MAC Support for Broadcast-based Ad-hoc Forwarding Schemes

A. Rahman and P. Gburzynski

University of Alberta
Department of Computing Science
Edmonton, Alberta CANADA T6G 2E8

Abstract – **As a collision avoidance scheme, IEEE 802.11 performs poorly in broadcast applications because its RTS/CTS-based virtual carrier sense mechanism is specifically tailored to point-to-point communication. The situation is aggravated by the lack of acknowledgments, which are difficult to implement in a broadcast environment. Consequently, most efforts in ad-hoc routing have focused on establishing explicit paths from source to destination using dynamic routing tables and point-to-point hops. We postulate a mechanism, dubbed fuzzy acknowledgments, to address the problem of reliability of broadcast communication in IEEE 802, which renders the scheme useful for broadcast-based ad-hoc routing protocols (akin to controlled flooding). We believe that such protocols may be useful for some classes of applications and that, in terms of their performance, they need not significantly yield to more complex schemes based on point-to-point transmission. This belief is supported by our recently proposed ad-hoc routing protocol, called TARP, which also benefits from the idea of fuzzy acknowledgments introduced in this paper.**

I. Introduction

Routing protocols for ad-hoc wireless networks can be broadly divided into proactive, i.e., ones that maintain up to date routing information at every node [1,2], and reactive that collect the necessary information only when it is useful for communication [3,4]. Recently, we proposed a simple reactive ad-hoc routing protocol, dubbed TARP[1] [5], in which routing is inherently broadcast-based. This means that intermediate stations do not care about the identity of their neighbors. Routing paths in TARP are built with the assistance of scalable heuristics that base their decisions solely on the identity of the source and destination and the perceived number of hops traveled by a packet. Depending on the amount of resources available at intermediate nodes, as well as on their mobility, these heuristics exhibit a better or worse convergence to the shortest path between source and destination.

In its original version introduced in [5], TARP visibly yielded in terms of performance to its more complex competitors. The primary reason for the protocol's deficiency was the inadequacy of the MAC scheme based on IEEE 802.11. In TARP, the role of a single forwarding hop is to pass the packet to *any* subset of neighbors who consider themselves "eligible" to take care of its next hop. The sender does not care about the identity of the neighbor that picks up the packet, but it would like to make sure that the packet has been picked up by at least one eligible neighbor. This *modus operandi* makes TARP different from other mainstream routing approaches and can be viewed as a defining property of a potentially large class of ad-hoc routing schemes, with TARP being one representative.

1 Tiny Ad-hoc Routing Protocol.

Note that the RTS/CTS/Data/ACK handshake of IEEE 802.11 is useless in the above variant of hot-potato routing. Its only sensible application would be to carry out a separate exchange of the same packet with each eligible neighbor. This, in addition to wasting bandwidth, would require a prior identification of those neighbors and would contradict the fundamental principle of TARP. With the traditional collision avoidance scheme of IEEE 802.11, the forwarding node has no means of determining or even guessing at the success of its broadcast transmission in the data-link layer. Consequently, the decision to retransmit the packet can only be made in higher layers based on some nebulous criteria, e.g., session timers, which results in increased delays and poor performance.

II. Defciencies of IEEE 802.11

The IEEE 802.11 family of MAC schemes is well equipped to handle point-to-point communication in the face of multiple parties trying to talk at about the same time [6]. With the four-way handshake, RTS/CTS/DATA/ACK, the sending node first verifies that the other party is present and willing to receive, and reserves bandwidth for the actual data exchange, and then, following the packet transmission, it receives an acknowledgment from the recipient. The RTS/CTS part of the handshake also accounts for the "hidden terminal" problem by including the recipient in the bandwidth reservation part of the complete exchange.

Unfortunately, none of these features is available for broadcast transmission, particularly in its flavor needed in TARP. First, the RTS/CTS part makes no sense because 1) the recipient's identity is unknown and unimportant, 2) there can be multiple legitimate recipients that do not know about each other. Second, even the two-way handshake, DATA/ACK, is not possible because of 2. Consequently, the only available option is to transmit blindly, following the standard DIFS delay and back-off procedure prescribed by the scheme. This completely ignores the hidden terminal problem, greatly increases the likelihood of a collision, and renders the data exchange highly unreliable.

