C and D initially conduct their RTS-CTS-DATA-ACK exchange at optimum power. Only after there is interference do the nodes react by identifying the source of the interference via the header fields in the interfering packets. Nodes will then calculate the optimal transmit power needed to reach the interferer, and send only the first frame of the interrupted conversation at optimal power to SHUSH the interferer *a posteriori*. All subsequent frames are transmitted again at the power level sufficient to communicate with node D. The benefits of this reactive approach are that only the optimal (minimum) floor space is reserved, instead of the maximum floor space, and the interferer is identified and suppressed. As a result, spatial reuse is increased and more concurrent conversations can take place. As we will show, the aggregate throughput is considerably enhanced despite interruptions to ongoing conversations. Moreover, this reactive approach conserves energy by sending every transmission at the minimal transmit power necessary. By SHUSHing the interferer after the fact, SHUSH’s principle is to respond to concrete information about interference rather than to anticipate worst-case interference.

We summarize the behavior of several transmit power controlled MAC protocols, including plain 802.11, BASIC, and SHUSH in Figure 2. In plain 802.11 without power control, RTS, CTS, DATA and ACK are all sent at a default maximum transmit power P_{max} . Asymmetric transmit powers allows variation across these four packet types. BASIC sends RTS and CTS at P_{max} and sends DATA and ACK at optimum power P_{opt} . OPC refers to a minimum-energy strategy of optimal

power control (OPC), in which all four packet types are sent at minimum transmit power needed to maintain connectivity. OPC was depicted in Figure 1 as being susceptible to the general hidden terminal problem.

A power controlled MAC (PCM) protocol was proposed for 802.11 [1]. PCM provides a mechanism to avoid collisions from nodes in the carrier sense zone. In PCM, RTS/CTS packets are transmitted at maximum power and data is periodically also transmitted at maximum power; the data is otherwise sent at optimum power to conserve energy. The periodic maximum power data transmissions enable nodes in the sender’s carrier sense zone to stay suppressed throughout the sender’s transmission. A drawback, as noted earlier, is that the RTS/CTS transmission at maximum power suffers from poor spatial reuse and degraded throughput in transmit power controlled environments.

A second critical issue encountered by MAC layer design given asymmetric transmit powers is *fairness* of medium access. There are two mechanisms of unfairness introduced by transmit power control. First, because of the general hidden terminal effect, high power conversations can inherently and arbitrarily interrupt hidden low power conversations, forcing retransmissions and unfairly degraded throughput for low power conversations. Second, this inherent unfairness favoring high power conversations over low power conversations is exacerbated by the backoff algorithm. After a collision, the backoff algorithm manages how long nodes delay before retransmitting, typically increasing the backoff interval after repeated collisions. For example, the IEEE 802.11 protocol employs the Binary Exponential Backoff (BEB) algorithm [2]. The unfairness caused by BEB in 802.11 is a well-known problem [5]. BEB favors the last node that was successful in transmission. This exacerbates the fairness problem in heavy traffic as the unsuccessful nodes do not gain access to the channel. In our case, the unsuccessful nodes are repeatedly the nodes transmitting at low power.

To address the unfairness issue, SHUSH employs a simple stateful strategy that is grafted into the reactive strategy of SHUSHing the interferer. SHUSH employs the principle that a node N that has been interrupted has a higher priority to access the medium than typical nodes, since node N had been communicating beforehand. Let us term this principle the *interruption principle*. Therefore, after an interruption of an ongoing conversation, the interrupted node should be able to access the medium *sooner* than other nodes. SHUSH employs the additional principle that the interrupted node waits until after the interrupter has finished before grabbing the channel sooner than other nodes. This second principle avoids immediately interrupting the interrupter, on the theory that continued interruptions will cause a domino effect leading to decreased throughput. Let us term this second principle the *patience principle*. A specific instantiation of these principles in 802.11 would be to allow an interrupted node to grab the channel during the DIFS interval immediately following completion of the interferer’s conversation. We demonstrate that SHUSH’s reactive strategy combined with SHUSH’s two principles of interruption and patience achieve fairness while maintaining high throughput in the presence of asymmetric

transmit powers.

