

High Performance Architectures for IP-based Multi-hop 802.11 networks

Arup Acharya, Archan Misra
IBM T J Watson Research Center

Sorav Bansal
Stanford University

Abstract

The concept of a *forwarding node*, which receives packets from upstream nodes and then transmits these packets to downstream nodes, is a key element of any multi-hop network, wired or wireless. While high-speed IP router architectures have been extensively studied for wired networks, the concept of a “wireless IP router” has not been addressed so far. In this paper, we examine the limitations of the IEEE 802.11 MAC protocol in supporting a low-latency and high-throughput IP **datapath** comprising multiple wireless LAN hops. We first propose a wireless IP forwarding architecture that uses MPLS with modifications to the 802.11 MAC to significantly improve the packet forwarding efficiency. We then study further enhancements to the 802.11 MAC that improve the system throughput by allowing a larger number of concurrent packet transmissions in multi-hop 802.11-based IP networks. With 802.11 poised to be the dominant technology for wireless LANs, we believe a combined approach to MAC, packet forwarding and transport layer protocols is needed to make high-performance multi-hop 802.11 networks practically viable.

1. INTRODUCTION

The deployment of wireless LANs, based on the IEEE 802.11 [1] family of standards, has witnessed the proverbial “exponential growth” over the past two years. As an access technology, short-range wireless LANs (wLANs) are now viewed as a legitimate complement, if not a threat, to the packet-based 3rd generation (3G) wide-area cellular networks that are witnessing a slow, but steady, worldwide deployment. The current wLAN architecture is however, self-limiting, since it only considers topologies where mobile devices access the wired backbone infrastructure over a single (last hop) wireless link. We believe that multi-hop IP-based wireless networks will prove to be a compelling access alternative in the future. The bulk of multi-hop wireless networking research has so far concentrated on the mobile ad-hoc networks (MANET) scenario, where the network topology exhibits rapid changes due to the movement of individual nodes. There are however, a variety of environments, where fixed (or relatively static) multi-hop wireless networks offer an extremely cost-effective, and easily deployable, access alternative. Potential examples of this include:

- *Community Networks* where rooftop antennas are used to connect a cluster of residential properties, with only designated buildings having a gateway to the Internet.
- *In-Building Networks* in malls, hotels and apartment blocks, where wireless links are used as relays between individual consumer devices and a wired “point of presence”, eliminating the need for expensive wiring and maintenance.

There are a number of research challenges in multi-hop 802.11 networks. The MANET community has

traditionally focused on the multi-hop *routing* problem, and concentrated on the development of routing protocols (e.g., AODV [2]) that aim to establish loop-free paths in rapidly varying network topologies. However, even for static multi-hop topologies, there are fundamental problems that need to be solved. For example, while the data transmission rates associated with the 802.11 family¹ are increasing rapidly (802.11a offers rates of 54 Mbps, while rates of up to 108 Mbps are already under development), the low throughput of packet-based 802.11-based multi-hop wireless networks remains a major performance bottleneck [3] even in static scenarios. The impact of 802.11 MAC on TCP throughput of flows that span multiple hops is an active area of current research. Unlike a wired multi-hop network, packets belonging to the same TCP flow contend with each for channel access on successive links, reducing end-to-end throughput [3][11]. In addition, the emergence of multi-rate wireless LANs (where the data rate varies with the transmission distance) will require modifications to conventional IP routing protocols: routing protocols will now have to choose between multiple high-speed hops vs. fewer lower-speed ones [12].

In this paper, we focus on challenges and architectural approaches for increasing the throughput in IP based multi-hop 802.11 networks by modifications at the MAC layer targeted for operation in a multi-hop environment and coupling the MAC layer enhancements with IP

¹ While there are different flavors of 802.11 such as a,b,e and g, these differ primarily in the modulation schemes employed, and the raw data rate achieved, at the physical layers. All these schemes however share a common medium access control (MAC) layer. Since we focus primarily on improvements at the MAC operation, our discussion does not distinguish between these different variants, and applies to all of them.

