Multicasting vs. Broadcasting: What are the Trade-offs?

Bora Karaoglu, Wendi Heinzelman
Department of Electrical and Computer Engineering
University of Rochester
Rochester, NY 14620
Email: {karaoglu, wheinzel}@ece.rochester.edu

***Abstract*—Network-wide broadcasting and multicasting are two important routing schemes used in group communications. In network-wide broadcasting, generated packets at the source node are distributed to all nodes in the network, while multicasting delivers the packets only to a subset of the nodes by creating and maintaining a data dissemination structure. Due to the overhead in multicasting, in certain situations, it is more efficient to use network-wide broadcasting instead of multicasting, even when the data is destined to a subset of the nodes. In this paper we analyze a specific broadcasting protocol and a specific multicasting protocol for two performance metrics, namely spectrum efficiency and energy efficiency, and determine the conditions that make one of them preferable over the other.**

I. Introduction

Group communications is essential for many applications in mobile ad hoc networks, including supporting electronic classrooms, tactical military communication, and communication in disaster recovery missions. One to many group communications are generally classified into two groups: network-wide broadcasting and multicasting. In network-wide broadcasting the objective is to distribute the generated data to all the nodes in the network. On the other hand, the objective of multicasting is to deliver the data to a subset of the nodes in the network. By using a data dissemination structure, multicasting protocols limit the diffusion of the data to a certain subset of the entire network, namely to the multicast members. [1] and [2] present recent surveys on various group communication protocols.

There is additional overhead incurred in multicasting protocols compared to broadcasting protocols. In certain scenarios, the cost of collecting and processing the additional information overwhelms the gains in limiting the data dissemination structure to the multicast members. As one might expect, in scenarios where the majority of the nodes are part of the multicast group, one can increase the efficiency by using a broadcasting protocol instead of using a multicasting protocol. In this paper, our objective is to investigate the trade-offs between multicasting and broadcasting in order to determine the conditions that make one of them preferable over the other.

There are similar studies that point out the trade-offs between multicasting and broadcasting. Researchers in [3] compare a selected multicast protocol, namely on demand multicast routing protocol (ODMRP), with a selected network-wide broadcast protocol, namely scalable broadcast algorithm (SBA). The authors report that while multicasting is preferable for small group sizes, as the group size increases, broadcasting becomes more efficient. However, the protocols considered in this work are quite different in nature. While the aim of SBA is to minimize the number of redundant transmissions, ODMRP follows a mesh-based approach in which redundant routes are created intentionally. Thus the comparison between those protocols does not provide a full understanding of the trade-offs between multicasting and broadcasting.

Towards the goal of investigating the trade-off between multicasting and broadcasting, we perform extensive simulation studies on a chosen protocol from each class: Network-wide broadcasting through time reservation using adaptive control for energy efficiency (NB-TRACE) [4] for broadcasting; and Multicasting through time reservation using adaptive control for energy efficiency (MC-TRACE) [5] for multicasting. The first reason for choosing these protocols is that they have been shown to outperform many other protocols in their class. Moreover, those protocols are built on top of the same MAC structure, and their sensitivity to MAC layer issues such as mobility and link errors are similar. Finally, the data maintained by the protocols are very similar to each other, and any additional burden of multicasting can directly be observed. Consequently, the protocols can be combined into a unique framework and coexist simultaneously. Ultimately, this approach potentially yields a unified protocol where the better approach (broadcasting or multicasting) can be used depending on the situation.

The rest of the paper is organized as follows. In Section II, the chosen candidate protocols are introduced. Then, we analyze the efficiency of the protocols for a sample scenario in Section III. The effect of node density on the relative efficiency of the protocols is investigated in Section IV. Finally, we conclude the paper in Section V.

II. Description of the Chosen Multicast and Broadcast Protocols

The main purpose of this section is to give the reader insight into the differences between the protocols. The details of the protocols, NB-TRACE and MC-TRACE, can be found in references [4] and [5], respectively.

This work was supported in part by the University of Rochester Center for Electronic Imaging Systems and in part by Harris Corporation, RF Communications Division.