In the rest of the paper, Section II explains the SHUSH MAC protocol in detail. Section III provides an in-depth analysis and evaluation of the goodput, fairness, spatial reuse, and energy consumption of SHUSH compared to 802.11, BASIC, PCM, and OPC transmit power controlled MAC protocols. Section IV provides the background related work and Section V concludes the discussion of SHUSH.

II. SHUSH PROTOCOL

As discussed in Section 1, transmit power control increases the number of hidden terminals in the network and exacerbates the unfairness of medium access for the low power links in the network. SHUSH provides the first comprehensive solution to the problem of general hidden terminals which is both fair as well as efficient. The main goals of the SHUSH protocol are to provide a unified solution that addresses hidden terminal effects while also achieving the following:

- Enhanced spatial reuse and improved aggregate throughput throughout the wireless network
- Fair sharing of the channel between low powered and high powered links
- Achieving the above without enforcing RTS/CTS based collision avoidance
- Minimizing overhead of the protocol in the absence of transmit power control and hidden nodes
- Reducing power consumption
- Minimal modifications to the standard

SHUSH maintains a conservative philosophy with respect to transmission power in an environment with asymmetric transmit powers. All conversations begin at the optimal transmit power, with the intent of minimizing interference and maximizing spatial reuse and aggregate throughput. The optimal power is calculated using the any of the protocols specified in [15] [8]. These protocols use either in-band(modify existing MAC headers) or out-of-band(special control packets) mechanisms to inform the transmitter of the signal strength at the receiver. Thus SHUSH is agnostic to the exact implementation of calculating optimal transmit powers and can leverage any of the existing techniques to achieve its power control goals.

In the event that an ongoing low power conversation is interrupted by a high power conversation that is unable to hear the hidden conversation, then SHUSH initiates a protocol for the interrupted node(s) to contact and SHUSH the interferer(s). By notifying the high power nodes of the interruption, the interrupted nodes gain access to the channel and reserve the channel for the duration of the incomplete communication. This notification can either take the form of a control message (RTS/CTS) or a DATA fragment packet, as discussed below.

Although SHUSH is not confined to the IEEE 802.11b standard, we evaluate the protocol and compare the protocol against already proposed transmit power control protocols based on 802.11b. The IEEE 802.11b standard defines 3 different inter-frame time intervals, *SIFS*, *DIFS* and *EIFS*. The *SIFS*, which is 10 μ sec, is the shortest of the inter-frame spaces and is used for a station to retain access of the channel once acquired. The *DIFS*, which is 50 μ sec, is used by stations

to transmit data/management frames. A station is allowed to transmit a frame only if it senses the medium to be free at the end of the *DIFS* interval and its backoff timer has expired. The *EIFS*, which is 364 μ sec, is used by stations which sense the carrier to be busy but cannot decode the frame. Hence nodes in the carrier sense zone (Figure 1) need to defer for *EIFS* rather than *DIFS* before transmitting the frame.

The IEEE 802.11 protocol specification defines a fragmentation mechanism in which network level packets greater than a pre-defined *fragmentation threshold* are partitioned into smaller MAC frames. These fragmented frames are transmitted in a *fragment burst* i.e. once the node has gained access to the channel it shall continue to send fragments until all fragments of a network level packet have been transmitted or it does not receive an acknowledgment.

In the presence of hidden terminals in the network, fragmentation assists in faster recovery as well as lower power consumption of the nodes that experience interference/collisions. However fragmentation adds some overhead as each data fragment frame needs to be separately acknowledged.

In order for SHUSH to operate correctly, SHUSH must be able to determine:

- the source address of the interferer whose transmission caused the collisions/interference
- the optimum power with which to SHUSH the interferer
- when the interrupting conversation has completed
- which interrupted node gets to SHUSH which interferer if there is more than one of each, and when

The following sub-sections explain how SHUSH determines the above elements by making use of the inter-frame spacing and fragmentation mechanism of the IEEE 802.11 protocol.

A. Obtaining information about the interferer

After experiencing interference, at least one of the interrupted nodes will be able to hear at least one of the interferer(s). This interrupted node must be able to acquire the address of the interferer, the optimum power needed to reach the interferer, and also the time when the interferer has completed its conversation.

