Multicast Routing in Wireless Mesh Networks: Minimum Cost Trees or Shortest Path Trees?

Uyen Trang Nguyen, Jin Xu
Department of Computer Science and Engineering
York University, Toronto, Ontario M3J 1P3, Canada
Email: {utn, jinxu}@cs.yorku.ca

***Abstract*— There exist two fundamental approaches to multicast routing: shortest path trees (SPTs) and minimum cost trees (MCTs). The SPT algorithms minimize the distance (or cost) from the sender to each receiver, while the MCT algorithms minimize the overall cost of the multicast tree. Due to the very large scale and unknown topology of the Internet, computing MCTs for multicast routing in the Internet is a very complex problem. As a result, the SPT approach is the more commonly used method for multicast routing in the Internet, because they are easy to implement and give minimum delay from the sender to each receiver, a property favored by many real-life applications. Unlike the Internet, a wireless mesh network (WMN) has a much smaller size, and its topology can be made known to all nodes in the network. This makes the MCT approach an equally viable candidate for multicast routing in WMNs. However it is not clear how the two types of trees compare when used in WMNs. In this article we present a simulation-based performance comparison of SPTs and MCTs in WMNs using the performance metrics such as packet delivery ratio, throughput, end-to-end delay, delay jitter, and traffic impacts on unicast flows in the same network.**

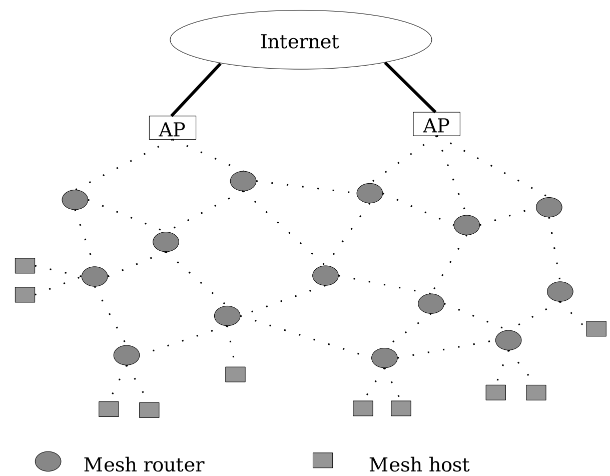

Fig. 1. Architecture of a WMN

I. Introduction

Wireless Mesh Networks (WMNs) form a new class of networks that has emerged recently. Major components of a WMN include wireless mesh routers, wireless mesh hosts (e.g., PCs, laptops, PDAs, and cell phones), and access points (or gateways) that act both as Internet routers and wireless mesh routers (Figure 1). The mesh routers in a WMN provides multi-hop connectivity from one mesh host to another, or to the Internet via access points. A WMN is dynamically self-organized and self-configured, with nodes in the network automatically establishing and maintaining mesh connectivity among themselves. This feature brings many benefits to WMNs such as low installation cost, large-scale deployment, reliability, and self-management.

Multicast is a form of communication that delivers information from a source to a set of destinations simultaneously in an efficient manner; the messages are delivered over each link of the network only once and only duplicated at branch points, where the links to the destinations split. Important applications of multicast include distribution of financial data, billing records, software, and newspapers; audio/video conferencing; distance education; and distributed interactive games. Although multicast is required to support many important applications, research on multicasting in WMNs is still in its infancy. In this article we address one of the most essential issues of multicast in WMNs – routing.

There are two fundamental multicast routing approaches: shortest path trees (SPTs) and minimum cost trees (MCTs). The goal in SPT algorithms is to compute a tree rooted at the sender and spanning all the receivers such that the distance between the sender and each receiver along the tree is minimum[1]. The two most well-known algorithms for computing SPTs are due to Bellman-Ford and Dijkstra [1]. These techniques are equivalent to applying point-to-point shortest path algorithm repeatedly, once for each sender-receiver pair. SPTs by definition are per sender. Therefore for many-to-many multicast, separate multicast trees need to be computed, one for each sender.

Unlike the SPT algorithms, which aim at minimizing the distance (or cost) from the sender to each receiver, the goal of MCT algorithms is to minimize the overall cost of the multicast tree[2]. MCT algorithms for multicast routing are based on the minimum Steiner tree problem, which is NP-complete. Thus several heuristics have been proposed to compute approximate Steiner trees [2], e.g., the 2-approximation heuristic proposed by Kou *et al.* [3], and the 11/6-approximation algorithm by Zelikovsky [4].

[1] Another way of interpreting this is to say that the SPT algorithms aim at minimizing the *end-to-end delay*.

[2] There also exist algorithms that try to optimize on both overall cost and source-to-destination distance (or end-to-end delay) [2].

Note that the total cost of a Steiner tree is less than the total cost of a corresponding SPT, by definition of Steiner trees. However the maximum distance between the sender and any receiver in a Steiner tree is typically longer than that in a SPT. This means that the average path length in a Steiner tree is more than that in a SPT.

Due to the complexity of computing Steiner trees in a distributed manner, the majority of the multicast routing protocols used in the Internet today are based on SPTs[3], such as DVMRP [5] and MOSPF [6]. The reason is that SPTs are easy to implement and offer minimum delay from the sender to each receiver, a desirable quality of service parameter for most of the real-life multicast applications [2].