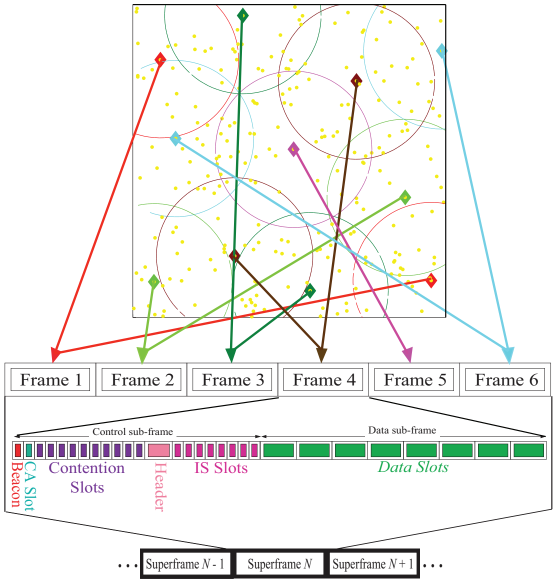

Fig. 1. Clustering and frame structures of MH-TRACE. Diamonds represent selected clusterheads and dots represent the nodes in the network. CH-frame matching together with the contents of each frame is depicted.

Both MC-TRACE and NB-TRACE are cross layer approaches where the MAC layer and the routing layer functionalities are implemented together in a unique framework. The MAC scheme of the protocols follows from MH-TRACE [6], where the network is organized into overlapping clusters, each managed by a clusterhead (CH). Each CH holds the channel access rights in its neighborhood and grants channel access to the neighbor nodes on demand. Clusters and CHs are chosen dynamically to ensure the existence of at least one CH in the neighborhood of each node. In other words, CHs form a dominating set. Fig. 1 depicts the clustering structure of MH-TRACE. Each CH chooses one of the frames and operates in that frame. Beacon and clusterhead announcement (CA) slots are used to create and maintain the clustering structure. Nodes pass their channel access request to a nearby CH using the contention slots. After receiving the channel access requests, the CH sends a header message that includes the transmission schedule that will be followed for the rest of the frame. Nodes that acquire channel access from the CH transmit an information summarization (IS) packet that contains the unique packet ID that will be transmitted together with the routing information prior to the data payload that is transmitted at the corresponding data slot. By listening to the relatively shorter IS packets, nodes become aware of the data that are going to be sent and may choose to sleep during the corresponding data slots. Acknowledgements are embedded into the IS messages.

A. NB-TRACE

Routing in NB-TRACE makes use of the clustering structure. The protocol sends a copy of each data packet to all of the CHs, and the CHs retransmit these packets to their cluster members. Each data session starts with an initial flooding stage where each rebroadcasting node implicitly acknowledges its upstream node through IS packets as a part of its transmission. In the case of the existence of more than one upstream node, only one of them is selected and announced in the IS packet. A node drops its relaying status and stops retransmitting the packets when it does not receive an acknowledgement for a certain amount of time. Only the CHs keep retransmitting the packets even when they do not receive any downstream acknowledgement. This behavior prunes the redundant retransmissions and creates a tree that starts from the source node and ends at the CHs. The dynamic behavior of the network is handled by a local branch repair mechanism.

B. MC-TRACE

MC-TRACE implements multicast routing on top of MH-TRACE using a mixed layer approach. Like NB-TRACE, MC-TRACE also starts with an initial flooding stage. Nodes that do not receive a downstream acknowledgement stop retransmitting. However, in MC-TRACE, CHs do not take a special role in routing. Instead, the member nodes keep sending an acknowledgement to their upstream node even when they do not receive any downstream acknowledgements. Therefore, the tree is kept alive directly by the group members.

Furthermore, in MC-TRACE, retransmitting nodes also choose and announce a downstream node in addition to their upstream node. The first node that sends an upstream acknowledgement is selected as the downstream node and announced in the following transmissions. The node that is announced as the downstream node is responsible for sending upstream acknowledgements and keeping the branch alive. With the help of this mechanism, in the case of more than one leaf member node receiving the data from the same branch, only one of them sends the acknowledgement messages. Although this mechanism eliminates redundant acknowledgements, the need for acknowledgements makes MC-TRACE consume considerably more resources compared to NB-TRACE when the cluster members are spread throughout the region.

