

On the Pitfalls of High-Throughput Multicast Metrics in Adversarial Wireless Mesh Networks

Jing Dong Reza Curtmola Cristina Nita-Rotaru
Department of Computer Science, Purdue University
{dongj,crix,crisn}@cs.purdue.edu

Abstract—Recent work in multicast routing for wireless mesh networks has focused on metrics that estimate link quality to maximize throughput. Nodes must collaborate in order to compute the path metric and forward data. The assumption that all nodes are honest and behave correctly during metric computation, propagation, and aggregation, as well as during data forwarding, leads to unexpected consequences in adversarial networks where compromised nodes act maliciously.

In this work we identify novel attacks against high-throughput multicast protocols in wireless mesh networks. The attacks exploit the local estimation and global aggregation of the metric to allow attackers to attract a large amount of traffic. We show that these attacks are very effective against multicast protocols based on high-throughput metrics. This leads us to conclude that aggressive path selection is a double-edged sword: it maximizes throughput, but in the absence of protection mechanisms it also increases attack effectiveness. Our approach to mitigate the identified attacks combines measurement-based detection and accusation-based reaction techniques. The solution also accommodates transient network variations and is resilient against attempts to exploit the defense mechanism itself. We demonstrate the attacks and our defense using ODMRP, a representative multicast protocol for wireless mesh networks, and SPP, an adaptation of the well-known ETX unicast metric to the multicast setting.

I. INTRODUCTION

Wireless mesh networks (WMNs) emerged as a promising technology that offers low-cost high-bandwidth community wireless services. A WMN consists of a set of stationary wireless routers that form a multi-hop backbone, and a set of mobile clients that communicate via the wireless backbone. Numerous applications envisioned to be deployed in WMNs, such as webcast, distance learning, online games, video conferencing, and multimedia broadcasting, follow a pattern where one or more sources disseminate data to a group of changing receivers. These applications can benefit from the service provided by multicast routing protocols.

Multicast routing protocols deliver data from a source to multiple destinations organized in a multicast group. In the last few years, several protocols [1]–[7] were proposed to provide multicast services for multi-hop wireless networks. Initially, these protocols were proposed for mobile ad hoc networks (MANETs), focusing primarily on network connectivity and using the number of hops between the source and receivers as the route selection metric. However, many of the applications that benefit from multicast services also have high-throughput requirements, and hop count was shown not to be the best metric for routing protocols that seek to maximize throughput [8], [9]. As a result, given the stationary nature and increased capabilities of nodes in mesh networks, recent research [10],

[11] focuses on maximizing path throughput by selecting paths based on metrics that capture the quality of the wireless links [9], [12]–[15]. We refer to such metrics as *link-quality* metrics or *high-throughput* metrics, and to protocols using such metrics as *high-throughput protocols*.

In a typical high-throughput multicast protocol, nodes periodically send probes to their neighbors to measure the quality of the links from their neighbors. During route discovery, a node estimates the cost of the path by combining its own measured metric of adjacent links with the route cost accumulated on the route discovery packet. The path with the best metric is then selected. Using high-throughput metrics requires the nodes to collaborate in order to derive the path metric. Although the selected path achieves a higher throughput, the assumption that all nodes are collaborative and behave correctly during metric computation and propagation leads to unexpected consequences in adversarial networks where some nodes act maliciously. In general, wireless networks are vulnerable to attacks coming from insiders and outsiders, due to the openness and shared nature of the medium, and the multi-hop characteristic of the communication. An aggressive path selection introduces new vulnerabilities and provides the attacker with an increased arsenal of attacks. For example, adversaries may manipulate the metrics in order to be selected on more paths and to draw more traffic, creating opportunities for attacks such as data dropping, mesh partitioning, or traffic analysis.

