RC22506 (W0207-012) July 2, 2002
Computer Science

IBM Research Report

Challenges in high performance data forwarding in multi-hop wireless networks

Arup Acharya, Archan Misra
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Sorav Bansal
IBM Solutions Research Center
Block 1, Indian Institute of Tech.
Hauz Khas, New Delhi 110016
India

Research Division
Almaden - Austin - Beijing - Delhi - Haifa - India - T. J. Watson - Tokyo - Zurich

Challenges in high performance data forwarding in multi-hop wireless networks

Arup Acharya
IBM TJ Watson Research Ctr
arup@us.ibm.com

Archan Misra
IBM TJ Watson Research Ctr
archan@us.ibm.com

Sorav Bansal
IBM India Research Labs
soravban@in.ibm.com

Abstract

We focus on challenges associated with creating a high-performance **datapath** comprising of multiple wireless LAN hops. We believe that 802.11 will be the dominant technology for wLANs and a combined approach to MAC, packet forwarding and transport layer protocols will be needed to make high-performance multi-hop 802.11 networks practically viable. The first challenge we see is to revamp the well-known MACA protocol used by 802.11 from a single-cell MAC in a direction that allows neighboring cells to operate simultaneously whenever possible, thereby increasing the overall system throughput. The second challenge we discuss is the notion of a "wireless router" or a forwarding node, whose primary function is to receive packets from one neighbor and transmit them to a second neighbour using the *same* wireless interface. This requires combining channel access functionality with that of next-hop address lookup within the network interface card without host participation. The third set of challenges arise from the effects of physical/MAC layer characteristics on network connectivity (i.e. whether two nodes are neighbors depends on the rate used), transport layer performance (i.e contention for the physical channel among neighboring hops lead to packets of the *same flow* contending with each other) and use of MAC contention mechanisms as a means for supporting transport-layer congestion control.

1. INTRODUCTION

Wireless networks have received an inordinate degree of attention from the research community over the last 5-7 years. The bulk of this research can be classified into the following distinct categories:

1. *Improvements in MAC protocols for single-hop wireless LANs.*

Enhancements such as QoS differentiation (e.g., [802qos]) and fair bandwidth sharing in IEEE 802.11, or the design of contention resolution schemes in the presence of uni-directional links, are designed primarily for the WLAN environment, where nodes attach to the wired backbone over a single wireless hop.

2. *Routing protocols for mobile ad-hoc networks.*

The emphasis here has been on the development of protocols (e.g., AODV [aodv], DSR [dsr]) that establish traffic routes in environments where node mobility results in rapid changes in network topology. Although such protocols sometime utilize the broadcast nature of the wireless medium (e.g., route snooping in DSR), their primary focus is on rapid recovery from link failures and the avoidance of long-lived routing loops in mobile environments.

The research community has, however, largely ignored the problem of efficient data forwarding in multi-hop, wireless environments. For example, we have only recently seen some research publications (e.g., [multi]) showing how TCP performs in multi-hop 802.11 networks. In this paper, we articulate the various challenges associated with efficient packet delivery in such multi-hop wireless networks, taking special care to explain how the wireless medium offers both unique opportunities and unique challenges which are absent in conventional wired networks. While the eventual solutions should undoubtedly be able to accommodate potential node mobility, it is important to realize that multi-hop wireless communication currently performs extremely poorly even when all nodes are purely static. For example, while the IEEE 802.11a standards offer maximum channel speeds as high as 54 Mbps, the throughput realized on experimental 802.11 networks is often O(100Kbps), even when data paths consist of a small number of wireless hops.

Our research goal is the development of algorithms and protocols for high-performance data delivery across multi-hop wireless paths in environments where all nodes use a common channel. Potential examples of such networks in commercial environments include *in-building* wireless networks in malls, hotels and apartment blocks, and *community* networks where rooftop antennas are used to create an ad-hoc wireless network in specific residential communities. We have identified at least three distinct areas of research:

- *Extend single-cell MAC protocols to operate in a multi-cell/multi-hop network by exploiting spatial diversity to significantly increase the number of permitted simultaneously active links.*
- *Devise a MAC that efficiently supports a packet "forwarding" operation (i.e packet reception followed by transmission as a combined operation). This requires the wireless NIC to be capable of determining the packet's next hop without invoking host processing and integrate channel access with next-hop address lookup.*
- *Impact of rate/distance tradeoffs on network topology, use of per-hop MAC control mechanisms to aid end-to-end transport-layer congestion control and the (wireless specific) phenomenon of packets of the same flow contending with each other for channel access on successive hops of an end-to-end path. effect of channel contention on successive packets of the same flow*.