III. Comparing Multicast and Broadcast

In general, multicasting protocols eliminate redundant retransmissions by confining the data dissemination to a limited area. However, this comes with the additional cost of overhead to keep the data distribution structure alive. Intuitively, while multicasting is expected to be a more efficient method of data distribution for small group sizes, broadcasting would be more efficient for large group sizes. In this section, we show that this is indeed the case by analyzing broadcasting and multicasting through extensive simulations of a select broadcasting protocol, NB-TRACE, and a select multicasting protocol, MC-TRACE.

In particular, the number of multicast group members beyond which NB-TRACE becomes more efficient for data dissemination, called the cross-over point, is determined for various scenarios. By comparing the simulation results, we can analyze the effect of the total number of nodes in the network and the size of the region in which the nodes are distributed on the value of this cross-over point.

Two performance metrics, energy efficiency and spectrum efficiency, are considered. Specifically, the average energy spent per node per generated packet and the the total number of transmissions per generated packet are measured for each scenario to compare the energy efficiency and the spectrum efficiency of the protocols, respectively.

We begin with describing the simulation environment and the parameters selected for the scenarios under concern. Then, for a network of 100 nodes distributed in a 1×1 km^2 area, the bandwidth efficiency of the protocols and their energy consumptions are compared. Finally, the analysis is extended by varying the number of nodes in the network and the size of the area in which nodes are distributed.

A. Simulation Environment

We conduct ns-2 simulations of NB-TRACE and MC-TRACE under different network scenarios. We used the default energy and propagation (two-ray ground) models in ns-2. Both path loss and interference are taken into account in determining a successful reception. The receiver can receive only those packets whose received power is above a certain threshold. For a successful reception, the receiver has to be within 250m of the transmitter with the given propagation model, reception threshold and transmission power. However, simultaneous transmissions interfere with each other and prevent successful reception. In the case of simultaneous transmissions, a successful reception is only possible if, at the receiver side, the power of one of the packets is 10 times larger than any other packet. The transceivers are fixed at 2 Mbits/sec data rate.

The packet generation model assumes a 16-bit voice coder that generates 100 byte packets every 25ms. The length of an IS slot is 12 bytes long in NB-TRACE, while it is 15 bytes long in MC-TRACE due to the extra routing information requirements. The superframe period is fixed to the packet generation period, and the number of frames per superframe is fixed at 6 for both NB-TRACE and MC-TRACE. However, due to the extra bits in the IS slots, MC-TRACE has 6 data slots per frame whereas NB-TRACE has 7. Each node transmits or relays the data packets of the stream using one of the available data slots.

In order to have a fair comparison in terms of energy consumption, the concept of group members is introduced to NB-TRACE. Nodes that do not belong to the group do not listen to the data slots of the stream.

The power spent by each node varies according to the operation performed by the node. During successful reception, collision and carrier sensing periods, the node consumes power at the rate of the reception power level. There is also an idle state where only the power needed to run the circuitry is dissipated without any actual packet receptions. The nodes are assumed to turn off any circuitry when they go into the sleep state, where the power consumption is minimal.

All the nodes except the source node are initially distributed according to a uniform random distribution, and during the course of the simulation the nodes move following the random way-point mobility model [7] [8] with node speeds chosen from a uniform random distribution between 0.0 m/s and 5.0 m/s with zero pause time. The source node starts from the center of the region and follows the same random way-point mobility model.

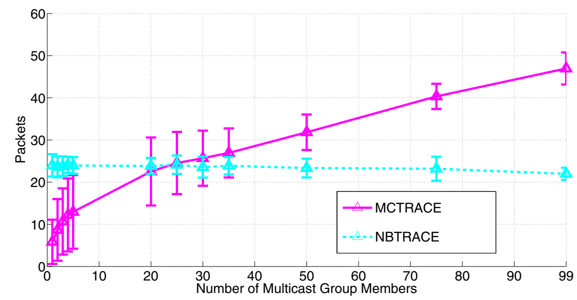

Fig. 2. Number of transmitted packets per generated packet as the number of group members is varied from 1 to 99 for a network of 100 nodes distributed on a 1×1 km^2 area.

