Forward Focus: Medium Access Control for Ad hoc Networks

Brent Ishibashi and Raouf Boutaba

University of Waterloo
School of Computer Science
Waterloo, Canada
{bkishiba, rboutaba}@uwaterloo.ca

***Abstract*— The wireless medium must be effectively shared in an ad hoc network. The multihop nature of the network demands it, with packets often requiring several transmissions in order to reach their destination. This paper describes a mechanism for improving the IEEE 802.11 MAC specifically for ad hoc networks. The Forward Focus mechanism uses the forwarding nature itself to improve the efficiency and effectiveness of the forwarding process. A simulation-based performance evaluation reveals that the approach can yield significant improvements in network throughput.**

I. Introduction

Sharing of the wireless medium takes on critical importance in an ad hoc network. Between the increased communication requirement created by multihop packet delivery, and the expense of routing protocol overhead, nodes must send and receive considerable load using a relatively scarce resource(Figure 1). This, along with a sometimes-dense concentration of nodes, makes an effective medium access control (MAC) protocol crucial to creating an effective ad hoc network.

While the concept of ad hoc networking is not specific to one technology, IEEE 802.11 has frequently been used to provide the underlying physical (PHY) and MAC specifications. Well-established in wireless LANs, the 802.11 MAC provides a random access protocol, designed to provide multiple nodes with a fair opportunity to access the medium. However, several concerns are apparent. In addition to capacity limitations[7][18], it has been revealed that the MAC protocol can in fact interfere with the operation of higher level ad hoc protocols. In particular, normal MAC-level issues can create routing failures and significant upper-layer throughput interruptions[23][26].

While other works have focused on adapting protocols to avoid the difficulties experienced in ad hoc networks, this work takes a different approach. Rather than attempting to reduce the negative effects created by the ad hoc environment, the unique characteristics of the ad hoc environment are leveraged in order to improve the efficiency and effectiveness of the MAC protocol. As a result, the MAC is better able to support the ad hoc routing protocol.

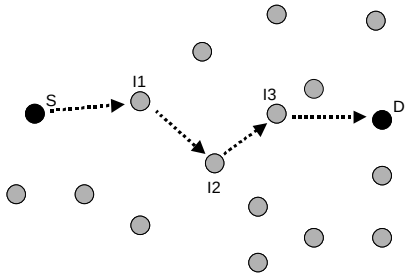

Fig. 1
Ad hoc network: Multiple transmissions are required to deliver data from S to D via intermediate nodes I1, I2, and I3.

Within this paper, the Forward Focus (FF) mechanism is proposed. FF allows the MAC to focus directly on the forwarding responsibilities that are critical to the network. It uses the forwarding nature to assist in fairly and efficiently re-using the wireless channel, while also rewarding nodes for serving as intermediates. As a result, the cost of forwarding is reduced, allowing traffic to be delivered more effectively to its destination. In addition, the overall efficiency of the MAC can be improved, resulting in improved network performance.

The remainder of the paper will be organized as

follows. In the next section, the IEEE 802.11 MAC is presented, followed by several important issues and proposed improvements. Section III highlights what we view to be key requirements for a MAC protocol in an ad hoc network. The FF mechanism is proposed in section IV, along with a protocol description for modifying the 802.11 MAC. Simulation results are presented in Section V, comparing 802.11's performance with and without the FF modifications. Finally, Section VI presents the conclusions we have drawn from this work.

II. Background

A. IEEE 802.11

IEEE 802.11[28][29] provides a mechanism for multiple stations to share a common wireless channel. Simultaneous transmissions may interfere with each other and prevent either from being correctly received. Therefore, the protocol works to avoid such collisions, which waste bandwidth and node resources. At the same time, the protocol aims to provide fair access to all nodes, in a manner that requires as little overhead as possible. In order to do this, 802.11 uses a carrier-sense multiple access with collision avoidance (CSMA/CA) mechanism.

Under this protocol, a node is not permitted to attempt to send if it senses that another node is already transmitting. If a node wishes to send, it must first sense the medium and determine that the medium is idle, before it can attempt its transmission. In order to do so, the medium must be sensed as idle for the DIFS period, the distributed interframe space. If the node senses that the medium is already being used, the node must wait.

Following the end of a transmission, a number of nodes may be waiting to access the medium. If nodes only wait for a DIFS idle period, all of them may attempt to send at the same time, increasing the probability of a collision. Instead, if a node senses the medium as busy, it initiates a backoff procedure. Under this procedure, the node must wait for the medium to be idle for a backoff period, chosen randomly from the interval of the contention window. Only after this backoff timer has expired can the node attempt its transmission.