III. Fuzzy Acknowledgments

We propose the following simple solution as an extension of the IEEE 802.11 MAC protocol. When a node receives a packet for which it considers itself eligible (i.e., the packet should be forwarded), it waits for a short amount of time, defined by the short inter-frame space (SIFS), and then sends an acknowledgment. When multiple recipients send their acknowledgments at (almost) the same time, the sender will not be able to recognize them as valid packets. However, the sender can interpret any activity (of a certain bounded duration) that follows the end if its last transmitted packet as an indication that the packet has been successfully forwarded. Although the value of this indication is inferior to that of a "true" acknowledgment, it may provide the kind of feedback needed by the data-link layer to assume that its responsibility for handling the packet has been fulfilled.

Thus, having completed a packet transmission, the sender immediately switches to listening mode and awaits a period of silence (of duration comparable to SIFS) followed by a burst of activity (of duration comparable to the duration of an acknowledgment packet). If such an event occurs within the prescribed interval, the sender assumes that the packet has been passed over; otherwise, it schedules a retransmission. With this approach, the acknowledgment packet (which carries no information other than its presence) can be made very short and consist of some characteristic pattern unlikely to be encountered in a regular packet.

IV. Incorporation into the Protocol

TARP (see [5] for a detailed description), is based on several heuristics (called rules) that drive path convergence in the face of the dynamic reconfiguration of the network. As the role of a fuzzy acknowledgment in TARP is to tell the sender that its packet has been forwarded towards the destination, it is important that only those recipients that are actually going to forward the packet (or the destination itself) send the acknowledgments. Consequently, acknowledgments cannot be sent mechanically in the data-link layer, and the incorporation of fuzzy acknowledgments into TARP requires some cross-layer coordination. For illustration, consider the configuration of nodes shown in Figure 1. Assume for simplicity that the network is static and that there is a session in progress between nodes 1 and 2, with the current converged path passing through nodes <1,10,15,20,22,2>. When node 1 sends a packet in the first hop, it will be received by nodes 6,7,10,12, and 13.

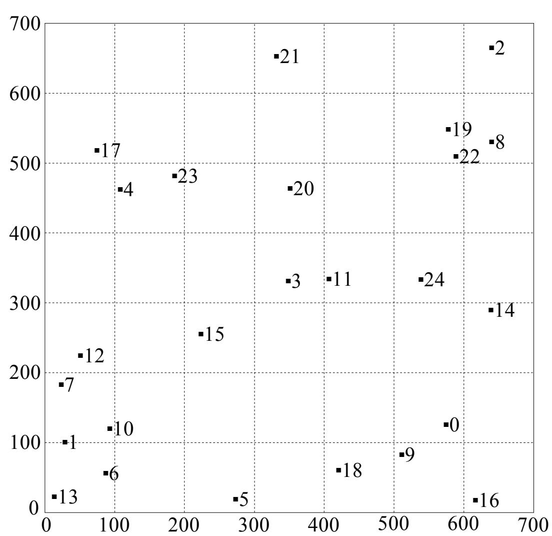

Fig. 1. A sample configuration of nodes.

In the next hop, node 10 is going to rebroadcast the packet because it lies on the converged path. But what will happen if the packet is received by some neighbors of node 1 but not by node 10, e.g., because of an interference. If any of those neighbors sends an acknowledgment that is subsequently received by node 1, then node 1 will conclude that the packet has been forwarded, which, of course, is not the case.

To see another problem, suppose that a packet sent from node 1 to node 2 has arrived at node 4, which, according to its current perception of path convergence, considers itself eligible. The recipients of the transmission of node 4 are nodes 17 and 23. If node 17 decides to forward this packet again, it will arrive back at 4 and 23, where it will be recognized as a duplicate. Thus, neither of the two nodes will find itself eligible and they will send no acknowledgments. Consequently, node 17 will keep retransmitting the packet over and over until it finally decides that the optimal path for the session lies elsewhere. This will create unnecessary activity in the neighborhood of node 17 contributing to the overall interference and reducing the amount of usable bandwidth.

By blurring the distinction between layers a bit further, we can avoid this problem and turn it into one more heuristic facilitating path convergence in TARP. Consider three nodes: the source *S*, the destination *D*, and an intermediate node *N*. Suppose that Δ refers to a time interval and define

$$RF_{SND}(\Delta)=n^{ack}_{SND}(\Delta)/n^{fwd}_{SND}(\Delta) \quad ,$$

where n^{fwd}_{SND} is the total number of packets between *S* and *D* passed through *N* within time Δ, and $n^{ack}_{SND} \le n^{fwd}_{SND}$ represents the portion of those packets for which *N* has received (fuzzy) acknowledgments. Depending on the setting of the interval Δ, *RF* (called the *relevance factor*) can be viewed as a measure of *N*'s relevance in forwarding packets between *S* and *D*. Alternatively, we can view *RF* as a measure of probability that *N* lies on the optimal path between the two end-nodes, at least as long as the configuration of nodes remains static.