Previous work showed vulnerabilities of unicast routing protocols that use hop count as a metric. Several unicast routing protocols were proposed to cope with outsider attacks [16]–[19] or insider attacks [18], [20]–[23]. Secure wireless multicast is less studied [24], [25] and focuses primarily on tree-based protocols using hop count as a path selection metric.

In this work, we study the security implications of using high-throughput metrics. We focus on multicast in a wireless mesh network environment because it is a representative environment in which high-throughput metrics will be beneficial. Although the attacks we identify can also be conducted in unicast, the multicast setting makes them more effective and, at the same time, more difficult to defend against. We focus on mesh-based multicast protocols as they have the potential to be more resilient to attacks. We use ODMRP [5] as a representative protocol for wireless mesh networks and SPP [10], a metric based on the well-known ETX [9] unicast metric, as a high-throughput multicast metric. We selected SPP since it was shown [10] to outperform all the other multicast metrics for ODMRP. To the best of our knowledge, this is the first paper to

examine vulnerabilities of high-throughput metrics in general, and in multicast protocols for wireless mesh networks in particular. We summarize our contributions:

- We identify attacks against multicast protocols that exploit the use of high-throughput metrics. The attacks consist of *local metric manipulation* (LMM) and *global metric manipulation* (GMM), and allow an attacker to attract significant traffic. We show that aggressive path selection is a double-edged sword: it leads to throughput maximization, but in the absence of protection mechanisms it also increases attack effectiveness. Our simulations using the ODMRP protocol and the SPP metric indicate that the GMM attack is the most damaging to the network; it requires less than half the number of attackers required by simple data dropping attack to create the same disruption in the multicast service.

- We identify a dangerous effect of the attacks, referred to as *metric poisoning*, which causes many honest nodes to have incorrect metrics. Consequently, any response mechanism cannot rely on poisoned metrics for local recovery and must either use a fallback procedure not relying on the metric or refresh the metric before starting recovery. As a small number of attackers can severely impede the protocol, an effective solution must identify and isolate the malicious nodes.

- We propose a defense strategy that combines measurement-based detection and accusation-based reaction techniques to mitigate the identified attacks. To accommodate transient network variations and prevent attackers from exploiting the defense mechanism itself, we use temporary accusations and limit the number of accusations that can be generated by a node. The duration of an accusation is proportional to the disruption created by the accused node. Simulations with ODMRP and the SPP metric show that our strategy is very effective in defending against the attacks and has low overhead.

II. HIGH-THROUGHPUT MESH-BASED MULTICAST ROUTING

We consider a multi-hop wireless network where nodes participate in the data forwarding process for other nodes. We assume a mesh-based multicast routing protocol, which maintains a mesh connecting multicast sources and receivers. Path selection is performed based on a metric designed to maximize throughput. Below, we provide an overview of high-throughput metrics for multicast, then describe in details how such metrics are integrated with mesh-based multicast protocols.

A. High-Throughput Metrics

Traditionally, routing protocols have used hop count as a path selection metric. In static networks however, this metric was shown to achieve sub-optimal throughput because paths tend to include lossy wireless links [9], [26]. As a result, in recent years the focus has shifted toward high-throughput metrics that seek to maximize throughput by selecting paths based on the quality of wireless links (*e.g.*, ETX [9], PP [14], [26], RTT [13]). In such metrics, the quality of the links to/from a node's neighbors is measured by periodic probing. The metric for an entire path

is obtained by aggregating the metrics reported by the nodes on the path.

Several high-throughput metrics for multicast were proposed in [10]. All of these metrics are adaptations of unicast metrics to the multicast setting by taking into account the fundamental differences between unicast and multicast communication. Transmissions in multicast are less reliable than in unicast for several reasons. In unicast, a packet is sent reliably using link-layer unicast transmission, which involves link-layer acknowledgments and possibly packet retransmissions; in multicast, a packet is sent unreliably using link-layer broadcast, which does not involve link layer acknowledgments or data retransmissions. Moreover, unicast transmissions are preceded by a RTS/CTS exchange; in multicast there is no RTS/CTS exchange, which increases collision probability and decreases transmission reliability. Many metrics for unicast routing focus on minimizing the medium access time, while metrics for multicast focus on capturing in different ways the packet delivery ratio.