A variety of interference scenarios may occur. Consider the topologies shown in Figure 3, which shows only three of the possible fifteen interference scenarios. Here nodes A and B are the high power sender (HPS) and high power receiver (HPR) respectively. Nodes C, C', C'', are the low power senders (LPS) and D, D', and D'' are the low power receivers (LPR) that are hidden from the nodes HPR and HPS. In the rest of this section we use the above terminology of HPR, HPS, LPR and LPS. C and D are within the transmission range of both nodes A and B. Nodes C' and D'' are within the transmission range of only node B. Nodes D' and C'' are outside the transmission range of both nodes A and B.

For reasons described in the next subsection, SHUSH needs to obtain the MAC address of the interferer. The interferer may be either the HPS or the HPR, or in some cases both. If the interferer is HPS, then every data frame header in the IEEE 802.11 MAC protocol would have the source address encoded in the frame. If the interferer is HPR, then ideally we would




Fig. 3. Example of three different topologies of the high power nodes A and B and low power nodes C, C', C'', D, D' and D''.

like the address of the HPR to be encoded in the ACK frames. However, 802.11 only includes the destination address in ACK frames, i.e. only the HPS' address is encoded in the ACK frames transmitted by HPR. While SHUSH is more general than 802.11, the specific implementation of 802.11 SHUSH in this paper uses the HPS' address in place of the HPR's address i.e. it uses the address of the node who transmitted the ACK frame.

Since the MAC address information is obtained from the MAC data and control frame headers, one caveat is that this approach depends upon there being at least two MAC frame headers transmitted between the HPR and HPS i.e. there must be at least one more frame header following the frame that caused the interference in order for the low power interrupted nodes to receive the MAC address information. In the case where the data frames are not fragmented, or the data is too small for fragmentation, then an extra data frame with zero payload (just the MAC header frame) is transmitted which is also ACKed by the HPR. Thus the default DATA-ACK handshake is modified to DATA-ACK-DATA_HEADER-ACK. This modified DATA-ACK sequence enables the interrupted low power node to gain all the information required to SHUSH the high power nodes. A 802.11b data frame header is 28 bytes long and the ACK frame is 14 bytes long and these frames unlike the RTS/CTS control frames are transmitted at the same rate as the data frames. The trailer header mechanism adds a very small overhead to the default DATA-ACK handshake. The semantics of transmitting the extra trailer header with its acknowledgment is similar to having fragmentation enabled, except that the data payload of the trailer fragment is set to 0. The 802.11 protocol already supports fragmentation, so we expect that SHUSH's trailer header mechanism requires minimal modifications to the existing protocol, while accruing significant benefits in tackling the aggravated hidden terminal problem, as shown in Section III.

In the case where fragmentation is enabled, then the extra trailer header is not required as the data fragment train and the corresponding ACKs provide the necessary information for the low power nodes to SHUSH the high power nodes.

SHUSH depends upon the MAC layer packets to have embedded information concerning the transmit power of each packet. Typically, the MAC layer header is modified to include field(s) required for transmit power control like RSSI and TxPower. Several transmit power control algorithms have proposed modifying the 802.11 header to incorporate such power control information [8], [11]. Given the transmit power of the interferer and the RSSI at the interrupted node, the

interrupted node can calculate the optimal transmit power required to SHUSH the interferer. SHUSH is agnostic to the specific implementation of transmit power control and can leverage off any of these protocols to achieve transmit power control.

The next important issue is to determine how the interrupted nodes know when the fragment burst between the high power interfering nodes has completed, i.e. exactly when the ACK frame for the last data fragment would be transmitted. In the case that HPS is the interferer, the 802.11 data fragments would have embedded within their headers information about the duration of the transmission. When HPR is the interferer, the 802.11 ACKs will also have the duration information embedded within their headers. Thus, the lower nodes will be able to determine when the interference from HPS/HPR will cease.

B. SHUSHing the interferer

The next task is to SHUSH the interferer and then resume the interrupted low power conversation. Normally, an 802.11 node that has been interrupted by interference will backoff using the BEB algorithm. As observed earlier, this leads to unfair medium access, especially for low power conversations. In SHUSH, the interrupted node does not exponentially back off, and instead waits until the HPS/HPR conversation has completed. At this point, the interrupted node is poised to SHUSH the interferer and gain control over the medium.