With the mobility included, the indications of *RF* become less accurate. This is not solely a problem of TARP: any information related to the configuration of paths tends to become outdated, if the nodes are allowed to change their location. Thus, the proper way to interpret the value of RF should be determined experimentally. One natural idea is to use a threshold. A value of *RF* above the threshold indicates that the node is relevant in sustaining the session.

A straightforward way to implement the new rule is to take advantage of the DD cache and flag those packets for which acknowledgments have been received with one extra bit. This approach automatically equates Δ with the DD cache expiration interval and rids the protocol of one parameter.

V. The Improvement

Figure 2 illustrates how the relevance factor *RF* affects the performance of TARP in terms of the packet delivery fraction. It shows that *RF* does influence the quality of routing and hints at the range between 0.6 and 0.75 as the suggested setting. Notably, the same range of values seems to be adequate for different mobility levels, which allows us to make *RF* a constant rather than a dynamically tunable parameter.

In Figure 3, we show how much improvement has been brought into TARP by the addition of the new rule. The upper

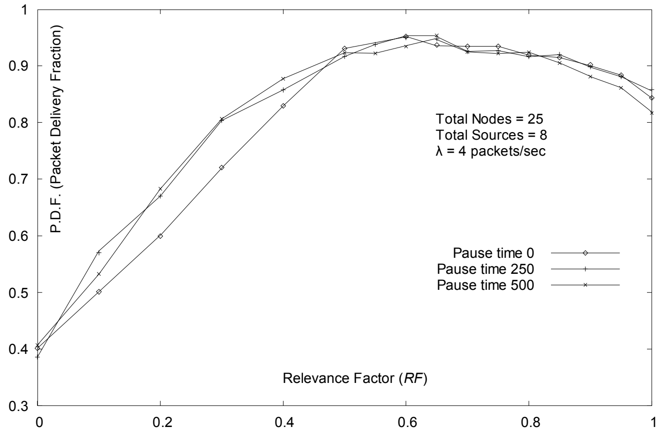

Fig. 2. Impact of RF threshold on packet delivery fraction.

portion presents a magnification of the two TARP curves from Figure with the inclusion of a new curve reflecting the fuzzy-acknowledgment version with *RF* = 0.7. The magnitude of the observed improvement has been consistent across different node densities and traffic levels. The bottom portion of Figure 3 compares the three variants of TARP for a larger number of sessions.

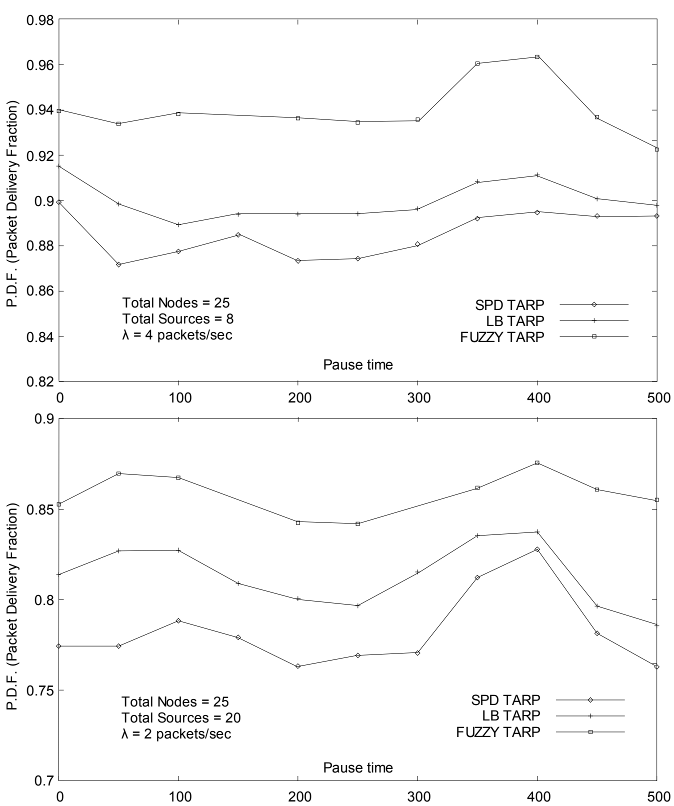

Fig. 3. The magnitude of performance improvement with the addition of fuzzy acknowledgments.

