Persistent Pseudo-Clearance Problem in IEEE802.11 Mesh Networks and its Multicast Based Solutions

Jian Zhang [1] Yuanzhu Peter Chen [2] Ivan Marsic[1]
[1] Department of Electrical and Computer Engineering, Rutgers University
[2] Department of Computer Science Memorial University of Newfoundland

***Abstract*— Wireless mesh networks are flexible solutions to extend services from wireless LANs. The current IEEE 802.11 Specification, however, needs to be modified in various ways to be a suitable technology for this purpose. In particular, in order to handle the well-known Hidden Node Problem (HNP), the Specification adopts MACAW by employing an RTS/CTS/DATA/ACK 4-way handshake. Some flaws of this scheme have been noticed, e.g. the Masked Node Problem (MNP). In this work, we identify a critical problem of the Specification's 4-way handshake, called persistent pseudo-clearance (PPC). PPC occurs when for two sender/receiver pairs a CTS from one pair's receiver collides with the DATA frames of the other pair. This logjam can persist for a period of time despite of the random backoff the senders employ. The persistent frame losses in PPC can cause more serious problems. The effect of giving up a frame transfer after reaching the maximum number of retries can propagate to upper layers, causing routing errors or TCP sender backoff. Multicast RTS (or MRTS) provides a good solution framework to break the cycle of losses and retransmissions between such peers. With minimal modification to MRTS, we provide an effective and efficient solution to PPC. Our experiments show that MRTS breaks the logjam of PPC while fully utilizing the network capacity.**

***Index Terms*— PPC, 802.11 MAC, collision avoidance, mesh network**

I. Introduction

Compared to single-hop access-point-based networks, i.e., wireless local area networks (wireless LANs), wireless mesh networks present greater complexity. Their features of arbitrary and diverse topologies and multi-hop connections lead to more complicated medium access and thus mutual interference. Therefore, wireless mesh networks need more sophisticated medium access control and interference management mechanism than conventional wireless LAN. The most dominant technology to realize wireless mesh networks is the IEEE 802.11 [10]. Its ubiquity has made it the first choice for this new generation of wireless access networks although it was originally designed for wireless LANs. Indeed, when directly employed in a wireless mesh network, IEEE 802.11 MAC does not utilize the full network capacity. Thus, improvements have been proposed to meet the needs of wireless mesh networks.

In this paper, we identify a problem of the RTS/CTS/DATA/ACK 4-way handshake of 802.11 MAC, called *persistent pseudo-clearance* (PPC), which causes a large number of data frame losses and retransmissions. Consequently, the link layer error may propagate to upper layers of the network protocol stack, resulting in routing errors or TCP sending agent back-off, and thus triggering further performance degradation.

Essentially, PPC occurs when two sending/receiving pairs interfere with each other so that repetitive frame losses at the MAC layer are experienced. We notice that, despite the random backoff behavior of the senders regulated by the 802.11 MAC, such a cycle of losses can well last long enough to exhaust the maximum number of retransmits. If a mesh Internet gateway [3] is one of the senders involved in PPC, its overall capacity can be significantly impaired since delivering PPC-related frames costs excessive resources and may block backlogged frames of other flows from being serviced. To fully utilize the relaying capacity of mesh gateways, we explore using a multicast-based RTS, i.e., MRTS [16][18][11]. MRTS is a powerful extension to the existing 802.11 MAC, originally proposed to transmit frames opportunistically and to overcome the head-of-line (HOL) blocking problem. Here, we show that MRTS can effectively cope with PPC. Further, we cycle-shift the receiver addresses contained in the MRTS frame to poll other available receivers alternatively. This solution is thus called *Shift-MRTS*, or SMRTS.

In the rest of this paper, we first describe PPC and the conditions for it to occur in Section II. We then review related work that attempts to solve problems similar to PPC but from the aspect of network performance improvement in Section III. In Section IV, we provide our solutions to PPC, MRTS and SMRTS, where we elaborate the rationale of extending RTS with multicasting capabilities. To verify the effectiveness of our methods, we design experiments and show the results. Potential extensions of this work are described in Section V.