All the high-throughput multicast metrics proposed in [10] showed improvement over the original path selection strategy. The SPP metric [10], an adaptation of the well-known ETX [9] unicast metric, was shown to outperform the other multicast metrics [10], [27]. Thus, in the remainder of the paper and in our experimental evaluation, we consider SPP for demonstrative purposes. Below, we first give an overview of ETX, then show how it was extended to SPP.

ETX Metric. The ETX metric [9] was proposed for unicast and estimates the expected number of transmissions needed to successfully deliver a unicast packet over a link, including retransmissions. Each node periodically broadcasts probe packets which include the number of probe packets received from each of its neighbors over a time interval. A pair of neighboring nodes, A and B , estimate the quality of the link $A \leftrightarrow B$ by using the formula $ETX = \frac{1}{d_f \times d_r}$, where d_f and d_r are the probabilities that a packet is sent successfully from A to B (forward direction) and from B to A (reverse direction), respectively. The value of ETX for a path of k links between a source S and a receiver R is $ETX_{S \rightarrow R} = \sum_{i=1}^k ETX_i$, where ETX_i is the ETX value of the i -th link on the path; $ETX_{S \rightarrow R}$ estimates the total number of transmissions by all nodes on the path to deliver a packet from a source to a receiver.

SPP Metric. ETX was adapted to the multicast setting by Roy *et al.* in the form of the SPP metric [10]. The value of SPP for a path of k links between a source S and a receiver R is $SPP_{S \rightarrow R} = \prod_{i=1}^k SPP_i$, where the metric for each link i on the path is $SPP_i = d_f$ and d_f is defined as in ETX. The rationale for defining SPP as above is twofold:

- Unlike in unicast, where a successful transmission over a link depends on the quality of both directions of that link, in multicast only the quality of the forward direction matters because there are no link layer acknowledgments. The quality of a link $A \rightarrow B$, as perceived by node B , is $SPP_i = d_f$ and represents the probability that B receives a packet successfully from A over the link $A \rightarrow B$. Node B obtains d_f by counting the probes received from A over a fixed time interval.

–Also unlike unicast, in which the individual link metrics are summed, in multicast they are multiplied. This reflects the fact that for SPP the probability of a packet being delivered over a path from a source to a receiver is the product of the probabilities that the packet is successfully delivered to each of the intermediate nodes on the path. If any of the intermediate nodes fails to receive the packet, this causes the transmission for the entire route to fail, since there are no retransmissions. $SPP_{S \rightarrow R}$ (in fact $1/SPP_{S \rightarrow R}$) estimates the expected number of transmissions needed at the source to successfully deliver a packet from a source to a receiver.

SPP takes values in the interval $[0, 1]$, with higher metric values being better. In particular, $SPP = 1$ denotes perfect reliability, while $SPP = 0$ denotes complete unreliability.

B. High-Throughput Mesh-Based Multicast Routing

Multicast protocols provide communication from sources to receivers organized in groups by establishing dissemination structures such as trees or meshes, dynamically updated as nodes join or leave the group. Tree-based multicast protocols (e.g., MAODV [6]) build optimized data paths, but require more complex operations to create and maintain the multicast tree, and are less resilient to failures. Mesh-based multicast protocols (e.g., ODMRP [5]) build mesh structures that provide path redundancy making them more resilient to failures, but have higher overhead due to redundant retransmissions.

We focus on ODMRP as a representative mesh-based multicast protocol for wireless networks. Below we first give an overview of ODMRP, then describe how it can be enhanced with any link-quality metric. The protocol extension to use a high-throughput metric was first described by Roy *et al.* [10], [27]. We refer to the ODMRP protocol using a high-throughput metric as ODMRP-HT in order to distinguish it from the original ODMRP protocol.