SHUSH takes advantage of the fact that there is a DIFS waiting interval enforced after the completed conversation i.e. after the ACK for the last data fragment is received, in which HPS and HPR are idle. After a successful transmission, the high power nodes defer for an interval of $DIFS+RBO$ where RBO is the random backoff. The RBO is calculated as:

$$RBO = Random() * aSlotTime \quad (1)$$

where $Random()$ is a pseudorandom integer drawn from a uniform distribution over the interval $[0, CW]$ and $aSlotTime$ is $20 \mu sec$ for the DSSS PHY specification of 802.11b. To regain access to the channel the low power node re-initiates the interrupted conversation during the DIFS interval, before the high power node starts the random backoff (RBO). On a successful transmission node A (HPS) would set its CW to CW_{min} which is defined to be 31 for the DSSS PHY. Hence the RBO generated ranges from 0 to $620 \mu sec$. To gain a fair share of the wireless medium, the low power nodes (node C) need to SHUSH the high power source nodes (node A) before the high power node starts the RBO, as the RBO could also be set to $0 \mu sec$. SHUSH adheres to the interruption principle, which gives priority to a node that has been interrupted by a high power conversation, i.e. it is only fair that the interrupted conversation be given first priority to resume.

SHUSH divides the DIFS interval into 2 parts $aSlotTime$ and $aSlotTime + SIFS$ as shown in figure 4. This division orders the SHUSH responses based on whether the LPS or LPR was interfered with. For example, in Figure 3, node C corresponds to LPS and node D corresponds to LPR. If LPS has been interrupted, it has all the information necessary to




Fig. 4. The synchronization during the DIFS time interval

SHUSH the interferer. LPS will then initiate a SHUSH within the first half of the DIFS interval. If instead only the LPR is interfered with, then the LPR does not know whether the LPS has been interfered with. For this reason the LPR waits during the first half of the DIFS to see if the LPS initiates a SHUSH. If no such SHUSH is detected, then the LPR initiates a SHUSH in the second half of the DIFS.

A random backoff from the interval of $[0-20]\mu s$ is selected before transmitting the SHUSH. This ensures that multiple low power nodes that have been interrupted do not transmit the SHUSH frame simultaneously. Without the random backoff, the SHUSH frames would result in a collision at the high power nodes as every low power interrupted node would transmit the SHUSH at the same instant. It is also possible for multiple low power nodes that are interrupted to be hidden from each other. In this case, the SHUSH frames transmitted by the LPR/LPS would collide at the HPS which would lead to the high power node to set its NAV to EIFS and backoff. Thus in this case the channel reservation of the low power nodes would not be communicated to the high power nodes.

The last issue concerns what kind of SHUSH message gets sent to the interferer. An important design constraint of SHUSH was to make minimal modifications to the existing MAC protocol. Rather than design a new control frame, we make use of the existing RTS, CTS, and DATA frames to signal a SHUSH. If LPS has been interrupted, then the message sent by LPS will depend on whether RTS/CTS is enabled. If RTS/CTS is enabled, then the SHUSH frame consists of an RTS sent at the optimum power. If RTS/CTS is not enabled, then the SHUSH frame consists of the data frame that was interfered with. Thus unlike BASIC and PCM, SHUSH does not enforce the use of the RTS/CTS based collision avoidance. If only LPR was interrupted, then the SHUSH frame consists of an unsolicited CTS sent at optimum power. All SHUSH frames contain a duration field that indicates how long to backoff. This will cause HPR/HPS to set their NAV fields and back off. The LPS on receiving this unsolicited CTS frame resumes the interrupted conversation at optimal power after a SIFS interval.

III. SHUSH PROTOCOL EVALUATION

To evaluate the SHUSH protocol, we compare the performance of the protocol with 4 other transmit power control 802.11 based MAC protocols. We implement and compare SHUSH with Optimal Power Control (OPC), BASIC [6] and Power Controlled MAC (PCM) [1]. We also compare the performance of these protocols with 802.11b (802.11) without any transmit power control. Figure 2 provides an overview of the behavior of the 5 different transmit power protocols that we evaluated.