To reduce the probability of a collision, the contention window is increased with each failed packet transmission. Each time a transmission fails, the contention window is doubled, up to a maximum value. Using this mechanism, under periods of congestion, where many nodes are attempting to use the medium, the contention window grows. Nodes are then, on average, forced to backoff for longer periods, allowing more nodes opportunity to access the medium with reduced probability of collision. The window is reset after a successful transmission.

Although this procedure reduces the probability of collision, it does not eliminate it. Therefore, in order to reduce the cost of a collision, 802.11 includes an optional request-to-send/clear-to-send (RTS/CTS) mechanism. RTS and CTS packets are short MAC protocol control packets, used to facilitate the contention process. If a collision occurs while using the RTS/CTS mechanism, the collision can be detected after minimal waste of resources, rather than waiting for the completion of the data packet transmission.

The RTS/CTS mechanism works as follows: when a node is permitted to access the medium, it transmits a RTS packet. If the receiver hears the RTS and is available to receive the data packet, it returns a CTS packet to the sender. On receiving the CTS, the sender then begins the data transmission, which, if received correctly, is acknowledged with an ACK from the receiver. All other nodes that hear either the RTS or CTS packet set a network allocation vector (NAV) timer, and remain silent until the exchange is completed. This use of virtual carrier-sensing, allows the mechanism to reduce the effects of the hidden-terminal problem.

An additional idle period must also be included between each pair of MAC-level frames. This short interframe space (SIFS) serves to separate consecutive packets, and allows for device processing and switching time. The complete packet exchanges, with and without RTS/CTS, are shown in Figure 2 and 3.

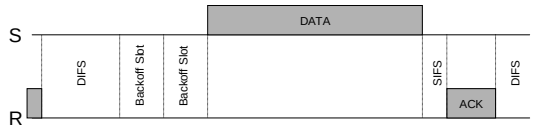

Fig. 2
Basic IEEE 802.11 Packet Exchange

B. Problems with IEEE 802.11 in Ad hoc Networks

Unfortunately, the protocol is somewhat inefficient. At a minimum, a node must remain idle for the DIFS period prior to attempting a transmission. However, if all nodes are in either a backoff or idle state, the channel will remain idle for additional slots, thereby wasting further network bandwidth. Longer backoff times, the

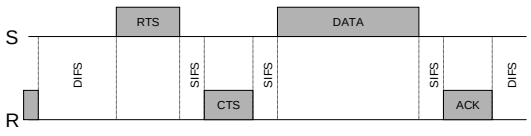

Fig. 3
PACKET EXCHANGE USING RTS-CTS MECHANISM

norm under congested conditions, aggravate this issue, increasing the possibility for idle contention slots.

The opposite extreme, reducing the size of the contention window will result in increased probability of collision. While the amount of time the medium will be idle may be decreased, increased collisions will reduce the effectiveness of the packet exchanges. Not only will bandwidth still be wasted, but nodes will also consume additional resources in attempting the transmissions.

While these properties are features of the protocol, some interesting effects become apparent when 802.11 is used in an ad hoc network. As a successful packet exchange requires both acquisition of the medium by the sender, and availability of the receiver, in an ad hoc network, many attempted transmissions fail. However, repeated failures within the MAC have been shown to cause significant problems in ad hoc routing. Inability to reach an adjacent node is interpreted as a failure in the link between two nodes. This link failure in turn forces a route breakage. As a result, link and route failures can occur, even in a network without mobility[26].

Additionally, a fairness issue has been revealed when 802.11 is used in ad hoc networks. Due to the resetting of the contention window after a successful transmission, successful nodes are favoured with shorter backoff times. Although this policy aims for fairness at the MAC-layer, it has been shown to be very disruptive to higher-level protocols. As a result, different traffic flows can experience vastly different throughputs. This suggests that fairness in an ad hoc network must be considered completely differently than in other wireless networks.

C. Improvements to the IEEE 802.11 MAC

While many works have focused on finding new technologies to be used in MAC protocols for ad hoc networking[3][9][11][22][25], the popularity and prevalence of 802.11 make it difficult to ignore. This has led to numerous attempts to improving the performance of 802.11-based ad hoc networks.

Instead of addressing the MAC itself, many have focused on adapting the other protocols for use in an ad hoc environment. For example, TCP's congestion controls reacted improperly to the route failures initiated by the MAC protocol. Several works fixed this problem in part by creating new versions of TCP that did not interpret all lost or delayed packets as an indication of congestion[8][16].

Several approaches have attempted to deliver Quality-of-Service (QoS) in an ad hoc network by providing some type of priority service at the MAC-level. This allows the differentiation of services between different nodes or traffic types, allowing medium access to be shared unequally, based on need. Several approaches have been proposed, including varying the length of data packets, changing the interframe spacing, and adjusting the contention window[1][15].