Recently, Ruiz and Gomez-Skarmeta explored the problem of multicast routing in wireless multi-hop networks in which nodes are static, e.g., WMNs [7]. The authors re-define the cost of a MCT by reasoning that in a broadcast medium, the transmission of a multicast data packet from a given node to any number of its neighbors can be done with a single data transmission. Thus, in a wireless multi-hop network the minimum cost tree is one which connects sources and receivers by issuing a minimum number of transmissions, rather than having a minimal edge cost as defined for traditional Minimum Steiner Trees (MSTs). In other words, the tree contains a minimum number of *multicast forwarding nodes*. The authors show that a Steiner tree typically needs a higher number of transmissions than their proposed Minimum Number of transmissions Trees (MNTs).

Figure 2 shows an example of a multicast group having four receivers and the three different routing trees constructed using the shortest path tree, the minimum Steiner tree, and the minimum number of transmissions tree algorithms. The example illustrates the typical characteristic of each kind of tree. Assuming the cost of every edge is one, the SPT provides the shortest average path length (2.25 hops); the MNT has the lowest cost (5 units); the MNT contains the least number of forwarding nodes (2 nodes), and requires the least number of transmissions per packet (3 transmissions/packet).

Ruiz and Gomez-Skarmeta demonstrated that the problem of computing MNTs is NP-complete and proposed enhanced heuristics to approximate such trees [7]. They presented experimental results that show that the MNT algorithm offers the least number of transmissions compared with the MST and the SPT algorithms. On the other hand, the mean path lengths given by the MST and the MNT algorithms are longer than that by the SPT algorithm, as expected. However, the presented experimental results do not indicate how the multicast groups perform in terms of packet loss rate (or packet delivery ratio) and throughput – the true performance measures of a multicast session – or end-to-end delay – an important performance metric for realtime multicast applications such as distribution of stock quotes, distributed interactive games and tele-conferencing.

One could argue that SPTs are best in an environment where the network topology is unknown and the multicast members may be geographically distributed over a very large area, such as the Internet. In a WMN where the topology is known and the network size is much smaller (e.g., less than 500 nodes), MCTs such as Steiner trees or MNT trees are no longer difficult to implement, and could potentially offer better performance because they typically consume less bandwidth than SPTs. On the other hand, MCT algorithms produce longer paths than SPT algorithms. In a wireless multi-hop network, the longer the path, the higher the probability that a packet will be lost due to packet collision or congestion, resulting in a throughput reduction. One could thus infer that SPTs could achieve higher throughput than MCTs. In short, it is not clear how the performance of SPTs and MCTs in a wireless multi-hop network compares.

In this article, we present a simulation-based performance comparison of minimum cost trees and shortest path trees. Specifically we examine the following two minimum cost tree algorithms: the minimum Steiner tree (MST) heuristic proposed by Kou *et al.* [3], and the MNT heuristic by Ruiz and Gomez-Skarmeta [7]. We use the Dijkstra's algorithm for computing SPTs. We measure the packet delivery ratio, throughput, end-to-end delay, and delay jitter of multicast trees built by the MST, MNT and Dijkstra's algorithms. We also study the effects of data traffic generated by the multicast nodes in MSTs, MNTs and SPTs on the packet loss rate of unicast flows in the same network.

In the remainder of this article, we first describe our simulation setting and define the performance metrics. We then present experimental results that compare the performance and traffic impacts of multicast trees constructed the MST, MNT and SPT algorithms. Finally we conclude the article.

II. Experiment Setting

Our experiments were carried out using QualNet [8], a software that provides scalable simulations of wireless networks and a commercial version of GloMoSim [9].

A. Simulation Parameters

Our simulation models a medium-size network of 100 wireless routers placed in a 2000m x 2000m terrain, and a large network of 300 routers placed in a 3000m x 3000m terrain. We will use the term "router" and "node" interchangeably in this article. The nodes are distributed uniformly over the sub-areas within a terrain, and the nodes within a sub-area are randomly placed in that sub-area. There are no network partitions throughout the simulation. The edge cost or the cost of each wireless link is assumed to be one.

A two-ray propagation model is used in our experiments, with free space path loss (2.0, 0.0) for near sight and plane earth path loss (4.0, 0.0) for far sight. The transmission power is set constant at 20dBm; the corresponding transmission range of the routers is 315m; and the data transmission rate at the physical layer is 11Mbps, according to the specifications of wireless routers manufactured by TROPOS [10]. We use PHY802.11b at the physical layer. We implement only one

[3]Some of the multicast routing protocols in the Internet proposed recently are based on sub-optimal shared trees as opposed to per-source SPTs. The reason for using such shared trees is to reduce state information to be stored in the routers rather than to minimize the overall cost [2].

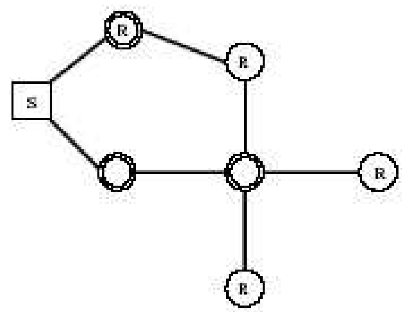

(a) Shortest path tree

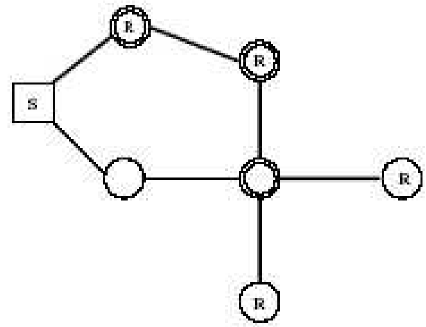

(b) Minimum Steiner tree

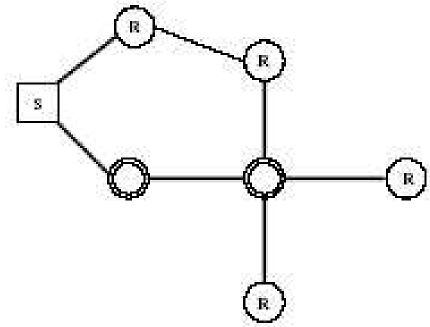

(c) Minimum number of trx tree

S: sender – R: receiver – Double-lined circles denote multicast forwarding nodes.