A. ns2 Simulation Setup

We implement the above mentioned protocols in the ns2 simulator (ns-2.26) with the CMU wireless extensions [16]. The channel bit rate was set at 2Mbps for the data frames and control frames were transmitted at the basic rate of 1 Mbps. Packet size was set to 800 bytes with the fragmentation threshold set at 400 bytes. Hence each network level packet was fragmented into 2 data fragments. With fragmentation enabled we evaluated the trace files generated and observe an improvement of around 6% in the goodput as compared to the trailer header mechanism. However, the fairness and capacity of the network are the same for both the techniques. We use a CBR traffic source for each pair of nodes (flow) in the network. The 2-Ray ground reflection model was used. We do not consider mobility in our simulations. We evaluate the protocols for a random topology over a flat grid of $500m^2$. The random topology was generated by setting up CBR traffic between any 2 random nodes in the network. The nodes were selected such that they are within transmission range of each other. The receive threshold for 1Mbps, 2Mbps, 5.5Mbps and 11Mbps were set at -92dB, -90dBm, -85dBm and -80dBm, respectively. As part of future work we would like to evaluate the protocol under different data traffic characteristics, like bursty web traffic, TCP flows, etc

The maximum distance between adjacent nodes for the random topology was set at 200m. This is the farthest that a node could communicate at the maximum transmit power. We used 10 different power levels similar to those used in [1]. At maximum transmit power a node can communicate with another node at a maximum distance of 200m.

B. Simulation Results

1) *Goodput*: We first compare the five protocols by measuring the aggregate goodput of the network. Goodput is a measure of the amount of useful application layer data received by the node, and hence ignoring all the MAC layer and network layer control signaling. We generated 10 different random topologies and for each topology we divide the total goodput of the network by the number of active flows in the network ($Kbits/sec/node$). We measure the aggregate goodput of the network for each of the five protocols. We also varied the number of nodes in the network from 20 to 80, i.e. varied the number of active flows in the network from 10 to 40.

Figure 5 shows the aggregate goodput of the network with RTS/CTS signaling enabled. OPC outperforms all the other protocols. This is because the high power transmissions experience no interference from the smaller low power transmissions. Thus the high power nodes always have their backoff window at CW_{min} and dominate the medium by saturating the link between them. On the other hand the low power nodes are constantly interrupted and their backoff window keeps on increasing exponentially. However, as the number of flows increase in the network, the aggregate goodput of OPC drops and approaches that of the other protocols. This is because as the number of flows increase, the number of high power transmissions that interfere with each other also increase. We

also observe that SHUSH does better than 802.11, BASIC and PCM. This is because 802.11, BASIC and PCM all transmit the RTS/CTS at maximum power and hence result in a lower spatial reuse. SHUSH transmits the RTS/CTS at optimal power and only on a collision does it require to step up the power as described in II.

Figure 6 shows the goodput of 802.11, SHUSH and OPC without RTS/CTS negotiation. Overall we observe that the goodput without RTS/CTS is higher than that of with RTS/CTS due to the absence of the extra control frames. PCM and BASIC enforce the use of RTS/CTS based collision avoidance and hence cannot be compared. We observe similar results here and also an improvement in the aggregate goodput of all the 3 protocols in the absence of RTS/CTS.




Fig. 5. Aggregate Goodput with RTS/CTS signaling (Error bars not shown)

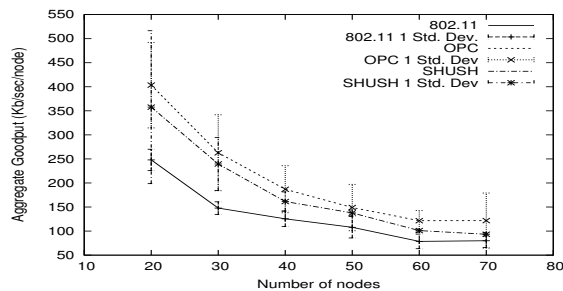


Fig. 6. Aggregate Goodput without RTS/CTS signaling with error bars of 1 std. dev.

2) *Fairness*: One of the objectives of the SHUSH protocol is to provide a fair share of the medium to low power and high power nodes in the network. To evaluate fairness, we measure the standard deviation of the goodput of the individual flows in the network, averaged over 10 different random topologies. Thus a high standard deviation implies that the protocol is unfair by providing unequal share of the medium across the individual flows. A low standard deviation implies that all flows in the network are provided nearly equal share of the medium and hence the protocol is more fair. Figure 7 shows the aggregate standard deviation with the RTS/CTS signaling for all five of the MAC protocols for increasing number of flows in the network. Clearly, OPC is the most unfair protocol in sharing the channel. This is because the high power transmissions dominate the channel, starving low power conversations. PCM, BASIC and SHUSH show comparable fairness. Figure 8 shows similar results of the aggregate standard deviation without RTS/CTS signaling

with OPC again being the most unfair and SHUSH performing better than 802.11 as well as OPC.