ODMRP overview. ODMRP is an on-demand multicast routing protocol for multi-hop wireless networks, which uses a mesh of nodes for each multicast group. Nodes are added to the mesh through a route selection and activation protocol. The source periodically recreates the mesh by flooding a JOIN QUERY message in the network in order to refresh the membership information and update the routes. We use the term *round* to denote the interval between two consecutive mesh creation events. JOIN QUERY messages are flooded using a *basic flood suppression* mechanism, in which nodes only process the first received copy of a flood message.

When a receiver node gets a JOIN QUERY message, it activates the path from itself to the source by constructing and broadcasting a JOIN REPLY message that contains entries for each multicast group it wants to join; each entry has a *next hop* field filled with the corresponding upstream node. When an intermediate node receives a JOIN REPLY message, it detects that it is on the path to the source if the next hop field of any of the entries in the message matches its own identifier. If so, it makes itself a node part of the mesh (the FORWARDING GROUP) and creates and broadcasts a new JOIN REPLY built upon the matched entries.




Fig. 1. An example of ODMRP-HT mesh creation for a multicast group with 2 sources (S_1, S_2) and 6 receivers (R_1, \dots, R_6). The label on each link represents the value of the link’s SPP metric. Note that receivers can also act as forwarding nodes for other nodes (e.g., node R_5).

Once the JOIN REPLY messages reach the source, the multicast receivers become connected to the source through a mesh of nodes (the FORWARDING GROUP) which ensures the delivery of multicast data. While a node is in the FORWARDING GROUP, it rebroadcasts any non-duplicate multicast data packets that it receives.

ODMRP takes a “soft state” approach in that nodes put a minimal effort to maintain the mesh. To leave the multicast group, receiver nodes are not required to explicitly send any message, instead they do not reply to JOIN QUERY messages. Also, a node’s participation in the FORWARDING GROUP expires if its forwarding-node status is not updated.

ODMRP-HT. We now describe how ODMRP can be enhanced with high-throughput metrics. The main differences between ODMRP-HT and ODMRP are: (1) instead of selecting routes based on minimum delay (which results in choosing the fastest routes), ODMRP-HT selects routes based on a link-quality metric, and (2) ODMRP-HT uses a *weighted flood suppression* mechanism to flood JOIN QUERY messages instead of a basic flood suppression.

As required by the link-quality metric, each node measures the quality of the links from its neighbors to itself, based on the periodic probes sent by its neighbors. The JOIN QUERY message is flooded periodically by a source S and contains a *route cost* field which accumulates the metric for the route on which the message travelled. Upon receiving a JOIN QUERY, a node updates the route cost field by accumulating the metric of the last link travelled by the message. Because different paths may have different metrics, JOIN QUERY messages are flooded using a *weighted flood suppression* mechanism, in which a node processes flood duplicates for a fixed interval of time and rebroadcasts flood messages that advertise a better metric (indicated by the route cost field)¹. Each node also records the node from which it received the JOIN QUERY with the best quality metric as its *upstream* node for the source S .

After waiting for a fixed interval of time, during which it may receive several JOIN QUERY packets that contain different route metrics, a multicast receiver records as its upstream for source S the neighbor that advertised the JOIN QUERY with the best metric. Just like in ODMRP, the receiver then

¹Several studies [25], [27] show that the overhead caused by rebroadcasting some of the flood packets is reasonable, validating the effectiveness of this weighted flood suppression strategy.

constructs a JOIN REPLY packet, which will be forwarded towards the source on the optimal path as defined by the metric and will activate the nodes on this path as part of the FORWARDING GROUP. In Fig. 1 we give an example of how ODMRP-HT selects the mesh of nodes in the FORWARDING GROUP based on the SPP link-quality metric.