Type of tree	Avg path length	Tree cost	Number of forwarding nodes	Number of trx per packet
Shortest path tree	2.25	6	3	4
Minimum Steiner tree	2.75	5	3	4
Minimum number of trx tree	2.50	6	2	3

Fig. 2. Different routing trees over the same network

channel on each wireless link. The MAC802.11 with Distributed Coordination Function (DCF) is chosen as the medium access control protocol. Multicast and broadcast transmissions use CSMA/CA. Unicast transmissions also use RTS/CTS in addition to CSMA/CA.

We assume that each sender or receiver is connected to a *different* router. (In practice, there can be many hosts communicating with a router, e.g., to form a wireless local area network). The sender and the receivers of a multicast group were selected randomly, and the same sender and receivers and the same network configuration were used for all three algorithms (MST, MNT and SPT) in order to obtain a fair comparison. All receivers joined a multicast group at the beginning and stayed until the whole group terminated. The sender of a multicast group transmits at a *constant bit rate* specified for each experiment.

The data packet size excluding the header size is 512 bytes. The size of the queue at every node is 50 Kbytes. The packets in a queue are scheduled on a first-in-first-out basis. We did not implement flow or congestion control in order to test the network performance under very high loads.

Each experiment executed for 600 seconds of simulated time. Each data point in the graphs was obtained from 10 runs using different randomly generated seed numbers, and the collected data were averaged over the 10 runs.

B. Performance Metrics

We use the following metrics to measure the performance of a multicast routing protocol:

- **Average packet delivery ratio** (PDR). The PDR of a receiver of the number of data packets actually delivered to the receiver versus the number of data packets supposed to be received. The average PDR of a multicast group is the average of the PDRs of all the receivers in the group.
- **Average end-to-end delay**. The end-to-end delay of every packet received at every receiver is recorded; the average over all the packets received is then computed.
- **Average throughput** . The *throughput* is defined as the total amount of data a receiver actually receives divided by the time between receiving the first packet and the last packet. The average taken over all the receivers is the average throughput of the multicast group, assuming that each group has one sender.
- **Average delay jitter** . Delay jitter is the variation (difference) of the inter-arrival intervals from one packet received to the next packet received. The per-receiver delay jitter at a receiver is the sum of all the absolute values of delay jitters from the first packet received to the last packet received divided by the total number of packets received. The average delay jitter is the average of the per-receiver delay jitters taken over all the receivers.
- **Average path length**: The path length is an indirect indicator of performance, as will be shown in the next section. The average path length is the average of the lengths of all source-to-destination paths in a multicast tree.

III. Experimental Results

We compare the performance of multicast sessions created by the SPT, MST and MNT algorithms using the above metrics. We also examine the impacts of the traffic created by the multicast nodes on unicast flows.

A. Average Path Lengths

In this set of experiments, we simulated different multicast groups by varying the number of receivers from 10 to 80. Each multicast group has one source. The graphs in Figure 3 show the average path lengths of the trees constructed by the three algorithms to be compared. The results confirm that the MST and MNT algorithms produce longer paths than the SPT algorithm in all cases. Furthermore, the larger the network, the wider the difference gap. For instance, in the case of 40 receivers, the MST/MNT average path length is about 20% longer than the SPT average path length in the network of 100 nodes, but about 40% longer in the network of 300 nodes.

B. Multicast Performance

In the network of 100 nodes, we examined two multicast groups having 20 and 40 receivers respectively. The results are illustrated in Figure 4. When the traffic load is light (under 30 packets/sec), the performance of the SPT, the MST and the MNT is comparable with respect to packet delivery ratio and throughput. When the traffic load is moderate or high, the SPT outperforms the MST and the MNT in all cases, and the difference can be significant. For example, when the number of receivers is 20 and the traffic load is 60 packets/sec, the PDRs of the SPT, the MST and the MNT are 99%, 87%, and 94%, respectively. The reason is due to longer path lengths of the MST and the MNT. The longer the path a packet has to travel, the higher its chance of getting damaged or lost due to collision and/or congestion, especially under high traffic load.

The average end-to-end delays incurred by the SPTs are also the lowest thanks to shorter source-to-destination paths. For example, in the case of 20 receivers and traffic load of 60 packets/sec, the average end-to-end delays given by the SPT and the MNT are 11.4ms and 15.8ms, respectively; in other words, the SPT average end-to-end delay is almost 40% lower. The MST provides the highest average end-to-end delay in this case. The SPTs also give the lowest average delay jitter values in almost all the cases.