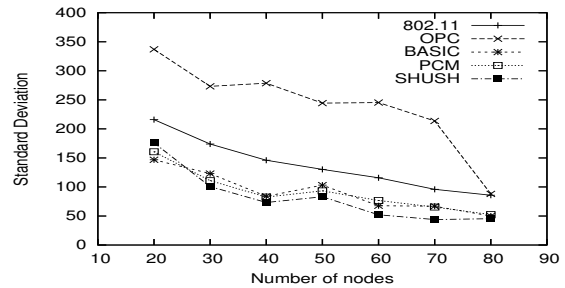


Fig. 7. Std. Dev. in goodput of individual flows with RTS/CTS signaling (Error bars not shown)

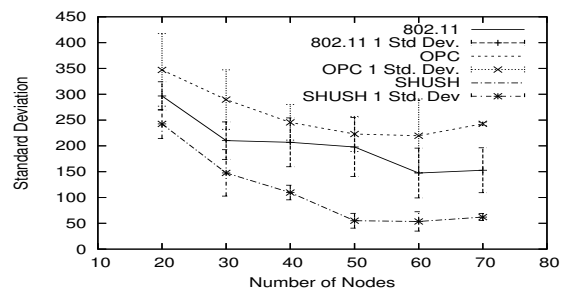


Fig. 8. Std. Dev. in goodput of individual flows without RTS/CTS signaling

As observed from the results above, even though OPC provides a higher aggregate goodput, it is the most unfair of all the protocols. Figure 9 shows the OPC goodput of the individual flows in a 20 node network (10 flows). The height of the bars demonstrate the goodput of the individual flows and the line graph at the bottom plots the distance between the source and destination node for every flow. Figure 10 shows the topology of this network. Clearly, flow 8 is the most powerful transmission and dominates flows around it i.e. flows 3,6,7,9. Similarly, flow 2 dominates flow 1 and 10. This explains the high goodput obtained by flow 2 and 8. These high power transmissions constantly force the less powerful nodes to backoff as explained in section II and hence the goodput of flows 3,6,7,9,1 and 10 is very low. Even though flow 5 is a low power transmission, it achieves a high goodput since it is isolated from all the other flows in the network. Hence, flow 5 does not experience any interference/collisions from the other flows. This clearly demonstrates the drawback of OPC in which hidden terminals are forced to backoff and are continuously interfered with when transmitting.

Figure 11 gives an estimate of the spatial reuse of the network obtained by SHUSH, OPC and 802.11. The spatial reuse also reflects the fairness of the MAC protocols. The spatial reuse is calculated as the number of active transmissions per unit time interval of the trace. The time interval was set to $0.5sec$ and for each time interval we measure the number of unique nodes that receive an application layer packet. The graph plots the spatial reuse of a 70 node network. The x-axis plots the time intervals and the y-axis measures the pairs of nodes that receive an application layer packet in that




Fig. 9. Goodput of a 20 Node random topology network using OPC




Fig. 10. Random topology of the 20 Node network

time interval. Clearly, SHUSH outperforms 802.11 due to the *patience principle* employed by SHUSH. Hence, unlike OPC which does not provide any chance for the low power nodes to transmit, SHUSH provides a much improved spatial reuse of the network. BASIC and PCM achieve similar spatial reuse as that obtained by 802.11.

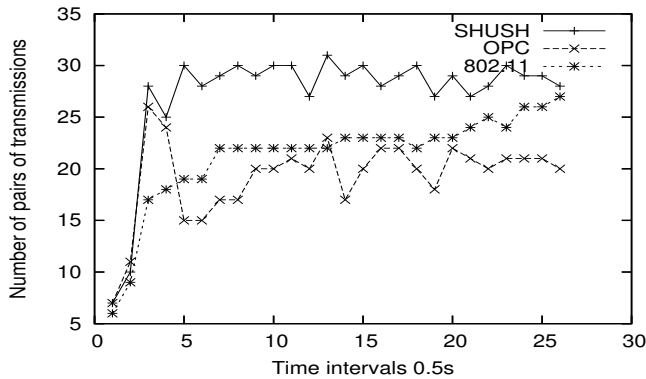


Fig. 11. Spatial reuse of the network (70 node random topology with 35 pairs)