III. ATTACKS AGAINST HIGH-THROUGHPUT MULTICAST

We present several attacks against high-throughput multicast protocols. The attacks exploit vulnerabilities introduced by the use of high-throughput metrics. They require little resource from the attacker, but can cause severe damage to the performance of the multicast protocol. We first present the adversarial model, followed by the targets and the details of the attacks.

A. Adversarial Model

Malicious nodes may exhibit Byzantine behavior, either alone or in collusion with other malicious nodes. We refer to any arbitrary action by authenticated nodes deviating from protocol specification as Byzantine behavior, and to such an adversary as a Byzantine adversary. Examples of Byzantine behavior include: dropping, injecting, modifying, replaying, or rushing packets, and creating wormholes.

This work considers attacks that target the network level and assumes that adversaries do not have control on lower layers such as the physical or MAC layers. We assume the physical layer can use jamming-resilient techniques such as direct sequence spread spectrum (DSSS) or frequency hopping spread spectrum (FHSS) (as in the case of 802.11). We do not consider Sybil attacks and we do not prevent traffic analysis. Instead, the goal of this work is to achieve survivable routing.

B. Attack Goals

In this work, we primarily focus on attacks that aim to disrupt the multicast data delivery. The two main attack targets that allow the attacker to achieve this goal are the path establishment and data forwarding phases of the protocol.

Path establishment attacks prevent receivers from connecting to multicast sources. In ODMRP-HT, since each receiver only activates a single path to each source, an attacker lying on that path can prevent path establishment by dropping the JOIN REPLY message. *Data forwarding attacks* disrupt the routing service by dropping data packets. In both cases, the attack effectiveness is directly related to the attackers' ability to control route selection and to be selected on routes. Traditionally, such ability can be achieved via wireless-specific attacks such as rushing and wormholes. The use of high-throughput metrics gives attackers additional opportunities to be included in the mesh by manipulating the routing metric. Rushing and wormholes are general attacks against wireless routing protocols that have been studied extensively [28]–[31]. Thus, below we focus on metric manipulation attacks, which require a little effort to execute, yet are extremely detrimental to the protocol performance.




Fig. 2. Metric manipulation attack during the propagation of the flood packet from the source S to receiver R . A label above a link is the link's real SPP metric; a label below a link is the link's metric falsely claimed by a node executing a LMM attack; a label below a node is the accumulated route metric advertised by the node.

C. Metric Manipulation Attacks

As discussed in Section II, multicast protocols using high-throughput metrics prefer paths to the source that are perceived as having high-quality, while trying to avoid low-quality paths. Thus, a good strategy for an attacker to increase its chances of being selected in the FORWARDING GROUP is to advertise artificially good metrics for routes to the source.

The use of high-throughput metrics requires each node to collect *local* information about its adjacent links based on periodic probes from its neighbors. This local information is accumulated in JOIN QUERY packets and propagated in the network, allowing nodes to obtain *global* information about the quality of the routes from the source. Consequently, adversaries can execute two types of metric manipulation attacks: *local metric manipulation* (LMM) and *global metric manipulation* (GMM). These attacks are Byzantine in nature, as they are conducted by nodes which are authorized to participate in the routing protocol.

Local Metric Manipulation (LMM) Attacks. An adversarial node artificially increases the quality of its adjacent links, distorting the neighbors' perception about these links. The falsely advertised "high-quality" links will be preferred and malicious nodes have better chances to be included on routes.

A node can claim a false value for the quality of the links towards itself. In Fig. 2 a malicious node C_1 claims that $SPP_{B_1 \rightarrow C_1} = 0.9$ instead of the correct metric $SPP_{B_1 \rightarrow C_1} = 0.6$. Thus, C_1 accumulates a false local metric for the link $B_1 \rightarrow C_1$ and advertises to R the metric $SPP_{S \rightarrow C_1} = 0.9$ instead of the correct metric $SPP_{S \rightarrow C_1} = 0.6$. The route $S-A_1-B_1-C_1-R$ will be chosen over the correct route $S-A_3-B_3-C_3-R$.