100 repetitions are performed for each scenario, and the presented results show the averages and the standard deviations of the results.

B. Bandwidth Efficiency

In the TRACE frame structure, there are fixed number of data slots per superframe. The data slots are used to transmit the payload of both the generated data and the relayed data. Any redundant use of the data slots wastes the available network resources and may prevent another stream's information from being disseminated over a region. Thus, efficient use of the data slots is one of the goals of both NB-TRACE and MC-TRACE.

The number of data packet transmissions per generated packet is depicted in Fig. 2 for both NB-TRACE and MC-TRACE. The change in the number of multicast members does not effect the number of data transmissions in NB-TRACE significantly since data is broadcasted to the entire network regardless of the locations and number of multicast members. The slight variation in NB-TRACE is due to the additional restriction we impose on the algorithm to make a fair comparison between NB-TRACE and MC-TRACE. We prevent non-group members from listening to the ongoing stream unless they are relaying the data to save energy. However, this reduces the efficiency of the recovery of link failures through local branch repair mechanisms, since the neighbor nodes that are not in the multicast tree do not have any data in their transmission queue.

Unlike NB-TRACE, in MC-TRACE, the number of transmissions increase as the number of group members increase. Nodes are distributed with an independent identical random uniform distribution[1]. Hence, as the number of group members

[1]It is well known that the random way point mobility model alters the uniform node distribution assumption as time elapses [9] [10]. However, this effect is negligible in our case since the simulation duration of 100 sec is short enough.

increases, the size of the multicasting tree increases. Hence, the number of transmissions per generated packet increases.

It can be observed from Fig. 2 that for small group sizes the data dissemination is more efficient in MC-TRACE compared to NB-TRACE, and vice versa. NB-TRACE requires 4 times more transmissions compared to MC-TRACE when there is only one group member. At the other extreme, when all the nodes are in the multicast group, the number of data transmissions required using NB-TRACE is half of the transmissions required when using MC-TRACE. The cross-over point of MC-TRACE and NB-TRACE occurs at a multicast group size of 24 nodes, above which NB-TRACE provides more efficient data dissemination and vice versa.

C. Energy Efficiency

In Section III-B, we compared the number of data transmissions for NB-TRACE and MC-TRACE. The number of transmissions and receptions of other packet types, namely beacon, CA, contention, and header are expected to be comparable. Thus, a similar trade-off that exists in Section III-B is also expected for the energy consumption metric.

On the other hand, since each IS slot is matched with its corresponding data slot, the difference in the number of data transmissions between protocols is also expected to be observed in the number of IS transmissions. Furthermore, since the length of IS slots in MC-TRACE are longer compared to the ones in NB-TRACE, the energy consumption of MC-TRACE is expected to be even higher. This behavior is expected to shift the energy consumption curve of MC-TRACE in such a way that the cross-over point will occur with fewer multicast group members.

The energy consumption per node per generated packet is depicted in Fig. 3 as the number of group members is varied from 1 to 99. The energy consumption increases with an increase in the number of multicast members for both protocols. Although with an increasing number of multicast group members there was a slight decrease in the number of transmitted data messages for NB-TRACE, shown in Fig. 2, the energy consumption increases. This is due to the fact that the energy consumption in the reception and transmission states of a node are of the same order. As the number of multicast group members increases, the number of receptions increase, which in turn increases the energy consumption.

It can be observed from Fig. 3 that for small multicast group sizes MC-TRACE is more energy efficient while for large group sizes NB-TRACE performs a more energy efficient operation. NB-TRACE consumes 13% more energy compared to MC-TRACE when there is only one group member. On the other hand, when all 99 nodes are in the multicast group, the energy consumption of NB-TRACE is 21% lower than the energy consumption of MC-TRACE.