In the larger network of 300 nodes, the performance differences between the SPT and the MST/MNT are even more pronounced, as illustrated by the graphs in Figure 5, which show the results of two multicast groups having 40 and 80 receivers respectively. For instance, when the number of receivers is 40 and the traffic load is 60 packets/sec, the PDRs of the SPT, the MST and the MNT in the network of 100 nodes are 98%, 93% and 94%, respectively, while the PDRs in the network of 300 nodes are are 96%, 72%, and 72%, respectively, a difference of 24%. In the same scenario, the average end-to-end delay given by the SPT is about 20% lower than those of the MST and the MNT in the smaller network, and about 65% lower in the larger network. In other words, given the same multicast group size, as the network size increases, the performance gain of SPTs over MSTs and MNTs also increases. The reason is that the larger the network, the bigger the difference in path length between the SPT and the MST/MNT as mentioned earlier.

C. Impacts of Multicast Traffic on Unicast Flows

This set of experiments evaluates the impacts of data traffic generated by the SPT, the MST and the MNT. We considered three scenarios: the network has one, three and five multicast groups (or senders), respectively. Each sender/group has 10 receivers. We measured the impacts of multicast data traffic on the packet loss rate, or the packet delivery ratio (PDR), of 20 unicast flows. Each unicast source transmits at a constant bit rate of 1 packet/sec, and the packet size is 512 bytes excluding the header size. We varied the traffic load generated by each multicast sender and measured the average PDR of the unicast flows. Figure 6 displays the results of this set of experiments on networks of 100 and 300 nodes.

When the number of multicast nodes in the networks is small or moderate (10 to 30 receivers), the SPT, the MST and the MNT have similar effects on the average PDR of the unicast flows. When the multicast group is large (e.g., 5 senders and 50 receivers) *and* the total multicast traffic load is high (150 packets/sec or more), the SPT causes more packet losses to the unicast flows than the MST and the MNT, from 1% to 3.5% more. This is due to the fact that the SPT involves more nodes in the data forwarding task than the MST and the MNT, and thus causes more packet collisions and more congestion to the unicast flows when the multicast sending rate is high.

IV. Discussion and Conclusions

We quantify the performance differences of minimum cost trees and shortest path trees in WMNs. Our simulation results show that SPTs offer significantly better performance to multicast flows than MCTs such as MSTs and MNTs. The average packet delivery ratio given by SPTs is higher by up to 24%, and the SPT average end-to-end delay is up to 65% lower in our experiments. The only drawback of SPTs is that when the multicast sending rate is high SPTs cause more packet losses to other flows than MCTs. The reason is that in a SPT typically more nodes are involved in the data forwarding task compared with a MCT. However, SPTs cause only 1% to 3.5% more packet loss to unicast flows than MCTs in our experiments, and only when the multicast traffic load is high. Under light or moderate traffic, SPTs and MCTs have similar effects on other flows in the network. In our opinion, the high multicast performance gain of SPTs outweights the above drawback.

Note that it is much easier to design a reliable transport protocol for unicast communications than for multicast communications in wireless multi-hop networks. Several TCP-based protocols have been proposed for reliable data delivery for unicast flows in wireless ad-hoc networks [11], which could be applied to WMNs. However the problem of reliable multicast in wireless ad-hoc networks still remains unsolved. This justifies the use of SPTs for multicast routing in WMNs as they offer noticeably higher packet delivery ratio to multicast receivers than MCTs. The unicast packets lost can be retransmitted and delivered using the existing TCP-based protocols. It is harder to recover the multicast packets lost, especially when the loss rate is high and there does not exist an effective reliable multicast protocol for wireless multi-hop networks.

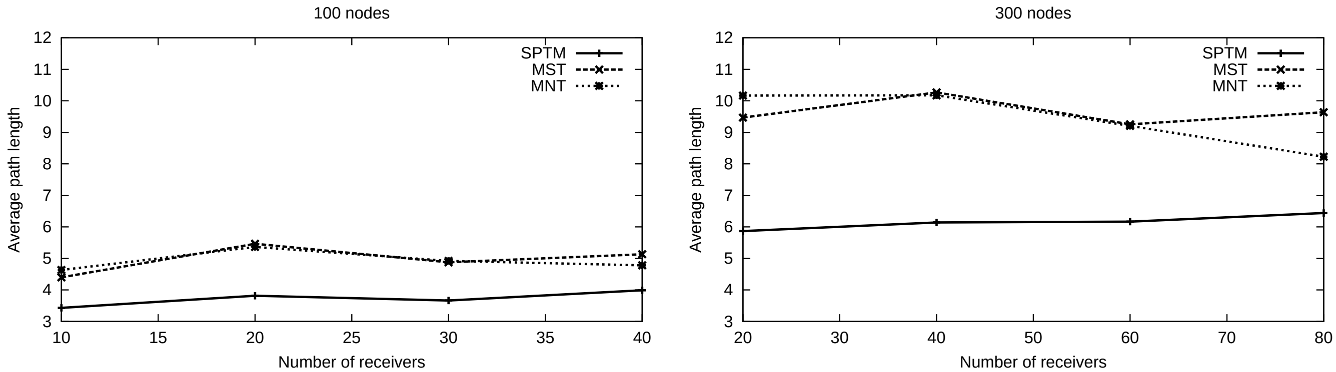

(a) Network of 100 nodes   (b) Network of 300 nodes

Fig. 3. Average path lengths

Our simulations assume one channel for all transmissions and IEEE 802.11 CSMA/CA medium access control for multicast transmissions. Until an effective throughput enhancement technique for multi-hop routing and/or an efficient MAC protocol for multicast are implemented in WMNs, SPTs offer the best tradeoff in terms of multicast performance and traffic overhead.