Finally, a few works have attempted to improve performance by acknowledging the value of an established channel. Considerable effort is required in order for a source to acquire the channel. Therefore, it should be used as effectively as possible before it is released again. While holding the channel indefinitely is also undesirable, sometimes using the medium for only a single packet is a waste of effort.

Instead, the channel can be re-used, in order to send multiple packets, up to a certain total length. Once a channel has been acquired and a connection determined to be strong, a sender may continue to transmit a subsequent contention-free burst of packets, as seen in Figure 4[19][20]. A similar approach sends packets consecutively, then collects positive and negative acknowledgements from each recipient[21]. Both these approaches must balance the efficiency gains of re-using the medium without contention, with the danger of preventing other nodes from accessing the medium for a prolonged length of time.

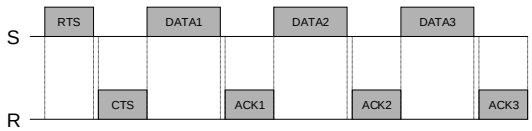

Fig. 4
SENDING MULTIPLE PACKETS

III. Ad hoc MAC Requirements

In developing a MAC protocol for ad hoc networks, the characteristic properties of ad hoc networks must be considered. From these characteristics, specific requirements arise. In this section, four key requirements for an ad hoc MAC are identified, and the concept of our approach is introduced.

A. Minimize protocol overhead

This is generally true in any network. However, in an ad hoc network, it takes on particular importance due to the scarcity of bandwidth. A slight increase in overhead consumes a proportionately higher percentage of available resources.

The 802.11 protocol includes a considerable amount of overhead within the packet exchange. The RTS, CTS, and ACK frames must all be transmitted, as well as the MAC header on the DATA packet. Additionally, the medium must remain idle during both the DIFS before an exchange can commence, and the SIFS separating consecutive frames. Within an ad hoc network, this overhead is multiplied. While the packet only counts once towards network throughput, multihop transmission incurs a multiple overhead penalty.

B. Reduce contention and collisions

A trade-off exists between longer contention periods and increased frequency of collisions. Either option is undesirable, each resulting in wasted network bandwidth. Therefore, the preferred solution is to avoid contention whenever possible. However, allowing one node to monopolize the channel is also detrimental to the proper operation of the network. Therefore, while contention and collisions should be avoided, it cannot come at the expense of fair access.

C. Consider fairness in ad hoc networks

The concept of fairness in an ad hoc network is not necessarily about providing equal access to the medium[24]. Transmitting a packet over the wireless interface does not in itself provide the node with any reward. In fact, forwarding packets for other nodes has an associated cost, not benefit.

If a node is forced to forward a large amount of traffic, it must also be able to use the channel frequently, in order to fulfil its responsibilities to the network. If it has the same shared access to the channel as a node that is entirely sending its own packets, not only will the forwarding node be disadvantaged, but the entire network may be ineffective.

D. Directly support ad hoc routing and forwarding

The fundamental characteristic of an ad hoc network is the dynamic multihop topology. This makes the routing protocol critical to network effectiveness. While it has been shown that the MAC can negatively impact the routing protocol[4], perhaps if designed properly, the MAC could also support and enhance routing. In order to do this, we propose to make the MAC aware of the forwarding process. By providing mechanisms to swiftly and efficiently forward packets along existing routes, neighbours remain reachable, links remain intact, and the route is maintained for as long as possible.

IV. Forward Focus Mechanism and Protocol Description

A. Mechanism

Following the transmission of an acknowledgement (thereby completing the packet exchange), in 802.11 the receiver then releases the channel to the next contention phase. It is explicitly stated in the specification that this must be done – neither the sender nor the receiver should ever maintain control of the channel. This is to prevent a sender or receiver from abusing the medium, by remaining in control of the medium, excluding all others. However, the protocol was designed primarily for use in wireless LANs, not ad hoc networks.

An ad hoc network is inherently cooperative, at least to some extent. The multihop nature of the network demands that, in order for the network to work effectively, nodes must forward traffic for other nodes. Therefore, although abuse is still a concern, a node using the medium to forward a packet is in fact only benefiting the network, drawing no direct reward for itself.

By utilizing the forwarding nature of the network, the receiving node can safely be allowed to re-use the medium. If the receiving node is serving as an intermediate node, it is deriving no benefit from receiving and subsequently re-transmitting the packet. Therefore, it can be allowed to immediately re-use the channel in order to transmit a packet of its own. In essence, the node is rewarded for receiving the packet to be forwarded, by allowing immediate access to the medium. We call this mechanism Forward Focus (FF), as it both utilizes and encourages the forwarding nature of an ad hoc network.

Note that this does not actually force the node to ever forward the packet it received. A misbehaving node could receive packets to be forwarded, drop them, and use the medium solely for its own traffic. Other mechanisms are required to ensure this does not happen.