3) *Energy Consumption*: We evaluate the energy consumption of the 5 different protocols based on the metric joules/byte transmitted. Figure 12 shows the energy consumption for the random topology. 802.11, which lacks power control, consumes the maximum energy and is constant across increasing nodes in the network. OPC, which sends every frame at the optimal transmit power has the minimum power consumption. Our results confirm with the results obtained in [1] where BASIC consumes more power as compared to PCM. However it is interesting to note that the power consumption for both BASIC and PCM remains constant even with increasing num-

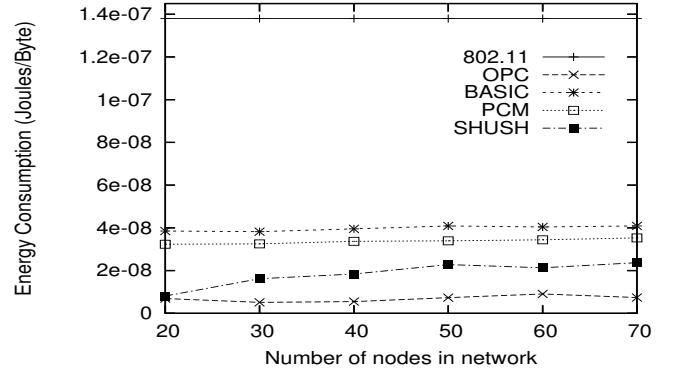


Fig. 12. Energy consumption for random topology

ber of nodes in the network. This is mainly because BASIC and PCM are not reactive to the collisions and interference caused by the high power nodes. BASIC and PCM always transmit the RTS/CTS at maximum power and PCM always steps up its transmit power periodically irrespective of the number of hidden terminals in the network. SHUSH on the other hand is reactive to the number of nodes in the network. When the number of flows in the network are small, SHUSH nearly performs as well as OPC. With increasing number of hidden terminals in the network, SHUSH needs to step up the transmit power and inform the high power nodes of the interference caused. Hence, the power consumption increases with increasing number of hidden terminals and thus SHUSH is more reactive as compared to PCM and BASIC. SHUSH consumes more power as compared to OPC, but the increase in power consumption results in improved fairness as compared to OPC.

IV. RELATED WORK

There has been little analysis in the literature that studies the effect of asymmetric transmission ranges on the MAC and network level. [19] derives an asymptotic expression and proves that a variable-range transmission outperforms a common-range transmission approach in terms of power savings and increased capacity of the network. The COMPOW [9] protocol is based on the common-range transmission approach. This approach has the drawback that a single distant node could cause the entire network to transmit at a higher power level.

Busy tone multiple access BTMA [13] proposes to combat hidden terminals in CSMA. BTMA requires a data channel and a control channel. The base station transmits a busy tone signal on the control channel as long as it senses a carrier on the data channel. [14], [11] introduce interference-limited MAC protocols, where nodes are allowed to transmit as long as their transmission does not corrupt the ongoing transmission. This is achieved by broadcasting a busy tone, which neighboring nodes use to calculate the interference-margin that can be tolerated by the on-going transmission. Though the above protocols do not tradeoff spatial reuse, the solution requires multiple channels and transceivers. [17] is a single channel solution which embeds the information about the interference

margin into the CTS frame. This information is used to bound the transmission power of potentially interfering nodes.

A variety of algorithms have been proposed to overcome the unfairness of the IEEE 802.11 protocol [4], [3], [10]. One of the seminal solutions was the Multiplicative Increase Linear Decrease (MILD) algorithm, proposed in the MACAW protocol [3]. The essence of the work in these algorithms is to achieve global state in the ad-hoc network by exchanging local information about fairness between neighboring nodes. These methods suffer from the overhead required to communicate fairness state.

MiSer [18] proposes a per frame transmit power control and rate control algorithm. It pre-computes the optimal {rate, power} table for a given {data payload, path loss, frame retry count} and the lookup is performed at runtime. However, this solution is restricted to 802.11a and 802.11h protocols and also requires the RTS/CTS based collision avoidance mechanism.