It is also worth mentioning another advantage of SPTs over MCTs. Although the mesh routers are static, mesh hosts attached to the end-routers such as cell phones, PDAs, and laptops may move from one end-router to another, and may join and leave a multicast session freely at will. It is much easier to support dynamic joins and leaves using SPTs than MCTs, because in a SPT each source-to-destination path is established independently of the other paths in the tree. In a MCT, a node joining or leaving the multicast session may require the whole tree to be re-computed in order to maintain the cost optimality (or the new tree would no longer be optimal). This implementation difficulty of MCTs is another factor that makes SPTs the more favorable multicast routing approach in WMNs.

REFERENCES

[1] D. Bertsekas and R. Gallagher, Data Networks, Prentice Hall, 1991.
[2] Sanjoy Paul, Multicast on the Internet and Its Applications, Kluwer Academic Publishers, 1998.
[3] L. Kou, G. Markowsky, and L. Berman, "A Fast Algorithm for Steiner Trees," *Acta Informatica*, no. 15, vol. 2, 1981, pp. 141145.
[4] A. Zelikovsky, "An 11/6-approximation Algorithm for the Network Steiner Problem," *Algorithmica*, no. 9, 1993, pp. 463470.
[5] D. Waitzman, C. Partridgeb, and S. Deering, "Distance Vector Multicast Routing Protocol (DVMRP)," Network Working Group, RFC-1075, November 1988.
[6] J. Moy, "Multicast Extensions to OSPF," RFC-1584, March 1994.
[7] Pedro M. Ruiz, Antonio F. Gomez-Skarmeta, "Approximating Optimal Multicast Trees in Wireless Multihop Networks," *10th IEEE Symposium on Computers and Communications*, 2005, pp. 686-691.
[8] QualNet Network Simulator version 3.9.1, http://www.scalable-networks.com/.
[9] GloMoSim Simulator, http://pcl.cs.ucla.edu/projects/glomosim/.
[10] TROPOS Networks, http://www.tropos.com/.
[11] C. Siva Ram Murthy and B.S. Manoj, Ad Hoc Wireless Networks: Architectures and Protocols, Prentice Hall, 2004.