However, if this occurs repeatedly, the route will break, and a new path will be formed, likely excluding this node.

This mechanism has several advantages. First, when a packet requires multihop delivery, it provides the opportunity for subsequent transmissions to be attempted in a contention-free manner. This reduces the likelihood of collisions. Second, because it eliminates the DIFS and contention period, it is in fact more efficient than 802.11. The channel is simply idle less of the time. Third, the mechanism considers the forwarding nature of the network, and encourages it. By reducing the burden on intermediate nodes, the network should be able to operate more effectively.

B. Protocol

Using the FF mechanism described in the previous section, we have designed a protocol that remains as close as possible to the MAC protocol in the 802.11 specification. It is identical to 802.11 in most respects. However, after receiving the data packet, the receiving node may sometimes acknowledge the packet, then maintain control of the channel. Rather than releasing the channel for a new contention period, the receiving node may then, if conditions are appropriate, attempt to send a packet of its own.

After receiving the data packet, the MAC must first check the network destination of the packet, by checking the destination IP address. This information is not typically used at the MAC-layer, however it is fairly readily attainable if required. Based on this IP address, the node can decide whether or not it is the destination for the packet. If it is, the protocol proceeds as per the 802.11 specification. However, if it is not, the node is an intermediate node, and this packet must be forwarded. The packet is passed up the protocol stack, as it would be normally, however the MAC is now granted permission to immediately reuse the channel.

The node must first wait the required SIFS (which must separate all consecutive MAC frames). Following this period, the MAC protocol may attempt to send the next packet available in the queue. If there is no packet to be sent, the channel is released normally. However, depending on the processing time at the node, there should be at a minimum one packet to be sent (the one just received).

The receiver then becomes a sender, attempting to initiate a new packet exchange. Depending on the packet to be sent, the intermediate node can either transmit it immediately, or initiate the RTS/CTS procedure. Both scenarios are shown in Figures 5 and 6. After the exchange has been initiated, it proceeds normally according to the 802.11 specification.

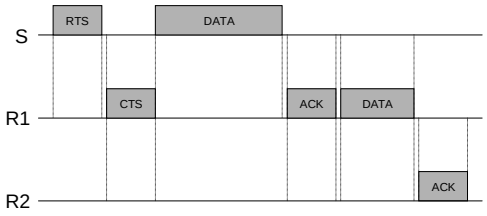

Fig. 5
FF SEQUENCE WITHOUT RTS-CTS

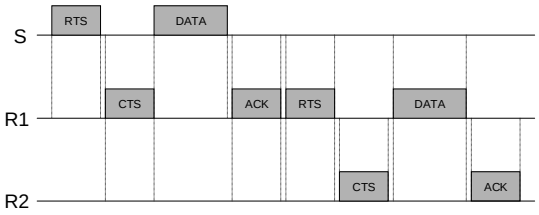

Fig. 6
FF SEQUENCE WITH RTS-CTS

V. PERFORMANCE EVALUATION

In order to evaluate the FF-based protocol, a simulation was constructed in the ns-2 simulator[27]. The results were compared with the results yielded by the ns-2 implementation of the 802.11 MAC. AODV[5][17] was used as the routing protocol.

Two different network scenarios have been used in this evaluation. The first is a linear topology network. This has been used as a proof of concept, in order to show the validity and potential effectiveness of the approach. The second scenario uses randomly generated topologies, and multiple flows, in order to evaluate the protocol in a more realistic situation.

A. Linear topology

For the purpose of evaluating the effect of FF on a single multihop flow, simple linear networks were

simulated. The lines of nodes ranged from 2 nodes (a direct link), to 10 nodes (requiring 9 hops). All of the nodes were spaced 200 metres apart, a value that ensured that nodes could reach the adjacent node or nodes along the line, but not the following node. Therefore, no node could reach more than two other nodes, with the end nodes only reachable by one other node. The two end nodes were used as source and destination, as shown in Figure 7.

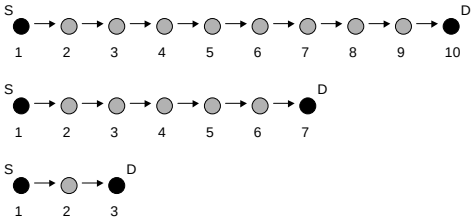

Fig. 7
3, 7, AND 10 NODE LINEAR TOPOLOGIES

This topology was tested first using a constant bit-rate (CBR) traffic source. It was configured to generate 512 byte (4 kb) packets at regular intervals. This interval was adjusted in order to vary the flow's offered load. The source was attached to node 1 and the sink to node n. UDP was used to deliver the packets from source to destination.