packet forwarding and routing. From a packet forwarding perspective, the wireless medium offers both unique opportunities and unique challenges, which are absent in conventional wired networks. We shall first describe the need for designing a ‘wireless router’ or a forwarding node whose primary function is to receive packets from one neighbor and transmit them to a second neighbor using the *same* wireless interface. This architecture of a forwarding node is geared towards execution of the packet forwarding operation entirely within the wireless network interface card (NIC), without participation of the host CPU by combining medium access (MAC) with next-hop address lookup using MPLS [16] labels. The second concept we discuss in this paper is the idea of a medium access protocol that is specifically targeted for a *multi-hop* environment, i.e., a MAC that allows for simultaneous data transmissions in neighboring cells, thereby increasing the overall system throughput. As we will see, a drawback of the RTS/CTS based 802.11 MAC is that transmissions in neighboring cells are disallowed due to what is commonly referred to as the “exposed node” problem. As discussed in [3], this exposed node problem can cause unexpected poor performance for IP transport protocols, such as TCP. In essence, TCP packets from the same flow, or from different flows, contend for the same channel in different hops, and end up getting dropped even though the overall channel utilization is very low. We will present a technique for achieving better spatial reuse through localized, but completely distributed, “synchronization” of packet transmissions between neighboring wireless forwarding nodes.

2. BRIEF OVERVIEW OF 802.11 MAC

To understand the performance limitations associated with the 802.11 MAC layer [4] in multi-hop wireless environments, we first provide an overview of the operation of the 802.11 DCF, the distributed access control mechanism that is typically used in multi-hop environments. The 802.11 contention resolution and access control mechanism is based on the Carrier-Sense Multiple-Access with Collision Avoidance (CSMA-CA) algorithm, which consists of two distinct steps:

- a) A node interested in sending an IP packet first senses the channel to see if the channel is idle. If the channel is busy (some activity from a neighboring node is detected), the node backs off and idles for a period of time before retrying. This exponential back-off can significantly increase the total time of a packet forwarding operation, especially in situations where a node is guaranteed uncontested access to the channel.
- b) On sensing an idle channel, the node transfers data to the neighboring recipient using a 4-way handshake consisting of the following steps:

- 1) Node A sends an RTS (request-to-send) to node B, effectively informing anyone within A’s neighborhood that the medium is “reserved” for the duration T_{RTS} .
- 2) Node B sends the CTS (clear-to-send) to node A, specifying the time interval T_{CTS} during which A is permitted to send this data—the CTS informs all neighbors of B that the channel is reserved for the duration T_{CTS} .
- 3) Node A then send the data packet itself— this data transfer phase immediately follows the reception of the CTS.
- 4) On successful reception, node B sends the final data ACK.

For contention resolution, 802.11 uses a timer-based exponential back-off scheme, For details of the 802.11 DCF protocol, the reader is referred to [4].

3. INTERACTIONS BETWEEN IP PACKET FORWARDING, MEDIUM ACCESS AND HOST PROCESSING

The data forwarding operation in the wireless environment differs from the corresponding function in wired networks in a fundamental way:

In a wired network, a forwarding node typically has at least two physical network interfaces, with the forwarding functionality consisting of receiving a packet over one physical interface and subsequently sending it out over a second interface². In contrast, a node N, with a single wireless interface, may act as a forwarding node simply by retransmitting a packet that it received, over the *same* interface. In effect, N acts as an intermediary for two nodes that are each within the communication range of N but not directly within the range of each other.

Accordingly, packet forwarding in the wireless environment does not typically imply the transfer of a packet between distinct interfaces on a single host. A conventional implementation of packet forwarding thus involves the reception of a packet on the wireless interface, transfer of the packet up the host’s protocol stack to the IP layer where a routing lookup is used to determine the IP (and MAC) address of the next hop, and subsequent transmission of the packet using the same wireless interface to the MAC address of the next hop. This form of forwarding suffers from two key deficiencies:

- a) The forwarding node is thus involved in two separate contention-based channel access attempts during the

² In high-end routers/switches, the packet is transferred from one interface to another via a dedicated switching fabric, while in software based routers, the packet is processed by the host CPU (e.g. route lookup) between packet reception on one interface and subsequent transmission on another.