The energy consumption of MC-TRACE increases faster than the increase in NB-TRACE and goes above the energy consumption of NB-TRACE at the cross-over point of 11 multicast group members. The cross-over point in Fig. 3 is

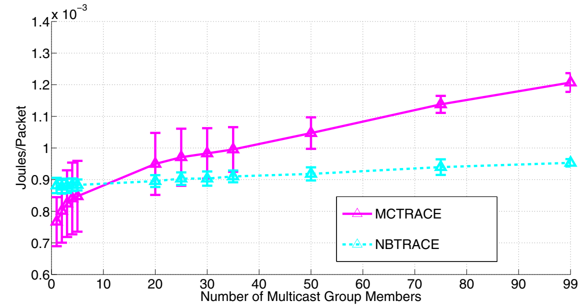

Fig. 3. Energy consumption per generated packet per node as the number of group members is varied from 1 to 99 for a network of 100 nodes distributed on a 1×1 km^2 area.

lower than the cross-over point observed in Fig. 2. This is expected since the length of the IS slots in MC-TRACE is higher than the length of the IS slots in NB-TRACE, which increases the energy consumption of MC-TRACE compared to NB-TRACE.

IV. Effect of Node Density

In Section III we observed that, depending on the number of multicast group members, NB-TRACE and MC-TRACE can be advantageous over the other one in terms of bandwidth and energy efficiency. We have identified the cross-over points for a network consisting of 100 nodes distributed on a 1×1 km^2 area. In this section, we extend our analysis to other scenarios to see the effects of node density on the efficiency of the protocols.

In considering the density, both the number of nodes and the size of the area in which the nodes are distributed are of concern. The nodes close to the edge of the target area have fewer neighbors and data dissemination close to the edges follows a pattern different than the data dissemination pattern of the nodes residing on the middle of the region. In order to investigate the edge effects accurately, variation in the number of nodes and the size of the area are considered separately.

Specifically, we investigate the cross-over points on the number of multicast group members above which NB-TRACE is more efficient compared to MC-TRACE and vice versa in terms of the number of data packet transmissions and energy consumption for a range of total number of nodes and area sizes. The observed cross-over points are presented in Table I.

The first observation from Table I is that for a given number of nodes, the cross-over point for the number of data transmissions increases with an increase in the size of the area. For larger areas, the separation between the source and the multicast members is larger. This translates into a larger number of hops between the source and a destination, and hence a higher number of transmissions in both NB-TRACE and MC-TRACE. However, the increase in the number of transmissions in NB-TRACE is larger than it is in MC-TRACE since the number of data transmissions in NB-TRACE is effected not only by an increase in the expected distance between the source and the multicast members but also by an increase in the expected separation between the source and

TABLE I
CROSS-OVER POINTS BETWEEN MC-TRACE AND NB-TRACE

Sim Area	**1000x1000**		**1500x1500**		**2000x2000**	
Number of Nodes	**100**	**200**	**200**	**400**	**200**	**400**
Number of Data TX	24	30	45	55	55	62
Energy Consumption	11	18	19	29	18	30

non-multicast members. As a result, the value of the cross-over point for the number of data transmissions metric increases with an increasing size of the network.

Table I also shows that, for a fixed network size, the network with more nodes has a higher cross-over point. Since the node locations are independent, the separation between the source and the multicast members is independent of the number of nodes in the network. Hence, the number of transmissions in MC-TRACE does not deviate significantly as the total number of nodes increases. On the other hand, as the total number of nodes in the network increases, they are expected to cover a larger region on the simulation area, and the number of CHs are expected to increase. As a result of this, the number of transmissions in NB-TRACE increases up to a limit where the entire area is covered as the total number of nodes increases. Thus, the cross-over point also increases for an increasing number of nodes.

The energy consumption of a node consists of a variable part that is incurred for data packet transmissions and receptions and a relatively constant part for the control messages. For a fixed network size and a fixed number of multicast members, the number of nodes that take part in the data dissemination tree created in MC-TRACE is independent of the total number of nodes in the network. Hence the variable energy consumption for data dissemination is also independent of the number of nodes. When the number of nodes in the network increases, the energy consumption per node in MC-TRACE decreases. On the other hand, for a fixed network size and a fixed multicast group size, the energy consumption in NB-TRACE is independent of the number of nodes in the network since the data dissemination tree covers the entire network. As a result, the cross-over point increases with increasing number of nodes in the network, as can be observed from Table I.