Figure 8 and 9 show the average throughput achieved as a function of path length for two offered loads. Note that both offered loads could be supported over a direct link (path length of 2 nodes), however as path length increases, the average throughput begins to decrease. For both offered load levels, the FF protocol gains a clear advantage at longer path lengths. In the low offered load case, FF successfully maintains a packet delivery ratio of over 90% throughout, while 802.11 falls below 90% at 5 hops, below 60% at 7 hops, and below 50% at 9 hops. By the 10-node path, throughput for 802.11 has fallen to approximately half that of FF. The higher offered load is well above what the flow can support, however the FF protocol again sustains a throughput well above that of 802.11.

Figure 10 shows the average throughput as function of offered load. In this figure, it is clear that the throughput rises as offered load increases, until the maximum capacity of the flow is reached. While there is only a slight

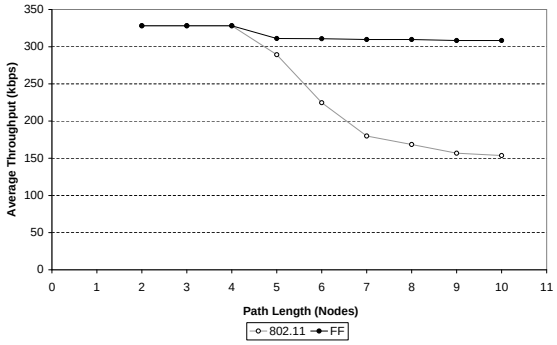

Fig. 8
THE EFFECT OF PATH LENGTH ON THROUGHPUT: 328 KBPS OFFERED LOAD

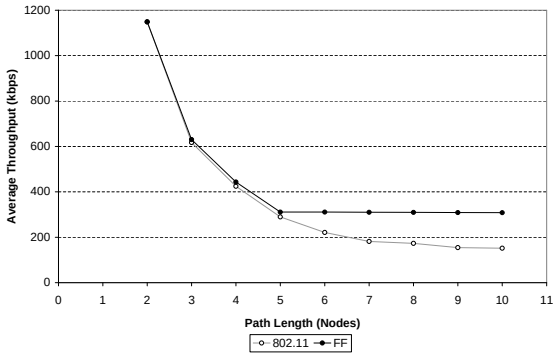

Fig. 9
THE EFFECT OF PATH LENGTH ON THROUGHPUT: 1148 KBPS OFFERED LOAD

difference in this maximum capacity for short paths (although FF still enjoys a small advantage), FF outpaces 802.11 by a considerable margin over the longer path.

For the shorter path, FF gains its slight advantage by being more efficient between transmissions. A small reduction in idle time between transmissions (a DIFS to a SIFS) allows a modestly increased throughput. However, the larger increase over the longer paths cannot be explained by this alone. Other effects must be involved to create such a significant difference.

Figures 11 and 12 illustrate the true reason for this improvement. Throughput in the 802.11-based network is extremely erratic, while the FF network sustains a remarkably steady flow. Within the 802.11 throughput trace (of a single test run), a number of periods are

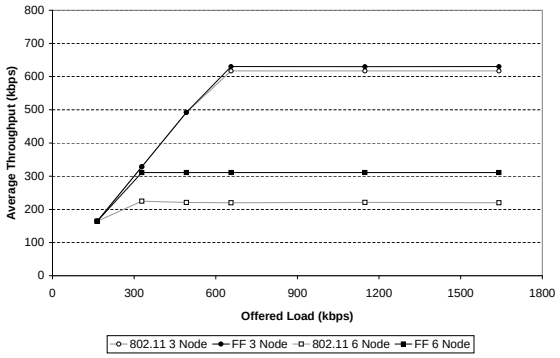

Fig. 10
THE EFFECT OF OFFERED LOAD ON THROUGHPUT

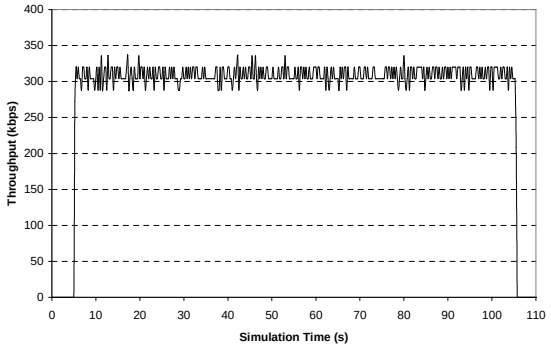

Fig. 12
THROUGHPUT TRACE FOR FF (310 KBPS AVERAGE)

visible where the throughput goes to zero, such as at t = 32 seconds. Frequently, these periods indicate route failures, during which time packets are delayed or dropped while a new route is discovered. The FF protocol avoids these failures, maintaining a steady flow throughout.