Global Metric Manipulation (GMM) Attacks. In a GMM attack, a malicious node arbitrarily changes the value of the route metric accumulated in the flood packet, before rebroadcasting this packet. A GMM attack allows a node to manipulate not only its own contribution to the path metric, but also the contributions of previous nodes that were accumulated in the path metric. For example, in Fig. 2 attacker C_2 should advertise a route metric of 0.25, but instead advertises a route metric of 0.9 to node R . This causes the route $S-A_2-B_2-C_2-R$ to be selected over the correct route $S-A_3-B_3-C_3-R$.

Impact of metric manipulation attacks on routing. The attacks we described allow attackers to attract and control traffic. In addition, the epidemic nature of metric derivation causes an epidemic attack propagation, which "poisons" the




Fig. 3. Metric manipulation attack in a network with one source (S), two receivers (R_1, R_2) and one attacker (A). The label on each link represents the value of the link’s SPP metric.

metrics of many nodes in the network. We exemplify both of these effects with the scenario in Fig. 3.

When no attackers are present (Fig. 3(a)), nodes B, C and D are activated as part of the FORWARDING GROUP. Consider that node A executes a metric manipulation attack (Fig. 3(b)): Upon receiving the JOIN QUERY, node A changes the metric and advertises a perfect metric with value 1. Consequently, both receivers R_1 and R_2 are “attracted” to it and only nodes B and C will be selected as part of the FORWARDING GROUP. The net effect is that both R_1 and R_2 are denied service since they do not have a path to the source.

The false metric advertisement by node A also poisons the metrics of many nodes in the network. For example, node C derives an incorrect metric of 0.9, and then propagates it to its neighbors, causing them to derive an incorrect metric as well. Besides distorting path establishment for data delivery, a severe side effect of the attack is that it introduces a significant challenge for attack recovery. For example, even if A ’s neighbors are able to detect A is an attacker, they cannot rely on the metric to find a new route to the source. Indeed, if node C detects that A is malicious, it can try to activate nodes F or B , which advertised the second best metric; however, routes through either F or B lead back to the attacker. The problem stems from the fact that the metric cannot be relied upon and nodes do not know the right direction to “break free” from the attraction of A . Hence, we make the observation that defense mechanisms cannot rely on the existing metric for recovery and have to either resort to a fallback procedure not using the metric or refresh the metric before starting recovery.

IV. SECURE HIGH-THROUGHPUT MULTICAST ROUTING

In this section, we present our secure multicast routing protocol (S-ODMRP) that accommodates high-throughput metrics.

A. Authentication Framework

We assume that each user authorized to be part of the mesh network has a pair of public and private keys and a *client certificate* that binds its public key to a unique user identifier. This defends against external attacks from users that are not part of the network. We assume source data is authenticated, so that receivers can distinguish authentic data from spurious data. Efficient source data authentication can be achieved with existing schemes such as TESLA [32]. Finally, we assume the existence of a secure neighbor discovery scheme [33].

B. S-ODMRP Overview

Our approach relies on the observation that regardless of the attack strategy, attackers do not affect the multicast protocol unless they cause a drop in the packet delivery ratio (PDR). We adopt a reactive approach in which attacker nodes are detected through a *measurement-based detection* protocol component, and then isolated through an *accusation-based reaction* protocol component. We next describe these two components.

Measurement-based attack detection. Whether by packet dropping alone or by combining it with metric manipulation to attract routes, the effect of an attack is that data is not delivered at a rate consistent with the advertised path quality. We propose a generic attack detection strategy that relies on the ability of honest nodes to detect the discrepancy between the *expected* PDR (ePDR) and the *perceived* PDR (pPDR). A node can estimate the ePDR of a route from the value of the metric for that route². For both variable and constant data rate source, the pPDR of a route can be determined by examining the continuity of the sequence number in received data packets, for example, by dividing the number of received packets by the number of packets sent by the source (derived from packet sequence numbers) over an interval of time. To allow nodes to determine pPDR even if the attacker nodes drops all data packets, we also include in JOIN QUERY packets the current sequence number authenticated by the source.