Proceedings of the 2007 15th IEEE Workshop on Local and Metropolitan Area Networks
1-4244-1100-9/07/$25.00 ©2007 IEEE

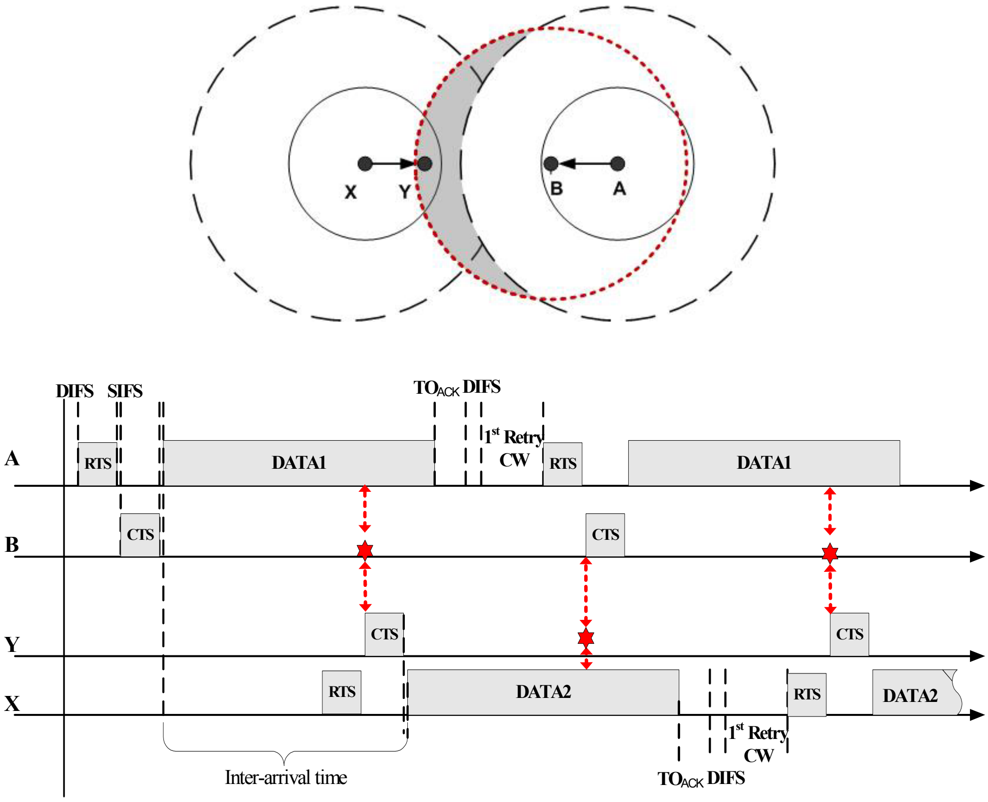

Figure 1. (a) Layout of PPC scenario; (b) Timeline of a cycle of retransmissions

II. Persistent Pseudo-Clearance Problem

Wireless mesh networks are operated in much more complex environments. Many problems have been addressed to provide link layer reliability of such networks. Among these, probably the best known is the hidden node problem (HNP). An effective solution to this problem is MACAW [5], adopted by the IEEE 802.11. It is essentially a 4-way handshake consisting of RTS/CTS/DATA/ACK. Unfortunately this scheme does not solve the problems for all possible scenarios. For example, the masked node problem (MNP) [15] was observed, where it is incorrect to assume that all nodes within the sender's transmission range can hear the RTS and all nodes within the receiver's transmission range can hear the CTS, even under perfect channel conditions. However, MNP in [15] is described under the assumption that the carrier sensing range and the interference range are the same as the transmission range, which is not particularly realistic [17]. With the capture effect taken into account and capture threshold set to, say, 10dB, the interference range is 1.78 times of the distance between the sender and the receiver. For example, in Figure 1(a), the interference range for transmission *AB* is shown as a red dotted circle centered at the receiver B. It is observed in [17][12] that this disparity between interference range and RTS/CTS transmission range may cause the failure of collision avoidance for the 802.11 4-way handshake scheme. In this paper, we identify PPC as another situation where the 4-way handshake malfunctions persistently. In fact, a large amount of data loss and cycles of collisions and retransmissions are observed when PPC is present.