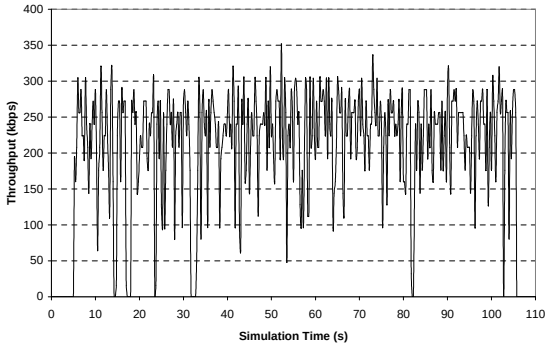

Fig. 11
THROUGHPUT TRACE FOR IEEE 802.11 (225 KBPS AVERAGE)

The linear topology was also used to evaluate the performance of the FF protocol for a bi-directional traffic flow. For this scenario, an FTP transfer was initiated over a TCP connection. This requires both the flow of packets from source to destination to be sustained, but also a reverse flow of acknowledgement packets, back to the source. The simulation attempted to transfer as much data as possible, within 100 seconds of simulated time.

The average throughput of the file transfer is shown in Figure 13. Again, there is the expected decrease in throughput as the path lengthens. The relative performance improvement of FF over 802.11 is shown in Figure 14. Although the improvement is not nearly as dramatic as for the CBR traffic, FF does gain up to a 5% improvement for long paths.

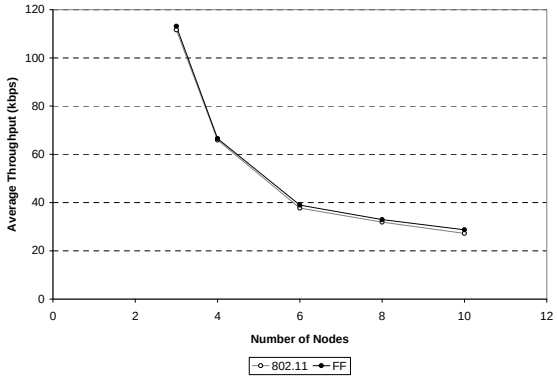

Fig. 13
AVERAGE THROUGHPUT FOR FF AND 802.11 FILE TRANSFERS

The smaller improvement can be explained as follows: in the CBR case, traffic is unidirectional, therefore the FF mechanism can move packets along as quickly as possible. As this is the only traffic in the system, packets tend to be forwarded immediately, and another packet is started along the chain as soon as the first one is out of range. However, for bi-directional traffic the forwards and reverse flows actually interfere with each other. Although the network still sees an improvement over 802.11, the flow is far less streamlined than when traffic flows in a single direction.

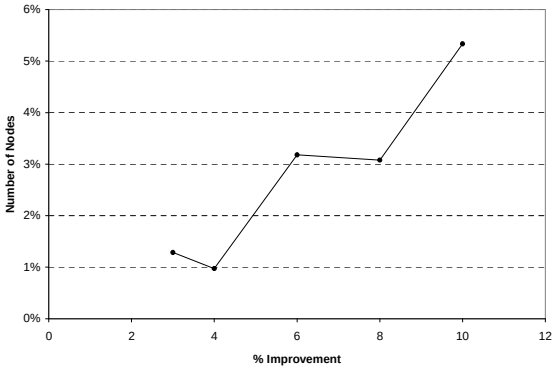

Fig. 14
% PERFORMANCE INCREASE OF FF OVER 802.11 FOR FILE TRANSFER)

B. Random network topology

This raises the question of what happens when there are multiple flows competing in different directions. In order to evaluate this scenario, simulations were performed using a static random topology. Nodes were distributed randomly over a 1000 metre square network area. For this simulation, 50 nodes were used. This configuration should typically yield an average path length of between 3 and 4 hops, with the longest paths occasionally reaching up to 10 hops. In order to avoid scenarios where nodes were unreachable, any disconnected network topologies were discarded.

For each run, 10 traffic streams were created, again using 512-byte CBR packets, transported by UDP. The 10 flows were started at one-tenth of a second intervals, in order to avoid overwhelming the network with a large burst of routing packets caused by simultaneous route discoveries. Each flow lasted exactly 60 seconds in duration.

In order to avoid overloading a particular source or destination node (thereby skewing the results), sources and destinations were paired 1 and 11, 2 and 12, and so on, up to 10 and 20. As the node locations were chosen randomly, this identification of nodes is in fact arbitrary. 5 different topologies were generated and combined with 2 seed values to provide 10 runs. The 10 data values were subsequently averaged.

Figure 15 illustrates the total network throughput. As load increases, FF gains up to a 15% increase over 802.11. Despite this increase in throughput, the packet delivery rate quickly falls below 90% all the way down to a mere 30% at the highest load, as seen in Figure 16. Interestingly, throughput continues to increase, due to the fact that some paths still have available capacity. These paths are likely either short (possibly adjacent), and/or isolated and not in direct competition for bandwidth with other paths.