It is interesting to note that, when considering energy-efficiency, the cross-over point is approximately constant for a fixed number of nodes as the network area increases. This is because the energy consumption of both NB-TRACE and MC-TRACE increases as the size of the area increases, mainly due to the increase in the number of clusters and in turn the additional control messages. Hence the increase in the average energy consumption is of the same order in NB-TRACE and MC-TRACE and does not alter the value of the cross-over point significantly for a fixed number of nodes.

To sum up, for the goal of minimizing the number of transmissions, both an increase in the number of nodes and an increase in the size of the network favor multicasting over network-wide broadcasting, making multicasting the optimal choice for a larger set of multicast members. Similarly, considering the goal of minimizing the energy consumption per node, an increase in the number of nodes in the network makes multicasting a better choice up to a larger number of multicast members. On the other hand, the relative efficiency of multicasting and network-wide broadcasting in energy consumption is independent of the size of the network.

V. CONCLUSION

In this paper, we examined the effect of the number of multicast members on the relative efficiency of multicasting and broadcasting. We consider two performance metrics: energy efficiency and bandwidth spectrum efficiency. We showed that for large multicast groups, using broadcasting instead of multicasting leads to 75% savings in the number of transmitted data packets and up to 21% savings in the average energy consumption. Similarly, for small multicast groups, multicasting reduces the number of data transmissions and the average energy consumption by 50% and 13%, respectively.

We also showed that an increase in the total number of nodes in the region decreases the relative efficiency of broadcasting compared to multicasting for both performance metrics, and hence the cross-over occurs at a larger number of multicast members. On the other hand, the increase in the size of the area does not effect the cross-over point significantly for the energy efficiency metric, while it increases the cross-over point for the bandwidth efficiency metric.

Due to the similarity between the protocols, NB-TRACE and MC-TRACE, they can be combined into a new unified protocol, U-TRACE. Based on the a-priori information about the number of nodes and the size of the network, the source node can choose the appropriate type of operation considering the number of nodes in the multicasting group.

REFERENCES

[1] C. de Morais Cordeiro, H. Gossain, and D. Agrawal, "Multicast over wireless mobile ad hoc networks: Present and future directions," *IEEE Network*, vol. 17, pp. 52–59, 2003.

[2] P. Mohapatra, C. Gui, and J. Li, "Group communications in mobile ad hoc networks," *IEEE Computer Magazine*, vol. 37, pp. 52–59, 2004.

[3] L. K. Law, S. V. Krishnamurthy, and M. Faloutsos, "Understanding and exploiting the trade-offs between broadcasting and multicasting in mobile ad hoc networks," *IEEE Transactions on Mobile Computing*, vol. 6, no. 3, pp. 264–279, 2007.

[4] B. Tavli and W. B. Heinzelman, "Energy and spatial reuse efficient network-wide real-time data broadcasting in mobile ad hoc networking," *accepted to IEEE Journal on Selected Areas of Communications*, 2005.

[5] B. Tavli and W. Heinzelman, "MC-TRACE: multicasting through time reservation using adaptive control for energy efficiency," in *Military Communications Conference, 2005. MILCOM 2005. IEEE*, 2005, pp. 2672–2678 Vol. 4.

[6] B. Tavli and W. B. Heinzelman, "MH-TRACE: Multi hop time reservation using adaptive control for energy efficiency," *IEEE Journal on Selected Areas of Comm.*, vol. 22, no. 5, pp. 942–953, June 2004.

[7] T. Camp, J. Boleng, and V. Davies, "A survey of mobility models for ad hoc network research," *Wireless Communications and Mobile Computing (WCMC)*, vol. 2, no. 5, pp. 483–502, 2002.

[8] C. Bettstetter and C. Wagner, "The spatial node distribution of the random waypoint mobility model," in *Mobile Ad-Hoc Netzwerke, WMAN 2002*. GI, 2002, pp. 41–58.

[9] C. Bettstetter, G. Resta, and P. Santi, "The node distribution of the random waypoint mobility model for wireless ad hoc networks," *IEEE Transactions on Mobile Computing*, vol. 2, pp. 257–269, 2003.

[10] J. Yoon, M. Liu, and B. Noble, "Sound mobility models," in *Proceedings of the ACM International Conference on Mobile Computing and Networking (MOBICOM)*, 2003, pp. 205–216.