Nodes in the FORWARDING GROUP and receiver nodes monitor the data rate of their upstream node. If $|ePDR - pPDR|$ for a route becomes larger than a threshold δ , then nodes suspect that the route is under attack because the route failed to deliver data at a level consistent with its claimed quality³.

Accusation-based attack reaction. We use a *controlled-accusation* mechanism in which a node that detects malicious behavior, temporarily accuses the suspected node by flooding in the network an ACCUSATION message containing the identity of the accused and the accuser nodes, as well as the duration of the accusation. As long as the accusation is valid, metrics advertised by the accused node will be ignored and the node will not be selected as part of the FORWARDING GROUP. This strategy also successfully handles attacks against path establishment. From the downstream node point of view, the dropping of a JOIN REPLY message causes exactly the same effect as the attacker dropping all data packets, thus the downstream nodes will react and accuse the attacker.

We use a temporary accusation strategy to cope with transient network variations: The accusation duration is calculated proportional to the observed discrepancy between ePDR and pPDR, so that accusations caused by metric inflation and malicious data dropping last longer, while accusations caused by transient network variations last shorter.

To prevent the abuse of the accusation mechanism by attackers, a node is not allowed to issue a new accusation

²For the SPP metric, a route’s ePDR is equal to the route’s metric.

³Note that this degradation in the route’s quality may also be caused by natural losses. We do not differentiate between losses caused by adversarial behavior and natural losses because lossy links must be avoided anyway. Moreover, it is difficult to differentiate between these two types of losses.

before its previously issued accusation expires. Accused nodes can still act as receivers even though they are excluded from the FORWARDING GROUP. Finally, we address attackers that strategically accuse certain nodes in order to disconnect the network (note that such attacks are effective and easy to execute). We make one exception from the rule that only non-accused nodes are included in the FORWARDING GROUP: If the best metric is advertised by an accused neighbor, a node also activates this neighbor (by sending a JOIN REPLY) in addition to the best non-accused neighbor. This ensures that good paths may still be utilized, even if honest nodes on these paths are falsely accused. In Sec. V-E, we show that this strategy only adds a very low overhead.

C. S-ODMRP Detailed Description

1) *Mesh Creation*: S-ODMRP mesh creation follows the same pattern of ODMRP-HT presented in Sec. II-B. The source node S periodically broadcasts to the entire network a JOIN QUERY message in order to refresh the membership information and to update the routes. The JOIN QUERY message is signed by S and is propagated using a weighted flood suppression mechanism. Nodes only process JOIN QUERY messages that have valid signatures and that are received from nodes not currently accused (indicated by an ACCUSATION TABLE maintained by each node). Nodes record the upstream node and the metric corresponding to the route with the best metric as bestUpstream and bestMetric.

The JOIN REPLY messages are then sent from receivers back to S along optimal paths as defined by the high-throughput metric, leading to the creation of the FORWARDING GROUP (the multicast mesh). After sending a JOIN REPLY to its bestUpstream, a node starts to monitor the PDR from its bestUpstream in order to measure its perceived PDR (pPDR).

2) *Attack Detection and Reaction*: FORWARDING GROUP nodes and receiver nodes use the previously described measurement-based mechanism to detect data-dropping attacks. When a node detects attack behavior (*i.e.*, $|ePDR - pPDR| > \delta$), it starts a *React_Timer* whose timeout is inversely proportional to its estimated PDR (ePDR). Since ePDR decreases monotonically along a multicast data path, nodes farther away from the source will have a larger timeout value for the *React_Timer*. This staggered timeout technique ensures nodes immediately below the attacker will take action first, before any of their downstream nodes mistakenly accuse their upstream node. When the *React_Timer* of a node N expires, N accuses its bestUpstream node and cancels the *React_Timer* on its downstream nodes with the following actions:

- create, sign, and flood an ACCUSATION message in the network, which contains N 's identity and the identity of N 's bestUpstream node. The message also contains a value accusationTime proportional to $|ePDR - pPDR|$, indicating the amount of time the accusation lasts.