Consider two pairs of nodes contending for the channel to fulfill their communication requests, say *X* to *Y* and *A* to *B*, respectively (Figure 1(a)). PPC occurs if the following conditions are satisfied. (i) *B* and *X* are beyond the carrier sensing range (the dotted circles) of each other, and so are *A* and *Y*; (ii) *B* and *Y* are beyond each other's transmission range; (iii) *B* and *Y* are within each other's interference ranges (which depends on the distances between the sender and the receiver). The exact conditions for PPC to occur depend on various factors including the locations of the senders and receivers, the transmission, carrier-sensing and interference ranges. Basically, the likelihood of occurrence, represented by the shaded area in Figure 1(a), tends to increase when a small carrier-sensing range is selected or when the distance between sender and receiver is long. The scenario in Figure 1(a) is based on the default settings of Network Simulator (ns-2) where carrier-sensing range is set to 2.2 times of the

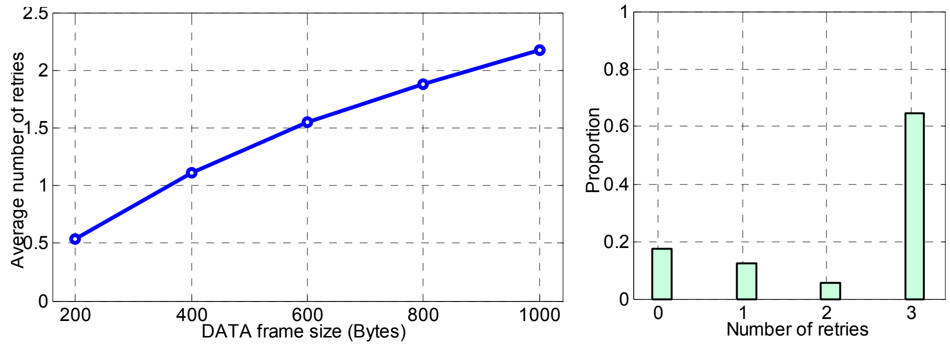

Figure 2. (a) Number of retries vs. DATA frame size; (b) Distribution of numbers of retries.

transmission range. It is considered too large/conservative with RTS/CTS handshake enabled. Exact calculation of its likelihood is beyond the scope of this paper. Readers are referred to [8], which includes a similar scenario.