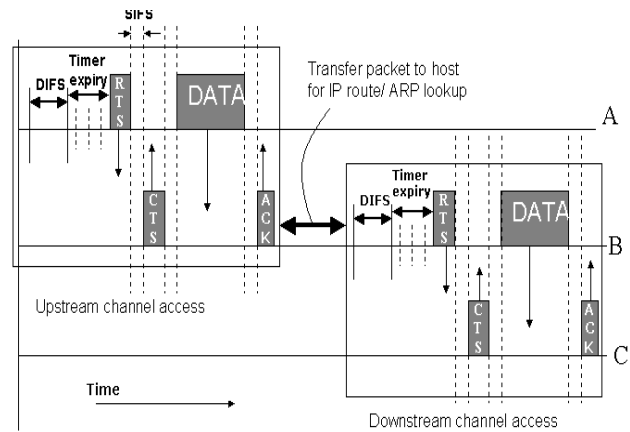
forwarding process: once to receive the packet (from the upstream node) and again to “forward” it (to the downstream node), and must thus suffer the contention resolution overhead twice.

b) The same IP packet makes an unnecessary roundtrip between the memory on the network interface card (NIC) and the host’s memory (accessed by the host software). This roundtrip not only loads the processor of the forwarding node, but also suffers from additional delays in transfers between the NIC and the host operating system.

The current IEEE 802.11 DCF MAC algorithm has been designed implicitly for either receiving or transmitting a packet, but not for a forwarding operation, i.e. receiving a packet from an upstream node and then immediately transmitting the packet to a downstream node as an atomic channel access operation. Figure 1 shows the 802.11 MAC-based forwarding operation of a node B, which is transferring an IP packet from an upstream node A to a downstream node C. After the IP lookup function in host A determines that B is the next hop of the DATA packet, the packet is transferred to A’s NIC. The MAC implementation on A’s NIC then performs a 4-way handshake (including any backoff timer-based countdown that may be needed to gain access to the channel) to forward the packet to B’s NIC. At B, the packet is transferred to the main memory from the NIC, and the host CPU is notified (e.g. via interrupts) for further processing of the packet by the IP protocol stack running on the host CPU. When the host software (IP protocol stack) picks up this packet from the head of the packet queue, it would execute the same steps as those executed at A, i.e., perform lookups to determine the IP address and then the MAC address of the next hop (C), insert the MAC-layer header (corresponding to next hop C) and transfer the packet to the NIC. This packet is now treated as an independent data transfer between the nodes B and C; accordingly, B performs the usual backoff timer countdown before initiating an RTS-CTS-DATA-ACK exchange with C. Once this handshake is successfully completed, the packet is received by C’s NIC, at which point the whole forwarding process is repeated. As with the initial data transfer (from A to B), node A is blocked from any activity (by the RTS sent by B) for the entire duration of the 4-way exchange between B and C.

3.1 Efficient Packet Forwarding using MPLS and 802.11 enhancements

The packet forwarding performance can be significantly improved if both a) the next hop for the packet (node C) can be determined within the NIC, without needing to perform the routing lookup in the host, and b) the packet can be routed from node B to C in an atomic fashion, without incurring the overheads of a separate transmission attempt. We now present our MPLS+DCMA solution



802.11 Based packet forwarding

Figure 1: Packet Forwarding Operation using 802.11

(full details are available in [7]) that addresses both these points.

To perform the next-hop lookup to be performed inside the NIC, we can use MPLS[5], a well-known IP-compatible technology to set up labels that enable us to determine the MAC address of the downstream node. In our solution, the network interface card is enhanced to store a label switching table, consisting of an incoming MAC address, an incoming label, an outgoing MAC address and an outgoing label, as shown in Figure 2. Labels are associated with routes or destinations, i.e., all entries in the label switching table that refer to the same route, will share the same outgoing MAC address (of the next hop) and outgoing label. For example, let an entry in the switching table of B be $\langle A, L_{AB}, C, L_{BC} \rangle$. The interpretation of this entry is that any packet received at B from A with a label L_{AB} will use C as the next downstream hop with a label L_{BC} . The combination of the outgoing label L_{BC} and the MAC address of the next hop node C, essentially defines a specific route to a destination, say Z. If B has another neighbor, say D, which uses B to reach another node Y, then there will be a corresponding entry in the label switching table $\langle D, L_{DB}, C, L_{BC} \rangle$. *The number of distinct outgoing labels is then equal to the number of destinations in the network.* Since each label is unique only to a single hop, the same label may be re-used by different nodes of the network.

The label-switching table is populated by a label distribution protocol running at the host in conjunction with a routing protocol. It is fairly straightforward to piggyback labels with route updates or run a standard separate label distribution protocol (as in wired networks) e.g., [5][6]. Accordingly, in the ideal case, an IP-to-label mapping is performed only at the source node of the

packet, with subsequent nodes forwarding the packet via label-swapping at their NICs.