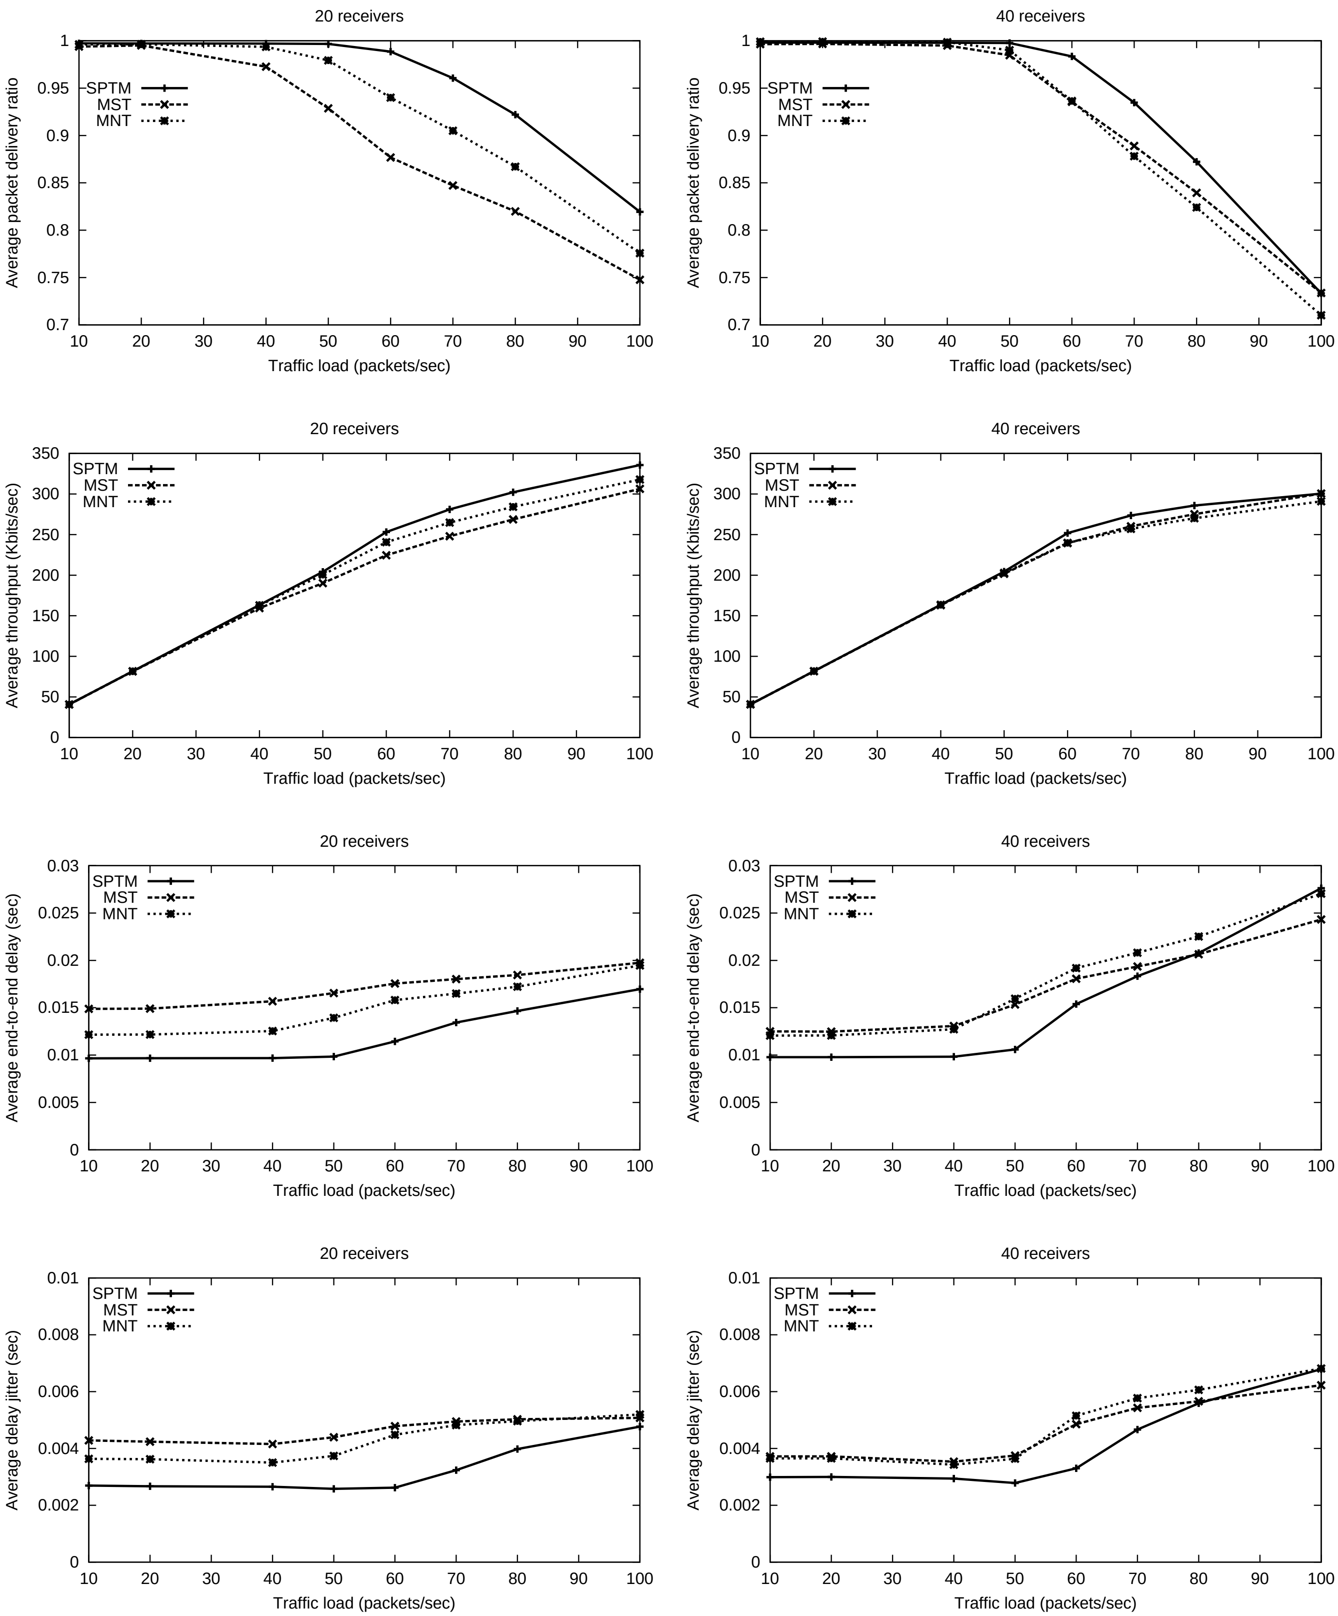

(a) One sender, 20 receivers

(b) One sender, 40 receivers

Fig. 4. Performance of multicast sessions in the network of 100 nodes