Figure 1(b) illustrates how PPC occurs as a cycle of retransmissions. Sender *A* has a frame for receiver *B*, denoted as DATA1, so it first transmits an RTS. *B* then transmits a CTS in reply. After the handshake, *A* starts transmitting DATA1. Sender *X* also obtains a frame for receiver *Y* slightly later, denoted as DATA2, so it transmits an RTS to *Y*. *Y* then replies with a CTS even though *A* is transmitting. This is possible because *A* is beyond *Y*'s carrier-sensing range and *Y* could not overhear the CTS sent by *B* previously. As a result, neither DATA1 nor *Y*'s CTS can be correctly received by *B* (a collision marked as the leftmost star in Figure 1(b)) while *Y*'s CTS returns to *X* successfully. In addition, receiver *B* cannot set its network allocation vector (NAV), hence missing the media reservation information carried in *Y*'s CTS. But on the other hand, sender *X* proceeds with the transmission of DATA2 to *Y*. After waiting until the ACK timeout and backing off for the period of contention window (CW), sender *A* resends the RTS to *B*. Without any NAV set in *B*, it replies with a CTS immediately. Unfortunately, the second collision occurs between this CTS and DATA2 at *Y*. This time, their roles are switched. As shown in Figure 1(b), a loop of interruptions and retransmissions is formed. Such a logjam can last for several interruption exchanges, depending on the actual backoff intervals selected by the senders. When the two senders are in the same backoff stage and their contention windows grow in a synchronized fashion, the logjam may last longer than the time needed for the senders to reach the maximum number of retrials. This is particularly true for large-size DATA frames or the low data-transmission-rate since the DATA frame duration is long relative to the control frames and inter-frame periods. We performed ns-2 simulations of the scenario in Figure 1(a) to determine the number of retries and its relation with the frame size. Figure 2(a) shows that, when data rate is fixed to 1Mbps, the average number of retransmissions for each frame increases from 0.54 to 2.17 as the DATA frame size increases from 200 bytes to 1000 bytes. For the case of 1000-byte frames, Figure 2(b) also shows the histogram of frames with different numbers of retries they experience. For about 62% of frames, the default transmission limit of 4 with RTS/CTS enabled has been reached (1 transmission + 3 retries). Thus, measures must be taken to avoid permanent loss of packets due to long logjams like this.

III. Related Work

There have been numerous research efforts to enhance the collision avoidance and to improve the efficiency of 802.11 MAC, but they are not specifically designed to address PPC. These methods can be categorized into two groups based on the number of transmissions they attempt to coordinate. The first category intends to coordinate the transmissions of two sender/receiver pairs that can potentially interfere with each other. The second category aims at more general settings, including multiple flows and the capacity of the whole network. We briefly review these methods and explain their insufficiency if they were directly applied to PPC.

MACA-P [2] is proposed to achieve parallel communication in a two-pair scenario. It tries to schedule in order to synchronize DATA/ACK transmissions of the two pairs so that the ACK frame from one receiver will not interrupt the DATA reception on the counterpart receiver. It assumes that the two receivers can hear each other and then exchange such scheduling information, which does not hold for PPC. In [9], the RTS/CTS frames are sent at the highest power level, and the DATA and ACK at a lower power level, for the purposes of power conservation and collision avoidance. With enlarged RTS/CTS transmission ranges, one receiver can be notified and silenced by the other receiver's CTS, hence avoiding the PPC problem. However, those high-power RTS/CTS frames may themselves further cause collisions, due to their longer-range interference. In addition, the enlarged blocking range of RTS/CTS reduces the channel spatial reuse substantially. Some proposals [12][14] enhance the power control and planning for general scenarios to achieve more spatial reuse and more effective collision avoidance. However, they either demand a very accurate propagation and interference model or require sophisticated power control technology that can be unrealistic for current 802.11 devices [1]. Instead of changing the power, extending the receivers' carrier sensing range can also avoid the collisions in PPC. The receivers, then, will not send CTS if they can sense a contending transmission. But again, this can undermine channel spatial reuse. Furthermore, finding a proper carrier sensing threshold for various scenarios can be very challenging. A dynamic self-learning carrier sensing scheme is proposed in [6] to handle hidden/exposed node problems. Here, the sender monitors the historical RTS/CTS success rate and its relationship with the measured signal strength level. Based on the success rate of RTS/CTS, it speculates on the channel availability and performs carrier sensing. However, for PPC, the statistics of RTS outcomes is not consistent and could be misleading when used as indicator of DATA losses. For example, in the sequence shown in Figure 1(b), the RTS frames are always responded to by their intended receivers, but the subsequent DATA frames are lost.