Since 802.11 is the dominant technology for wireless LANs for now and the near future, we believe that response to these challenges should be guided by how the 802.11 protocols can be improved, and as necessary, in a radical fashion.

2. MEDIUM ACCESS CONTROL (MAC) FOR MULTI-CELL/ MULTI-HOP NETWORKS

The broadcast nature of the wireless medium implies that a transmitter and receiver node can communicate effectively as long as the MAC layer ensures adherence to the following fundamental constraint: *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 (fig 1*).

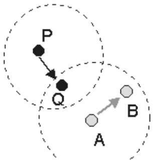

Fig. 1

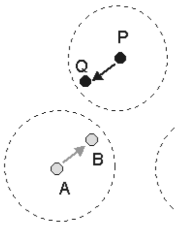

Fig. 2a

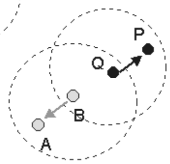

Fig. 2b

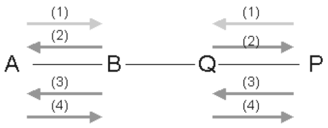

Fig. 2c

Current work on 802.11 (research[macaw] and standardization) however impose a more rigorous constraint: the MAC layer effectively ensures that *no node that is a one-hop neighbor of either the sender or the receiver of a data packet may be engaged in any communication activity (either transmitting or receiving) during the entire 4-way (RTS-CTS-DATA-ACK) exchange*. To consider the differences between these two constraints, see Fig.2 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.2c 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. For case (1), however, since A's transmission range does not include Q and P's transmission range does not include B (Fig. 2a), the two transmissions can proceed in parallel; a similar argument applies to case (2) as well (Fig. 2b). The 802.11 MAC is unduly restrictive and prohibits cases (1) and (2) essentially because 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 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 is effectively silenced during the entire duration of the 4-way handshake.

Significant improvement to the overall system throughput of multi-hop networks can however be realized if this constraint on concurrent packet exchanges can be relaxed or modified. In fact, the initial papers on CSMA-CA (.g., [macaw][dfwmac]) alluded to the possible exploitation of spatial diversity for parallel (concurrent) transmissions, but did not proceed with research in that direction. Recent attempts at improving the spatial reuse of the multi-hop network typically focus on two approaches, both of which fundamentally aim to *reduce the size of the one-hop neighborhood* and thus allow the network to be *partitioned* into a greater number of zones of concurrent transmissions:

a) Power control algorithms, e.g. [pcma]
b) Use of directional antennas, e.g. [direct].

While simulation studies indicate that both approaches can significantly improve the aggregate channel capacity of multi-hop networks, they do suffer from certain drawbacks. Distributed versions of power-control protocols require nodes to include and decipher the transmission power levels in the header of MAC control packets. In real-life situations, where the interference range is larger than the actual packet reception range, a node can suffer interference effects from neighboring transmitters even though it cannot correctly receive their packets (and thus cannot correctly perform the appropriate power-level computations). Directional antennas, on the other hand, use sophisticated hardware and phase-modulation strategies, and may not prove to be cost-effective solutions for large-scale deployment, especially in pervasive and mobile devices. Indeed, the focus of these current approaches is on simply increasing the number of disjoint network segments that can proceed in parallel, rather than on fundamentally trying to relax 802.11's constraint.