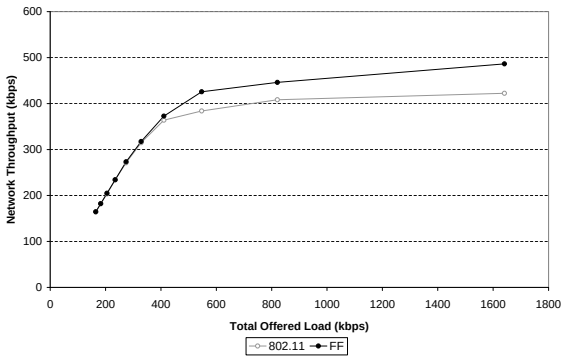

Fig. 15
THROUGHPUT IN A RANDOM NETWORK

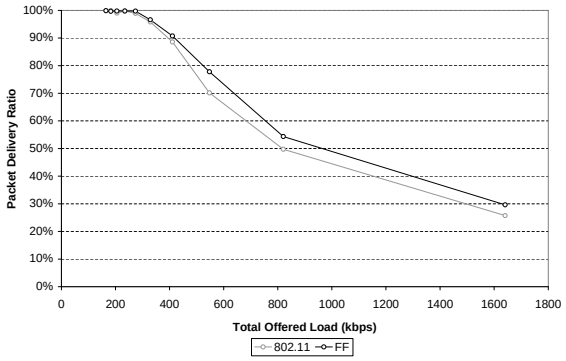

Fig. 16
PACKET DELIVERY RATIO IN A RANDOM NETWORK

Why doesn't FF outperform 802.11 in the random network by as wide a margin as in the linear case? First, the path lengths are shorter, and the individual offered loads are lower. Even in the linear case, FF and 802.11 were relatively close under those conditions. Second, there is increased contention in the longer paths, where FF would have an advantage. Therefore, the shorter paths comprise more of the traffic, reducing the potential gains of the modified protocol.

VI. Conclusions

This paper has focused on developing a MAC protocol to directly support ad hoc routing and forwarding. While the MAC protocol's effects on routing have been well documented, this work has investigated whether that the forwarding nature of the network can be used to our advantage. Using this concept, a simple mechanism for improving the IEEE 802.11 MAC was identified. This mechanism focuses on forwarding packets, reducing the cost of performing these ad hoc-specific duties. It works by allowing intermediate nodes to immediately re-use the medium in order to send packets of their own.

The intermediate node is able to attempt its packet transmission before any other node is allowed to access the medium. This provides the node with contention-free medium access, with very little chance of collision. It also requires less idle time than typical 802.11 packet exchanges.

Simulations were performed using an implementation of the protocol in the ns-2 simulator. A linear network topology allowed the study of the protocol's effect on a single flow. Both CBR and FTP generated traffic were used. For CBR (unidirectional) traffic, FF showed remarkable throughput increases over 802.11, particularly at longer path lengths and higher loads. This is where the advantages of the FF protocol really stood out. For FTP (bidirectional) traffic, FF still outperformed 802.11, but by a much smaller margin. At longer path lengths, FF's throughput advantage grew to approximately 5%.

A random topology scenario was also generated. Multiple traffic sources were used in order to simulate a more realistic ad hoc network. Again, FF's throughput surpassed 802.11, although it occurred in a range where packet delivery range was rather low. However, improvements of up to 15% were achieved.

Much of this gain was achieved by finding a mechanism that removed some of the instability experienced at the MAC-level. Even with the network static, 802.11 experienced link and route failures. By eliminating these unnecessary breakages, FF eliminated much of the volatility from the overall data flow. This relatively simple mechanism utilized the properties of the ad hoc network that have otherwise been treated as disadvantages.

FF also demonstrates that utilizing even a little bit of extra information can prove beneficial to the network. By allowing the MAC protocol access to even small amounts of routing information (in this case, the IP destination of the packet), a considerable gain was produced. The additional information allowed the MAC to assign access to the medium in a manner that better suited the traffic patterns. By doing so, it was able to provide better service to the routing protocol. It is possible that with further interaction or integration of the MAC and routing protocols, further improvements can be made.

References

[1] I. Aad, and C. Castelluccia, "Differentiation Mechanisms for IEEE 802.11," *Proc. INFOCOM 2001*, Apr. 2001.

[2] G. Anastasi and L. Lenzini, "QoS Provided by the IEEE 802.11 wireless LAN to advanced data applications: a simulation analysis," *Wireless Networks 6*, Baltzer, 2000.

[3] L. Bao and J.J. Garcia-Luna-Aceves, "A New Approach to Channel Access Scheduling for Ad Hoc Networks," *ACM MOBICOM'01*, 2001.