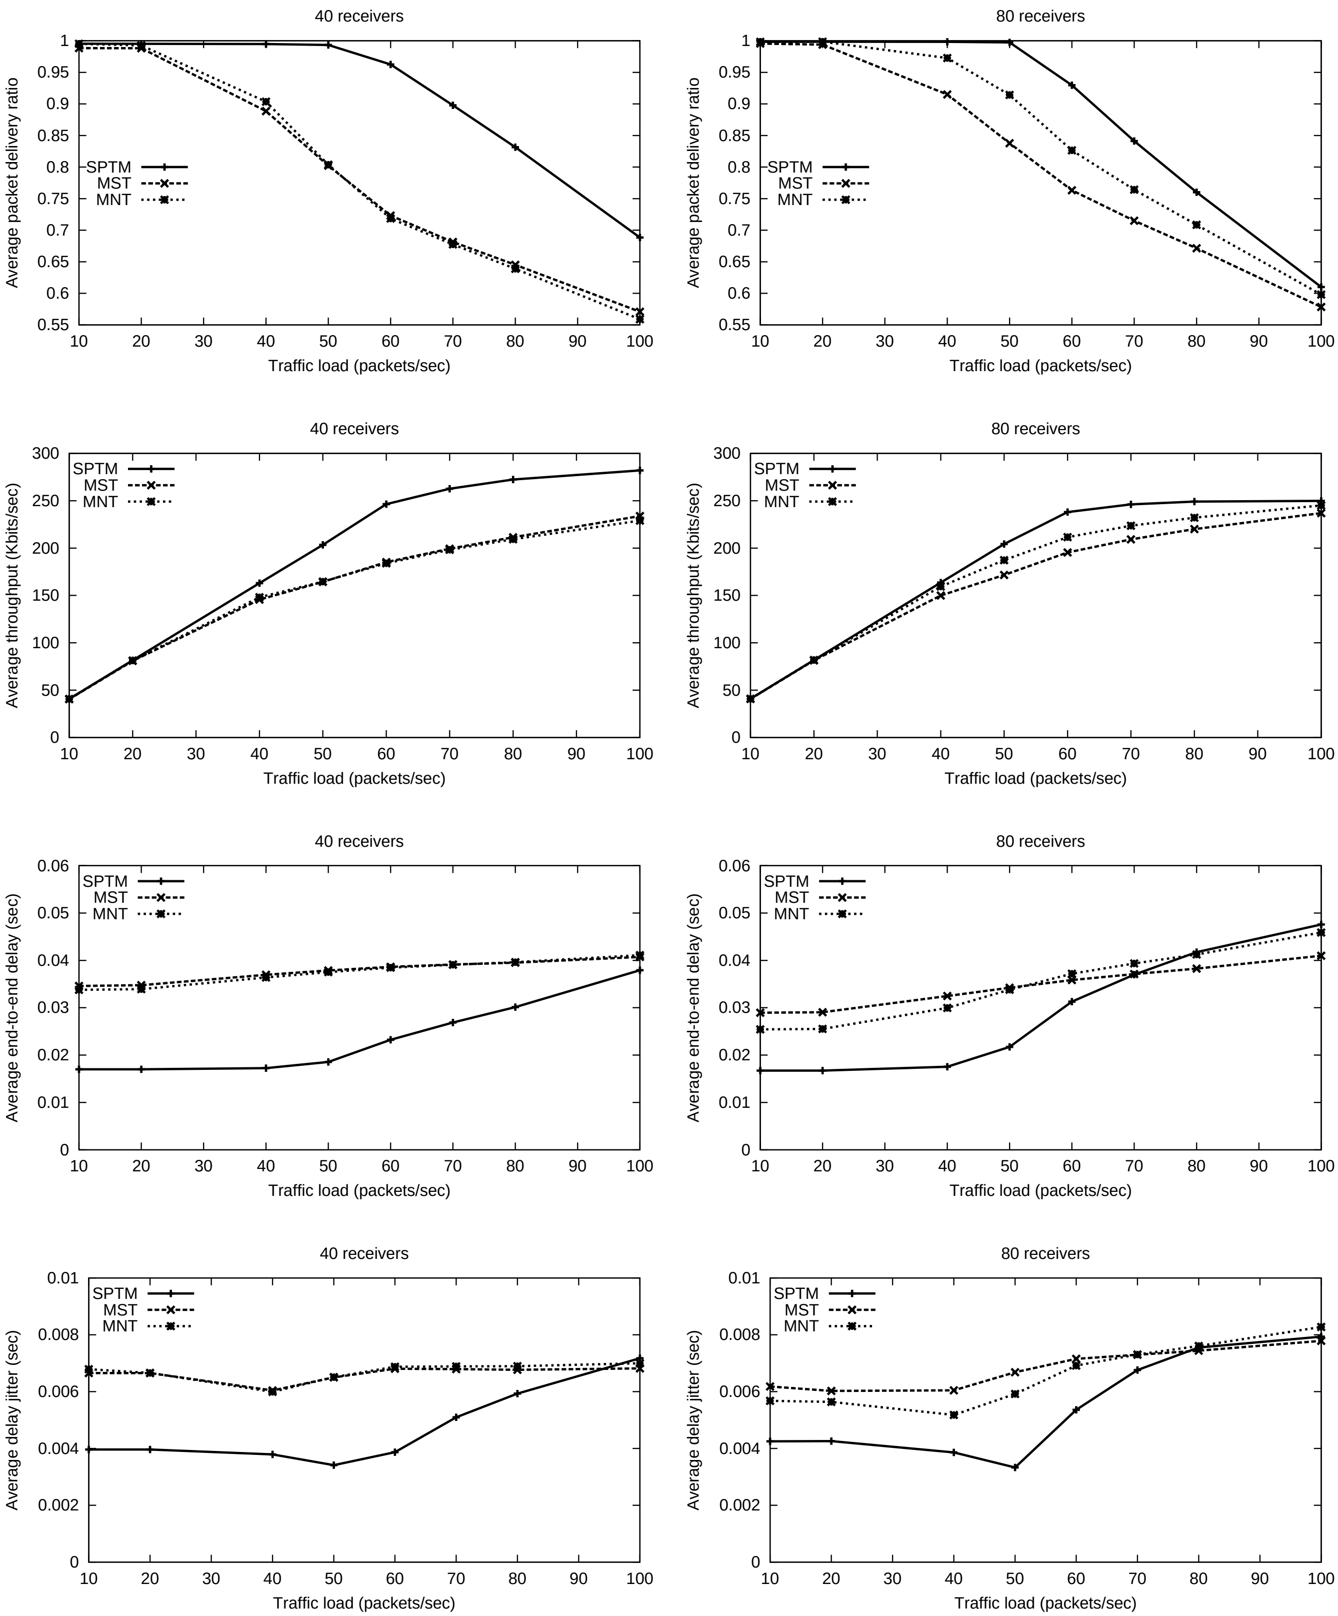

(a) One sender, 40 receivers

(b) One sender, 80 receivers

Fig. 5. Performance of multicast sessions in the network of 300 nodes