We believe that a fresh look is needed for medium access control that is inherently targeted for multi-cell, concurrent operation for high-performance multi-hop wireless networks to be realized in practice. To that end, we propose an approach where the 802.11 constraint (of silent neighborhoods of both sender and receiver) is replaced with the fundamental constraint, i.e a receiver is not in the neighborhood of more than one transmitter, and the 802.11 MAC is redesigned to support the fundamental constraint. The 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 and thus avoid the problem of packet collisions. This objective can be achieved without any basic changes to the 802.11 4-way handshake, by introducing a *variable control gap* between the RTS/CTS exchange and the DATA and ACK phases. One node generates a master transmission schedule and other neighboring nodes synchronize to that schedule (the interim control gap provides neighbors an opportunity to set up their individual transmissions). To support such a MAC for concurrent transmissions, there are a number of related problems that need to be addressed:

a) Since concurrent transmissions are never exactly synchronized (due to propagation delays and clock drifts), radios must be capable of capturing a packet with a sufficiently stronger signal even if arrives *later* than a packet with a much weaker signal. Current implementations of common-channel wireless cards can typically retrieve only the earliest arriving packet from a set of overlapping packet receptions or otherwise declare collision errors.
b) Since wireless environments have appreciably differing reception and interference ranges, any synchronizing scheme must allow for the possibility where a node cannot update its NAV (network allocation vector) with information about an impending transmission schedule from a transmitter beyond its transmission but within its interference range. MAC algorithms must therefore incorporate some form of adaptive learning where nodes dynamically learn about the feasible set of concurrent transmissions; the learning rate must be fast enough to incorporate appropriate levels of node mobility and topology changes.
c) Experimental and analytical studies show that purely de-centralized mechanisms for improving concurrency provide only moderate improvements in system throughput. To obtain dramatically higher throughput increases, the transmission schedules must be orchestrated over a larger neighborhood (than simply one hop neighbors) to avoid situations where the vast majority of nodes lie next to multiple master transmissions and are thus unable to exploit possible synchronized transmission opportunities (we have seen examples of this phenomenon on a grid-like arrangement of nodes).

3. FORWARDING NODE ARCHITECTURE (wireless router)

A router in wired network typically requires multiple network interfaces to act as a router or a forwarding node. In a multi-hop wireless network on the other hand, any node with a wireless network interface card can operate as a router or a forwarding node, since it can receive a packet from a neighboring node, do a route lookup based on the packet's destination IP address and then transmit the packet to another neighboring node using the same wireless interface. Medium access schemes to date, such as IEEE 802.11, have 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. Our second challenge in multi-hop wireless networks is a combined medium access and next-hop address lookup based that enables the entire packet forwarding operation to be executed within the wireless NIC without the intervention of the host protocol stack.

The motivation for integrating MAC with forwarding functionality arises out of one fundamental difference between wireless and wired networks: *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 interface. In contrast, a node N, with a single wireless interface, may act as a forwarding node by transmitting a packet to a node other than from which received the packet. 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*.

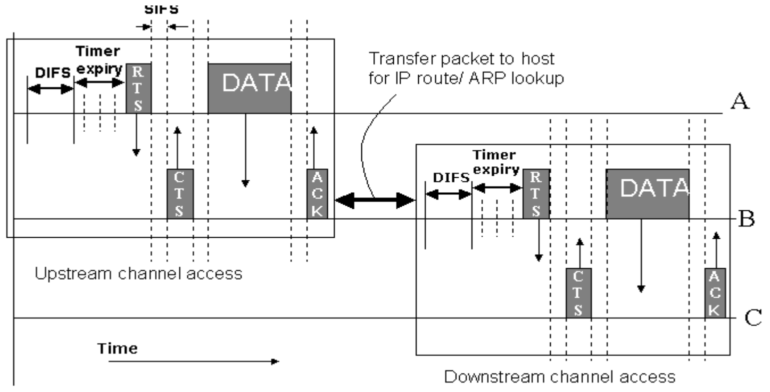

802.11 Based packet forwarding

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. The forwarding node is thus involved in two separate channel access attempts during the forwarding process: once to receive the packet and again to "forward" it. Moreover, the actual forwarding path involves two separate transfers of data between the memory on the network interface card (NIC) and the host's memory (accessed by the host software).