Figure 2: Label-based Forwarding (Host/NIC components)

We eliminate the overheads of multiple channel accesses by defining the Data-driven Cut-through Medium-Access Protocol (DCMA) [7] as a simple extension of the 802.11 DCF. DCMA combines the ACK (to the upstream node) with the RTS (to the downstream node) in a single ACK/RTS packet that is sent to the MAC broadcast address. The payload of the ACK/RTS packet contains the MAC address of the upstream node, the MAC address of the downstream node and a label intended for use by the downstream node to figure its next hop. The reservation for the downstream hop is attempted only after successfully receiving the DATA packet from the upstream node. Since the downstream node (and all other neighboring nodes of the forwarding node) is assured to be silent till the completion of the ACK from the forwarding node, our piggybacking scheme guarantees preferential (and contention-free) channel access to the downstream transmission. Cut-through in DCMA fails when the downstream node fails to respond to the ACK/(RTS) with a positive CTS; the forwarding node then simply queues the packet in the NIC queue and resumes the normal 802.11 channel access. The timing diagram in Figure 3 explains the operation of DCMA. Assume that node A has a packet to send to node D. A sends a RTS to B, which includes a label L_{AB} associated with the route to D. Assuming that its NAV⁴ is not busy for the proposed transmission duration, B replies with a CTS. B receives the DATA packet, and then sends a RTS/ACK control packet, with the ACK part addressed to A, and the RTS part addressed to C, along with a label L_{BC} . C's actions would be analogous to B, except that it uses the label L_{CD} in its RTS/ACK message.

⁴ NAV or the Network Allocation Vector is a data structure used by 802.11 DCF at each node to track if there is an existing reservation of the channel.

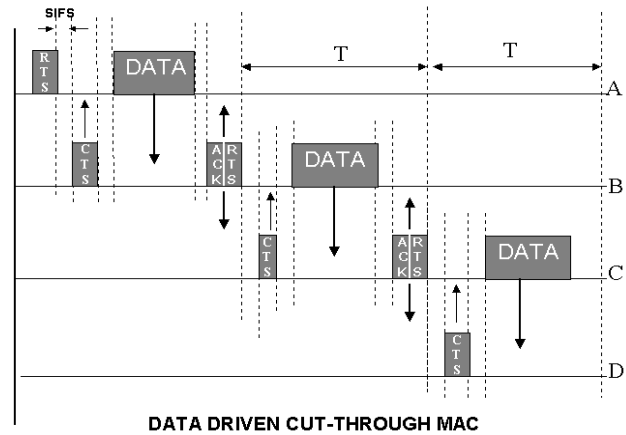


Figure 3: Data-driven cut-through MAC operation

Simulation results [7] have shown how DCMA can provide almost a 50% reduction in forwarding latency over a conventional 2Mbps 802.11b MAC for a 7-hop chain topology, even when we ignore the overheads of the NIC-host packet transfers at each intermediate node.

4. CONCURRENT PACKET TRANSMISSIONS IN MULTI-HOP ENVIRONMENTS

The 802.11 MAC protocol, and its variants, are primarily designed for a single-hop wireless environment, where nodes typically form a clique and communication always takes place over a single wireless hop (often to a base station providing connectivity to the wired infrastructure). In such a “single-cell” environment, the 802.11 MAC contention resolution mechanism focuses primarily on ensuring that only a single sender-receiver node pair receives collision-free access to the channel at any single instant. The 802.11 MAC does not seek to exploit the spatial diversity inherent in multi-hop networks, where different sets of nodes are able to concurrently communicate with different sets of neighbors. By exploiting this spatial diversity, we should be able to significantly increase the number of concurrent transmissions, by distinct sender-recipient node pairs that are “spaced sufficiently apart”.