The other camp takes a different perspective and focuses on not only the two directly involved traffics but the capacity in the senders' proximity. It is noticed that a sender engaged in multiple flows to a set of receivers should not persist with a particular receiver when this receiver is currently subject to high error rate of reception. For example, if an Internet gateway in a mesh network and one of its next-hop receivers are involved with PPC, the frames for that receiver may experience a long logjam, thus blocking the subsequent frames of other flows from being transmitted. As a result, the aggregate throughput of the mesh gateway, and thus the capacity of the mesh network, can be significantly degraded. In [4], CSDPS is proposed to defer the transmission and retransmission of a lost packet when it suspects a bad link state so that more resources can be assigned to other traffics with good link state. Fragouli *et al.* [7] enhances CSDPS by monitoring the history of RTS/CTS attempts and using that to limit the maximum number of RTS retries. Again, due to the failure of the collision avoidance of 802.11 in scenarios like PPC, RTS/CTS history is not a consistently reliable factor to predict DATA frame losses.

Our proposed method in Section IV below is based on the following critical observation. If a sender/receiver pair involved in PPC reduces its retrial intensity, the contending pair will experience a shorter interruption cycle and hence a higher throughput. In addition, if such reduction allows the former pair's sender to service instead other flows for which it has backlogged packets, then the entire system's throughput is increased. And this can be achieved via the multicast extension of RTS/CTS, i.e. MRTS.

IV. MRTS and SMRTS Solutions

Here, we briefly describe the Mutlicast RTS (MRTS). Interested readers can refer to [16][18] for more details. MRTS features a scheduler which can reorder the frames in sender's queue based on the status of its next-hop stations. To obtain the state information of the next-hop neighbors, an MRTS frame, in contrast to a unicast RTS frame as in conventional 802.11 MAC, is addressed to a set of receivers.

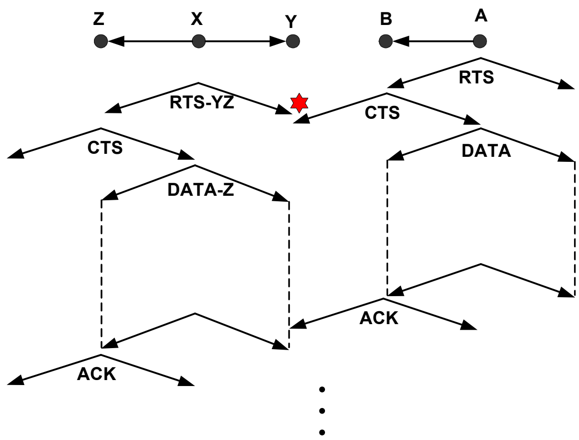

Figure 3. A scenario of MRTS employed in PPC

That is, an MRTS frame contains a list of next-hop receivers for which the sender has DATA frames currently queued. Each element of the list contains a receiver's address and the NAV of the corresponding frame. The priority among different receivers is decided by the order of their address appearance in the MRTS frame. That is, the earlier a receiver's address appears on the MRTS list, the sooner this receiver can return a CTS. This helps to avoid the collision of CTS frames returned by the receivers. The earliest receiver in a good channel state and having not set its NAV answers with a CTS. The sender finds the address of the responding receiver from the returned CTS. Then, the sender retrieves the corresponding frame from the queue and transmits it to that receiver.

Although MRTS is originally designed to exploit the channel diversity among the receivers, we find it a very good starting point to solve PPC. Figure 3 illustrates how MRTS copes with PPC. With frames of two flows backlogged at sender *X*, *X* transmits an MRTS addressed to both *Y* and *Z*. Due to a collision with *B*'s CTS or ACK frame, node *Y* will not return a CTS. Thanks to the multicast feature, node *Z* can instead reply to the MRTS. Upon receiving *Z*'s CTS, *X* starts transmitting the DATA frame to *Z*. As a result, both DATA frames from *X* to *Z* and *A* to *B* can be delivered simultaneously since *B*'s CTS or ACK will not interfere with *Z*'s DATA receptions, or vice versa. Note that, with the regular RTS/CTS scheme as in PPC, such collision of RTS may trigger a sequence of collisions at nodes *Y* and *B* (Figure 1(b)). When using MRTS, whenever there is collision of *X*'s MRTS frame at *Y*, node *X* automatically defers its transmission to *Y* and transmits a frame to *Z* first.