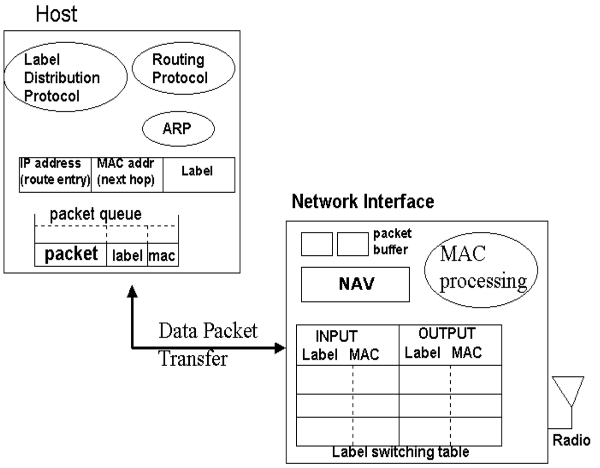

As an answer to the above challenge, we propose an architecture for a *forwarding node* that shifts the next-hop address lookup and packet forwarding functionality from the host processor to the wireless network interface card (NIC) by combining medium access control (MAC) for packet reception and subsequent transmission with address lookup in the interface card itself, using fixed-length addressing labels in the MAC control packets. 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 a outgoing label. Labels are associated with routes or destinations, such as in MPLS [mpls], with a separate label-distribution algorithm, e.g. LDP [ldp] used to distribute labels to appropriately reflect the traffic routes. This allows packet forwarding to be confined *entirely to the NIC*, which matches the label of an incoming packet with an entry in the data structure to determine the MAC address of the next hop node and the label to be used for that hop.

The above cut-through packet forwarding can be supported by extending the RTS/CTS/DATA/ACK exchange of 802.11 Distributed Coordination Function (DCF) mode of channel access in the following way, which we call DCMA (Data-driven Cut-through Medium Access). 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.

T		
ACK Flag	MAC address	
RTS Flag	MAC address	(out) Label

Payload of ACK/RTS control packet

The operation of DCMA can be understood by following the timing diagram shown to the right. Assume that node A has a packet to send to node D. A[1] sends a RTS to B, which includes a label L_{AB} associated with the route to D. Assuming that its NAV[2] 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.

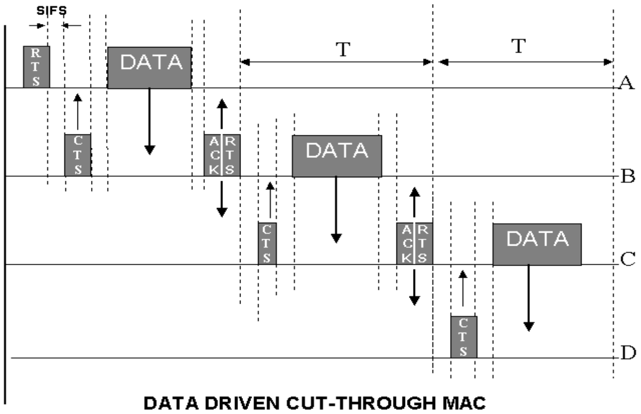

DATA DRIVEN CUT-THROUGH MAC

4. CONGESTION MANAGEMENT AND FLOW CONTROL

Congestion control in traditional wired Internet environments is managed by a combination of end-to-end flow control mechanisms (such as TCP congestion avoidance) and notification from intermediate nodes (through mechanisms such as RED and ECN) on the data path. We believe that significant wireless-specific improvements in both end-host behavior and intermediate-node congestion management mechanisms are needed before we can achieve acceptable throughput performance for reliable multi-hop communication.

It is now well established that transport protocols such as TCP perform very poorly over multi-hop wireless environments. While the initial reason was thought to be the higher loss rates of the wireless channel, recent research (e.g., [multi]) have demonstrated why such poor performance is often due to indirect packet forwarding failures caused by MAC-layer collisions between successive packets on adjacent hops. Initial research in this area reveals that such intra-session packet collisions are exacerbated by TCP's implicit rate control mechanism, which causes data traffic to occur in bursts. We are only beginning to see the first set of investigative results in this area, which often show how simple changes to TCP parameters, such as the slow-start rate or the size of the maximum congestion window, can significantly improve end-to-end performance. For example, the figure on the right shows the result of our simulation studies on an 802.11 chain, which show how the TCP throughput varies with changes to the maximum congestion window. We believe that TCP-based transport protocols are open to significant innovations that essentially smoothen the injection of packets into the network to ensure that consecutive packets do not cause MAC-layer contention. We also need to investigate in detail additional techniques,

---

[1] We assume the initial IP address to label mapping is done by the host and the label to be used, MAC address of the next hop and the packet is moved over to wireless interface card.