VI. Future Work

Our current work on TARP is focused on three issues. First, we would like to eliminate the detrimental diversity of alternative paths with the same shortest length, which is the primary source of resource wastage in the present version of TARP. As it turns out, there is a way to provide the requisite feedback to a forwarding node without affecting the spirit of the protocol

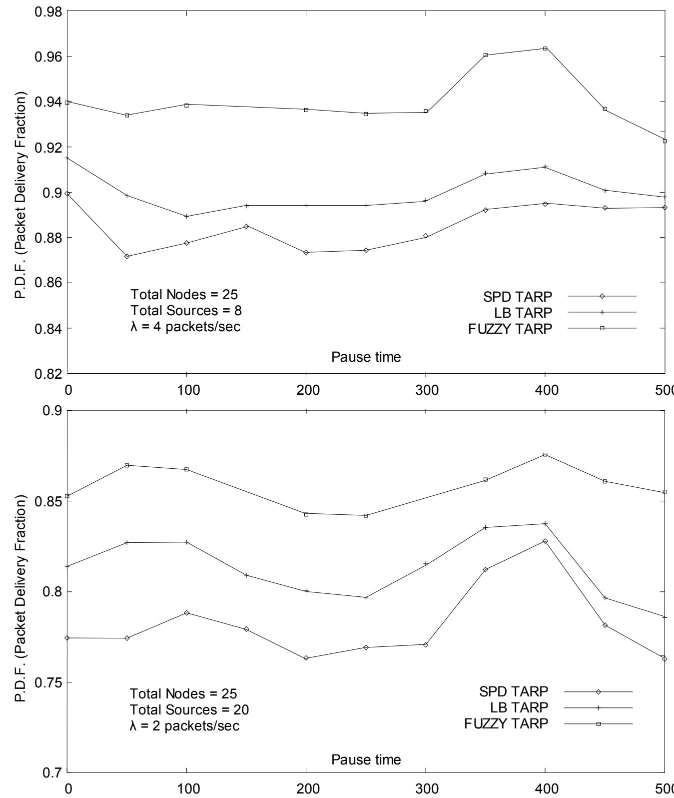

Fig. 4. The magnitude of performance improvement with the addition of fuzzy acknowledgments.

and giving up its underlying forwarding paradigm. We are currently experimenting with the *Fourth Rule* aimed at selecting one of the multiple shortest paths – the one that provides locally best service, as viewed by an eligible node.

One promising avenue is to try a fuzzy variant of the RTS/CTS handshake, whereby the forwarding node announces its intentions with a brief RTS-like packet that contains, instead of the recipient address, the signature of the session packet about to be forwarded. In response to this packet, all eligible nodes would send a *fuzzy* CTS response. Intuitively, this solution may work better than the fuzzy acknowledgments discussed in this paper because it would also account for the hidden terminal problem.

References

[1] Chen, T-W. and Gerla, M. "Global State Routing: a New Routing Scheme for Ad-hoc Wireless Networks." Proceedings of ICC'98, June, 1998.

[2] Perkins, C.E. and Bhagwat, P. "Highly dynamic Destination-Sequenced Distance Vector Routing (DSDV)," Proceedings of SIGCOMM'94, pp. 234-244, August, 1993.

[3] Perkins, C.E. and Royer, E.M. "Ad-hoc On-demand Distance Vector Routing (AODV)," Proceedings of WMCSA'99, pp. 90-100, 1999.

[4] Perkins, C.E., Royers, E.M., and Das, S.R. "Ad-hoc On-demand Distance Vector Routing (AODV)," Internet Draft: draft-ietf-manet-aodv-13.txt, IETF, February, 2003.

[5] Olesinski, W., Rahman, A., and Gburzynski, P. "TARP: a Tiny Ad-hoc Routing Protocol for Wireless Networks," Proceeding of ANTAC'03, Melbourne, Australia, December, 2003.

[6] Weinmiller, J., Schlager, M., Festag, A., and Wolisz, A., "Performance Study of Access Control in Wireless LANs IEEE 802.11 DFWMAC and ETSI RES 10 HIPERLAN," Mobile Networks and Applications, vol. 2, pp. 55-67, 1997.