We can further rearrange the order of the receiver addresses contained in MRTS frame when a DATA transmission failure is detected through an ACK timeout. Such a treatment is offered by SMRTS. A DATA frame loss occurs either due to channel error or ineffectiveness of the RTS/CTS handshake, indicating a non-trivial probability of poor channel condition or unavoidable interference at a receiver. With this in mind,

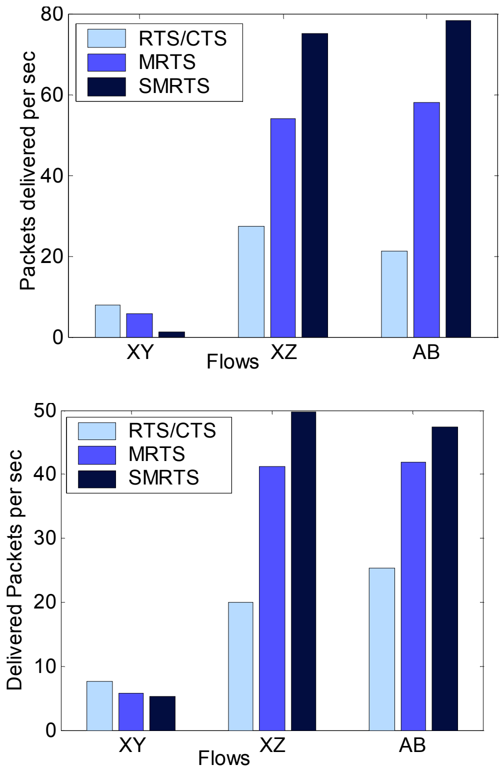

Figure 4. (a) Throughput performance comparison for the saturated case; (b) the unsaturated case

we further reduce the service time allocated for that receiver. This is achieved by changing the receiver's position in the next MRTS list. That is, while initiating a retransmission using another MRTS, the sender moves the address of that receiver to the end of the receiver list and shifts other receivers to the front. Thus, other flows will have better chances to be served in the next round. There is a trade-off between the aggregate throughput and absolute fair resource allocation among all the flows. This can be adjusted by choosing when to cycle-shift the receiver list, i.e., after how many failed DATA transmits. For example, the sender may wait for n consecutive DATA losses, instead of one, before shifting the MRTS list. The smaller the value of n, the more acutely it reacts to the DATA loss. To do this, we need to set and maintain a retry counter for each of the receivers. Whenever a frame for a receiver is retransmitted, we increment the counter of the corresponding receiver. And we reset it when a frame is successfully delivered or dropped for reaching the retry limit.

We test the performance of MRTS and SMRTS for the 5-node 3-flow scenario in Figure 3 by ns-2 simulation, and compare them with the original 802.11 RTS/CTS scheme. The results are shown in Figure 4. First, we generate a high-rate UDP traffic for each flow to saturate the network. The results shown in Figure 4(a) indicate that the throughputs of flows *XZ* and *AB* are significantly improved by MRTS with a small penalty on the throughput of flow *XY*. The performance is further improved by around 30% when using SMRTS. That is, less system resource is wasted by repetitive failed retransmissions of flow *XY* to achieve higher capacity in the area near the senders. Then we reduce the UDP traffic rate to investigate an unsaturated case. The results in Figure 4(b) show that the penalty on flow *XY* is negligible but the improvement on aggregated throughput is high.