[2] 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.

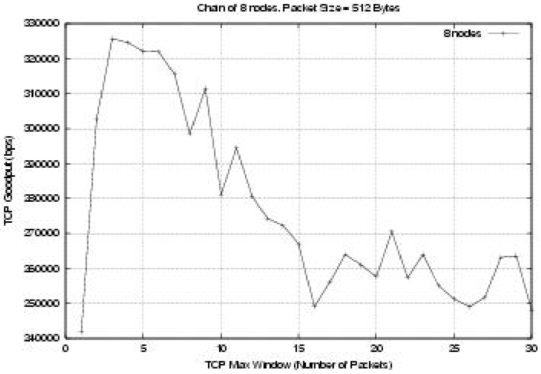

such as the use of asymmetric paths for TCP data and acknowledgement packets, or the interleaving of consecutive packets on multiple disjoint paths, that can further reduce the potential for MAC-induced losses. On a more general level, these possibilities illustrate our belief that high-performance multi-hop wireless networks require tight coupling between the MAC, routing and transport layers.

Intermediate forwarding nodes can also use wireless-specific mechanisms to significantly improve congestion management in such bandwidth-constrained environments. As new modulation schemes, such as IEEE 802.11a or 802.11g become available, nodes will have the ability to perform a *tradeoff* between their communication range (which indirectly affects the *network connectivity*) and the *transmission rate on downstream links*—in general, the faster the rate, the smaller the transmission radius. The ability to modify the network connectivity and transmission rates in response to the build up of queuing delays at intermediate nodes is a very promising tool for reducing the incidence of congestive bottlenecks. The development of decentralized algorithms that modify link transmission rates (and thus network topologies) in response to changing traffic patterns appears to be a very promising problem with no counterpart in traditional wired networks. We believe that additional MAC-layer enhancements will also prove to be necessary in this regard. For example, a downstream node that is experiencing significant queue buildup may use appropriate priorities in the MAC access algorithm to gain preferential access to the channel, and thus implicitly throttle the arrival of packets from upstream nodes. These examples should be evidence of the scope for significantly innovative research in integrated congestion control strategies in multi-hop wireless networks.

5. References

[802qos] IEEE 802.11 WG. *Draft Supplement to Standard for Telecommunications and Information Exchange between System, Medium Access Control (MAC) Enhancements for Quality of Service (QoS)*, IEEE 802.11e/D2.0, November, 2001.

[aodv] 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).

[dsr] D. B. Johnson and D. A. Maltz. *Dynamic Source Routing in Ad Hoc Wireless Networ*ks. In Mobile Computing, Chapter 5, pages 153-181, Kluwer Academic Publishers, 1996.

[multi] 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.

[macaw] V. Bhargavan, A. Demers, S. Shenker and L. Zhang. *MACAW : A Media Access Protocol for Wireless LAN's*. Proceedings of ACM SIGCOMM '94.

[dfwmac] J. Weinmuller, H. Woesner, J. Ebert, A. Wolisz. *Analyzing the RTS/CTS Mechanism in the DFWMAC Media Access Protocol for Wireless LANs*. IFIP TC6 Workshop on Personal Wireless Communications, April '95.

[pcma] J. P . Monks, V. Bhargavan and W. Hwu. *A Power Controlled Multiple Access Protocol for Wireless Packet Networks*. Proceedings of IEEE Infocom 2001.

[direct] Y.Ko, V. Shankarkumar and N.H. Vaidya, Medium Access Control Protocols using Directional Antennas in Ad Hoc Networks. IEEE INFOCOM'2000, March 2000.