The fundamental MAC constraint in a shared wireless medium is that: *no receiving node can be within the reception range of more than one simultaneously transmitting node*, since such concurrent transmissions will lead to collision and incorrect reception at the receiver. This is shown in Fig. 5a, where the parallel transmissions ($P \rightarrow Q$ and $A \rightarrow B$) would lead to a collision at a common receiver (B). While any MAC should prohibit only such parallel transmissions, the

802.11 DCF MAC however, imposes a more rigorous constraint: any node that is a neighbor of either participant (sender or receiver) in an ongoing RTS/CTS/DATA/ACK exchange must remain quiet for the entire duration of the 4-way exchange, i.e. it can neither initiate a parallel transmission (by sending a RTS) nor respond with a CTS during this period. To consider the differences between these two constraints, see Figs 5b and 5c, where Q and B are one-hop neighbors, and A's transmission range does not include Q (and vice versa), and P's transmission range does not include B (and vice versa). It is clear that the transmission patterns shown in cases (3) and (4) shown in Fig. 5d are not inherently feasible: in case (3), B's transmission to A would collide with P's transmission at Q, while in case (4), A's transmission (to B) would collide with Q's transmission (to P) at B. Now consider cases (1) and (2). For case (1), since A's transmission range does not include Q and P's transmission range does not include B (this is illustrated in Fig. 5b), the two transmissions should be allowed to proceed in parallel; a similar argument applies to case (2) as well (illustrated in Fig. 5b). The 802.11 MAC however prohibits both these scenarios, since any node that overhears a RTS or CTS must be silent for the interval specified by the RTS/CTS. So, for example, in Fig. 5b, after the RTS/CTS exchange takes place between A and B, 802.11 disallows node Q from responding to a RTS from P (during the subsequent DATA/ACK periods). This is commonly referred to as the "exposed node" problem. The restrictive behavior is needed since both the sender and the recipient of a data packet *revert between transmitting and receiving roles multiple times* over a continuous interval during the packet transfer. Since the DATA packet recipient acts as a receiver during the RTS and DATA portions, and the sender acts as a receiver during the CTS and ACK portions, the entire neighborhood of both nodes needs to be effectively silenced during the entire duration of the 4-way handshake (to preempt the possibility of a collision when either node is in a reception phase).




Fig. 5: 802.11 MAC's failure to exploit spatial reuse

As a result of this restrictive operation, 802.11-based multi-hop IP networks exhibit a very low degree of concurrency, and hence observe a sharp drop in the overall achievable system throughput.

4.1 Increasing Concurrent Transmissions Through Better Spatial Reuse

Three distinct approaches towards improving the spatial reuse in multi-hop wireless networks have been proposed. The first approach involves the use of *power-control algorithms*, where the fundamental goal is to reduce the zone of interference by lowering the transmission power of a transmitting node to the minimum needed to sustain communication with the intended neighboring destination node. One such representative power-control algorithm is PCMA [8], where both the sender and recipient nodes always ensure that their transmission power level never exceeds a value that would cause unacceptable interference at any of their neighbors.

The second approach involves the use of *directional antennas* (e.g., [9]), where a wireless node uses multiple antennas, each "tuned" to a specific portion (cone) of the total three-dimensional space. By limiting the zone of interference of each antenna to a specific orientation, the MAC layer can then allow for neighbors to be engaged in concurrent transmissions, as long as their transmissions are not directed towards one another. Unlike the omnidirectional 802.11 protocol, directional MAC protocols cannot guarantee that all neighbors of the data recipient are idle through the entire data transmission. Accordingly, it is more likely that the cut-through DCMA protocol presented in section 3 would be less successful in directional antenna-based environments, since neighboring nodes are more likely to be engaged in parallel (not necessarily synchronized) communications. *We can thus see that, from the viewpoint of the IP layer, multi-hop wireless networks are likely to present an interesting choice between an improvement in the overall system throughput or a reduction in the forwarding latency of an individual flow.* Both the power-control and directional antenna approaches can work directly with the 802.11 DCF MAC, since they essentially logically partition links into non-interfering sets.

A third approach, which we adopt, is to intelligently modify the MAC itself, and thus relax the unduly harsh restrictions of the 802.11 MAC. Our approach, called MACA-P [10], is one such 802.11-like protocol that provides synchronized parallel transmissions as long as the following fundamental constraint is not violated. In other words, it permits cases (1) and (2) in Figure 5, while still prohibiting cases (3) and (4). MACA-P's key idea is to allow neighboring nodes to *synchronize* their reception periods, *so that one-hop neighbors switch between transmitting and receiving roles in unison at explicitly*