V. Conclusions & Future Work

The 802.11 MAC provides a matured foundation for implementing wireless mesh networks. In order to solve problems emerging from such multi-hop networks, improvements have been proposed over the original wireless LAN centered design. In this paper, we identify the persistent pseudo-clearance (PPC) problem and present concise solutions to it. PPC can seriously undermine the per-hop reliability intended by the MAC layer and can propagate to upper protocol layers. Our solution, SMRTS, is essentially a multicast-based enhancement of the RTS/CTS dialog. It breaks the logjam between the senders involved in PPC effectively and efficiently, and requires minimal modification to the existing Standard. The investigation in PPC and SMRTS can be extended potentially in various directions, such as analytical studies of the likelihood of PPC occurrence and the expected length of the logjam. In addition, the shifting strategy of SMRTS may be further fine-tuned to balance the system throughput and resource allocation fairness under various network load settings.

References

[1] F.B. Abdesslem, L. Iannone, M.D. de Amorim, K. Kabassanov and S. Fdida. "On the feasibility of power control in current IEEE 802.11 devices," in *Proc. of IEEE PERCOMW,* 2006.

[2] A. Acharya, A. Misra and S. Bansal. "MACA-P: a MAC for concurrent transmissions in multi-hop wireless networks," in *Proc of PERCOM,* 2003.

[3] D. Aguayo, J. Bicket, S. Biswas, G. Judd and R. Morris. "Link-level measurements from an 802.11b mesh network," in *Proc. of ACM SIGCOMM,* 2004.

[4] P. Bhagwat, P. Bhattacharya, A. Krishna, and S. K. Tripathi. "Enhancing throughput over wireless LANs using channel state dependent packet scheduling," in *Proc. of INFOCOM,* 1996.

[5] V. Bharghavan, S. Demers, S. Shenker, and L. Zhang. "MACAW: A media access protocol for wireless LANs," in *Proc. of ACM SIGCOMM,* 1994.

[6] C. Chen, E. Seo, H. Kim and H. Luo. "Self-learning collision avoidance for wireless networks," in *Proc. of IEEE INFOCOM*, 2006.

[7] C. Fragouli, V. Sivaraman, and M. B. Srivastava. "Controlled multimedia wireless link sharing via enhanced class-based queuing with channel-state dependent packet scheduling," in *Proc. of INFOCOM,* 1998.

[8] M. Garetto, J. Shi and E.W. Knightly. "Modeling media access in embedded two-flow topologies of multi-hop wireless networks," in *Proc of MobiCom,* 2005.

[9] J. Gomez, A. T. Campbell, M. Naghshineh and C. Bisdikian. "Conserving transmission power in wireless ad hoc networks," in *Proc. of* ICNP, 2001.

[10] *IEEE. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications,* IEEE standard 802.11, 1997.

[11] S. Jain and S. R. Das. "Exploting path diversity in the link layer in wireless ad hoc networks," in *Proc. of IEEE WoWMoM Symposium,* 2005.

[12] E. Jung and N. Vaidya. "A power control MAC protocol for ad hoc networks," in *Proc. of ACM MobiCom,* 2002.

[13] T.-S. Kim, H. Lim, and J.C. Hou. "A coordinate-based approach for exploiting temporal-spatial diversity in wireless mesh networks," in *Proc. of ACM MobiCom*, 2006.

[14] A. Muqattash and M. Krunz. "A single-channel solution for transmission power control in wireless ad hoc networks," in *Proc. of ACM MobiHoc*, 2004.

[15] R. Saikat, C. Jeffrey B. and S. Se David. "Evaluation of the masked node problem in ad hoc wireless LANs," *IEEE Transactions on mobile computing,* Vol. 4, 2005.

[16] J. Wang, H. Zhai, and Y. Fang. "Opportunistic packet scheduling and media access control for wireless LANs and multi-hop ad hoc networks," in *Proc. of IEEE WCNC*, 2004.

[17] X. Yang and N. Vaidya. "On the physical carrier sense in wireless ad-hoc networks," in *Proc. of IEEE INFOCOM*, 2005.

[18] J. Zhang, Y.P. Chen and I. Marsic. "Adaptive MAC scheduling using channel state diversity for wireless networks," in *Proc. of IEEE WiCOM*, 2006.