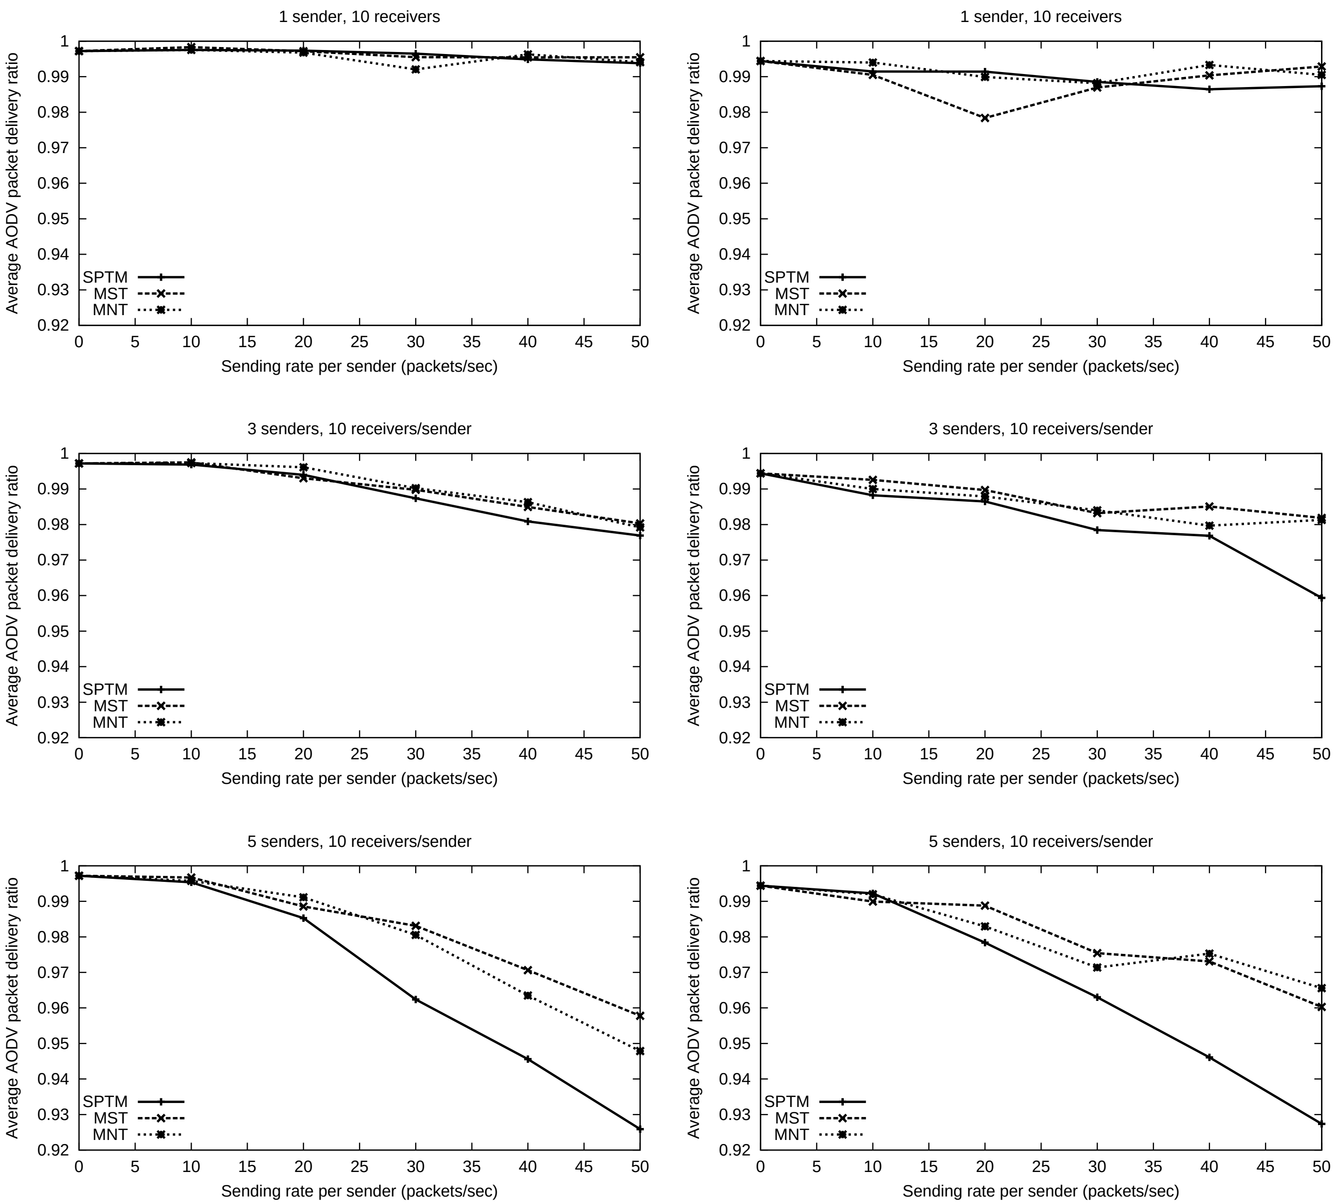

(a) Network of 100 nodes (b) Network of 300 nodes

Fig. 6. Impacts of multicast traffic on unicast flows