Figure 6: Concurrent Transmissions in MACA-P

defined instants, and thus avoid the problem of packet collisions. MACA-P modifies the DCF handshake by introducing a variable “control gap” between the RTS/CTS exchange and the DATA and ACK phases; this gap allows neighboring nodes an opportunity to synchronize their DATA/ACK phases. We add additional information in the RTS/CTS messages to explicitly delineate the DATA and ACK transmission intervals, which may be non-contiguous. The basic operation of MACA-P can be understood from the timing diagram in Figure 6 (corresponding to case (2) in Figure 5) where the RTS from B to A sets up the “master transmission” schedule. The subsequent control gap is then utilized by B’s neighbor, Q, to set up an overlapping “slave transmission” to P. Note that MACA-P only aligns the starts of the DATA and ACK phases, and thus allows variable-sized packet transmissions to proceed in parallel. In a similar manner, MACA-P also allows case (1) in Figure 5: when Q receives an RTS from P, it sends a CTS that alters the suggested transmission schedule by P to align it with the previously scheduled transmission A→B. In many cases, MACA-P can be combined with the pipelined DCMA protocol described in Section 3, since a node receiving a data packet is still assured of silent neighbors (at the end of the ACK transmission). Simulation results [10] have shown a 200% improvement in throughput for MACA-P when compared to 1Mbps 802.11b for certain topologies. This significant increase in overall system throughput is achieved by dramatically increasing the degree of concurrency in the system.

5. CONCLUSIONS

In this paper, we discussed some challenges associated with creating a high-performance multi-hop 802.11-based wireless **datapath** for IP traffic, and showed how some features of the 802.11 MAC contribute to either high forwarding delay or poor system throughput in multi-hop wireless environments. We first introduced the concept of a *wireless router*, and presented enhancements to the 802.11 MAC along with the use of MPLS labels that

allow packet forwarding to be executed entirely within the wireless NIC card. Next, we discussed the limitations of the 802.11 MAC for supporting concurrent transmissions in neighboring cells and discussed an approach for improving the system throughput through better spatial reuse. These approaches show how some IP-layer functions, such as next-hop lookup, may be performed instead at the MAC layer to improve the performance of TCP/IP-based applications in wireless networks.

Several other interactions between the 802.11 MAC layer and IP protocols need further study. Since 802.11 will be the dominant technology for wLANs, a fresh look at integrating the IP stack and the wireless MAC is justified to realize the vision of an all-wireless IP network.

7. References

- [1] IEEE Computer Society. 802.11 : Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, June 1997.
- [2] C. Perkins, E. M. Belding-Royer, and S. Das. *Ad Hoc On Demand Distance Vector (AODV) Routing*. IETF Internet draft, draft-ietf-manet-aodv-10.txt, March 2002 (Work in Progress).
- [3] S. Xu and T. Saadawi. Does the IEEE 802.11 MAC Protocol work well in multihop wireless ad hoc networks? *IEEE Communications Magazine*, June 2001.
- [4] B. Crow, I. Widjaja, J. G. Kim and P. T. Sakai. *IEEE 802.11 Wireless Local Area Networks*. *IEEE Communications Magazine*, Sept '99.
- [5] E. Rosen, A. Viswanathan, R. Callon. *Multiprotocol Label Switching Architecture*, IETF RFC 3031, Internet Engineering Task Force, January 2001.
- [6] L. Andersson, et al. LDP Specification , IETF RFC 3036, Internet Engineering Task Force, January 2001.
- [7] A. Acharya, A. Misra and S. Bansal, A Label-switching Packet Forwarding Architecture for Multi-hop Wireless LANs, *ACM WoWMoM*, September 2002.
- [8] J. P . Monks, V. Bhargavan and W. Hwu. *A Power Controlled Multiple Access Protocol for Wireless Packet Networks*. Proceedings of IEEE INFOCOM 2001, June 2001.
- [9] Y.Ko, V. Shankarkumar and N.H. Vaidya, Medium Access Control Protocols using Directional Antennas in Ad Hoc Networks, Proceedings of IEEE INFOCOM'2000, March 2000.
- [10] A. Acharya, A. Misra and S. Bansal, *MACA-P: A MAC for Concurrent Transmissions in Multi-hop Wireless Networks*, Proceedings of IEEE Pervasive Communications and Computing Conference (PERCOM), March 2003.
- [11] Z. Fu, et al, On TCP Performance in Multihop Wireless Networks, Proceedings of IEEE INFOCOM 2003, March 2003.
- [12] B. Awerbuch, D. Holmer and H. Rubens, High Throughput Route Selection in Multi-Rate Ad Hoc Wireless Networks, Johns Hopkins University Technical Report, March 2003.