V. CONCLUSIONS

SHUSH addresses two critical issues that are introduced by asymmetric transmit powers: the general hidden terminal problem; and unfairness of medium access. Unlike existing MAC protocols, SHUSH adopts a reactive mechanism to deal with hidden terminals. Nodes engage in conversations with neighbors at the optimum. Only after an interruption does a node that has been interfered with transmit at a temporarily elevated power to SHUSH the interferer. This efficiently reserves the minimum floor space necessary. Fairness is achieved by giving priority to the interrupted low power nodes in the network over the interfering high power nodes. We compared SHUSH against four other transmit power controlled MAC protocols, and demonstrated that the resulting SHUSH protocol achieves superior aggregate throughput, spatial reuse, fairness, and minimal energy consumption in almost all cases.

REFERENCES

- [1] E.-S. Jung and N. H. Vaidya, "A power control MAC protocol for ad-hoc networks," in ACM MOBICOM, 2002.
- [2] P. Karn, "MACA, a New Channel Access Method for Packet Radio", ARRL/CRRL Amateur Radio 9th Computer Networking Conference, 1990.
- [3] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, "MACAW: A Media Access Protocol for Wireless LAN's," ACM SIGCOMM, 1994.
- [4] Z. J. Haas and J. Deng, "On Optimizing the Backoff Interval for Random Access Schemes," IEEE Transactions on Communications, vol. 51, no. 12, pp. 2081-2090, December 2003.
- [5] C. E. Koksal, H. I. Kassab, H. Balakrishnan, "An Analysis of Short-Term Fairness in Wireless Media Access Protocols," Extended version of short paper to appear in Proc. ACM SIGMETRICS, June 2000.
- [6] J. Gomez, A. T. Campbell, M. Naghshineh, and C. Bisdikian, "Conserving transmission power in wireless ad hoc networks," in Proc. of IEEE Conference on Network Protocols (ICNP'01), Nov. 2001.
- [7] S. Agarwal, S. Krishnamurthy, R. H. Katz, and S. K. Dao. "Distributed Power Control in Ad-hoc Wireless Networks" In PIMRC01, 2001.
- [8] J. Ebert and A. Wolisz, "Combined tuning of rf power and medium access control for wlans," Mobile Networks and Applications (Monet), vol. 1, pp. 1-11, July 2000.
- [9] S. Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar, "Power Control in Ad-Hoc Networks: Theory, Architecture, Algorithm and Implementation of the COMPOW Protocol", In Proceedings of the European Wireless Conference pages 156-162, Florence, Italy, Feb. 25-28 2002.
- [10] Zuyuan Fang, Brahim Bensaou, Yu Wang "Performance evaluation of a fair backoff algorithm for IEEE 802.11 DFWMAC". MobiHoc 2002: 48-57
- [11] J. Monks, V. Bharghavan, and W. Hwu, "A power controlled multiple access protocol for wireless packet networks," in Proceedings of IEEE Conference on Computer Communications (INFOCOM), Apr. 2001, vol. 1, pp. 1-11.
- [12] A. Muqattash and M. Krunz, "Power Controlled DualChannel (PCDC) Medium Access Protocol for Wireless AdHoc Networks", INFOCOM 2003
- [13] C. Wu and V. O. K. Li, "Receiver-initiated busy-tone multiple access in packet radio networks", in Proc. ACM SIGCOMM 87, 1987, pp. 336-342.
- [14] 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 Area in Communications, 18(9):1647-57, 2000.
- [15] A. Sheth, R. Han "Adaptive Power Control and Selective Radio Activation for Low-Power Infrastructure-Mode 802.11 LANs." ICDCS Workshops 2003 Pages 812-817
- [16] The CMU Monarch Project. The CMU Monarch Project's Wireless and Mobility extensions to NS.
- [17] A. Muqattash and M. Krunz, "A single-channel solution for transmission power control in wireless ad hoc networks" in Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing MobiHoc 2004, 210-221, 2004.
- [18] D. Qiao, S. Choi, A. Jain, and Kang G. Shin, "MiSer: An Optimal Low-Energy Transmission Strategy for IEEE 802.11 a/h", in ACM MobiCom'2003, San Diego, CA, September 14-19, 2003.
- [19] J. Gomez and A. Campbell, "A Case for Variable-Range Transmission Power Control in Wireless Multihop Networks", In Proceedings of the IEEE Conference on Computer Communications (INFOCOM), HongKong, March 2004