[4] C. Barrett, M. Drozda, A. Marathe, and M.V. Marathe, "Characterizing the Interaction Between Routing and MAC Protocols in Ad-Hoc Networks," *Proc. ACM Mobihoc'02*, June 2002.

[5] E.M. Belding-Royer and C.E. Perkins, "Evolution and future directions of the ad hoc on-demand distance-vector routing protocol," *Ad Hoc Networks*, vol. 1, no. 1, Elsevier, July 2003.

[6] B. Bensaou, Y. Wang, and C.C. Ko, "Fair Media Access in 802.11 Based Wireless Ad-hoc Networks," *Proc. IEEE/ACM Mobihoc 2000*, Aug. 2000.

[7] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE JSAC*, 2000.

[8] K. Chandran, S. Raghunathan, S. Venkatesan, and R. Prakash, "A Feedback-Based Scheme for Improving TCP Performance in Ad Hoc Wireless Networks," *IEEE Personal Communications*, Feb. 2001.

[9] J.J. Garcia-Luna-Aceves and A. Tzamaloukas, "Reversing the Collision-Avoidance Handshake in Wireless Networks," *Mobile Computing and Networking*, 1999.

[10] M. Gerla, K. Tang, and R. Bagrodia, "TCP Performance in Wireless Multihop Networks," *Proc. IEEE WMCSA '99*, Feb. 1999.

[11] Z.J. Haas and J. Deng, "Dual busy tone multiple access (DBTMA) — a multiple access control scheme for ad hoc networks," *IEEE Trans. on Communications*, June 2002.

[12] Z.J. Haas, J. Deng, and S. Tabrizi, "Collision-free Medium Access Control Scheme for Ad Hoc Networks," *Proc. IEEE MILCOM*, 1999.

[13] V. Kanodia, C. Li, A. Sabharwal, B. Sadeghi, and E. Knightly, "Distributed priority scheduling and medium access in ad hoc networks," *Wireless Networks*, vol. 8, no. 5, Kluwer, 2002.

[14] J. Li, C. Blake, D. Couto, H. Lee, and R. Morris, "Capacity of ad hoc wireless networks," *Proc. ACM MobiCom'01*, July 2001.

[15] A. Lindgren, A. Almquist, and O. Schel, "Quality of service schemes for IEEE 802.11 wireless LANs: an evaluation," *Mobile Network Applications*, vol. 8, no. 3, Kluwer, 2003.

[16] J. Liu and S. Singh, "ATCP: TCP for Mobile Ad Hoc Networks," *IEEE JSAC*, July 2001.

[17] C. Perkins, E. Belding-Royer, and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing," IETF MANET WG RFC 3561, July 2003.

[18] E.M. Royer, S.-J. Lee, and C.E. Perkins, "The effects of MAC protocols on ad hoc network communication," *IEEE WCNC 2000*, Sept. 2000.

[19] B. Sadeghi, V. Kanodia, A. Subharwal, and E. Knightly, "Opportunistic Media Access for Multirate Ad Hoc Networks," *Proc. ACM MOBICOM'02*, Sept. 2002.

[20] S. Choi, J. del Prado, S. Shankar, and S. Mangold, "IEEE 802.11e Contention-Based Channel Access (EDCF) Performance Evaluation," *Proc. IEEE ICC '03*, May 2003.

[21] S.-T. Sheu, T. Chen, J. Chen, and F. Ye, "An Improved Data Flushing MAC Protocol for IEEE 802.11 Wireless Ad Hoc Network," *Proc. IEEE VTC '02*, Sept. 2002.

[22] F. Talucci, M. Gerla, and L. Fratta, "MACA-BI (MACA by invitation)-A Receiver Oriented Access Protocol for Wireless Multihop Networks," *Proc. IEEE PIMRC '97*, 1997.

[23] K. Tang, M. Correa, and M. Gerla, "Effects of Ad Hoc MAC Layer Medium Access Mechanisms under TCP," *Mobile Networks and Applications 6*, Kluwer, 2001.

[24] K. Tang and M. Gerla, "Fair Sharing of MAC under TCP in Wireless Ad-hoc Networks," *Proc. IEEE MMT '99*, Oct. 1999.

[25] 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 JSAC*, Sept. 2000.

[26] 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.

[27] ns-2 Network Simulator, The VINT Project. http://www.isi.edu/nsnam/ns/

[28] Wireless LAN Media Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE Standard 802.11, June 1999.

[29] Part 11: Wireless LAN Medium Access Control(MAC) and Physical Layer (PHY) Specifications: Higher-Speed Physical Layer Extension in the 2.4 GHz Band, IEEE Standard 802.11b-1999, Sept. 1999.