- create, sign, and send to its downstream nodes a RECOVERY message, which contains the ACCUSATION message. This message serves the role of canceling the *React_Timer* of nodes in

N 's subtree and activating the fallback procedure on receivers in N 's subtree (see Sec. IV-C3).

Upon receipt of an ACCUSATION message, a node checks if the signature on the message is valid and if it does not have an unexpired accusation from the same accuser node. If both checks pass, the node adds a corresponding entry to its ACCUSATION TABLE. Accusations are removed from the ACCUSATION TABLE after the accusationTime has elapsed.

Upon receipt of a RECOVERY message rr from its bestUpstream node, a FORWARDING GROUP node N verifies the signature on the message and checks that it does not have an unexpired accusation from the same accuser node (except for the duplicate from the flooded ACCUSATION message). In addition, to prevent colluding attackers from accusing each other only for a short amount of time even though they cause a large PDR drop, the node also checks that the accusationTime in the message is at least as much as its own observed discrepancy (the $|ePDR - pPDR|$ value). If all checks pass, it cancels its pending *React_Timer*, forwards rr to its downstream nodes, and if it is a receiver, activates the recovery procedure (see below).

3) *Fallback Recovery*: When an attack is detected during a round, the receiver nodes in the subtree of the attacker need to find alternative routes to “salvage” data for the rest of the round. As shown in Sec. III-C, a side effect of metric manipulation attacks is *metric poisoning*, which prevents recovery by relying on the metrics in the round the attack is detected. We address this inability by falling back to the fastest route for routing during the remainder of the round⁴. Specifically, during the JOIN QUERY flooding, besides recording the bestUpstream node, each node also records the upstream for the fastest route as fastestUpstream. To recover from an attack, a receiver sends a special JOIN REPLY message to its fastestUpstream node. Each node on the fastest route forwards the special JOIN REPLY message to their fastestUpstream node and becomes part of the FORWARDING GROUP.

V. EXPERIMENTAL EVALUATION

In this section, we demonstrate through experiments the vulnerability of metric enhanced multicast protocol by examining the impact of different attacks, and investigate the effectiveness of our defense mechanisms and its associated overhead.

A. Experimental Methodology

Simulation Setup. We implemented ODMRP-HT and S-ODMRP using the ODMRP version available in the Glosim [34] simulator. Nodes were set to use 802.11 radios with 2 Mbps bandwidth and 250m nominal range. We simulate environments representative of mesh networks deployments by using the two-ray radio propagation model with the Rayleigh loss model, which models environments with large reflectors, *e.g.*, trees and buildings, where the receiver is not in the line-of-sight of the sender.

⁴This strategy is not attack-proof, as the fastest route may include malicious nodes. However, since the route is only used for the remainder of the round, it is preferable to use an efficient procedure than to find adversarial-free path, which is itself a challenging task and usually requires expensive protocols [23].




Fig. 4. The effectiveness of metric attacks on ODMRP-HT. For comparison we also depict attacks against ODMRP which does not use high-throughput metrics.

The network consists of 100 nodes randomly placed in a $1500\text{m} \times 1500\text{m}$ area. We randomly select 20 nodes as multicast group members and among the 20 members one node is randomly selected as the data source. Group members join the group at the beginning of the experiment. At second 100, the source starts to multicast data packets for 400 seconds at a rate of 20 packets per second, each packet of 512 bytes. When attackers are present, they are randomly selected among nodes that are not group members. For S-ODMRP, we use RSA signatures with 1024-bit keys, simulating delays to approximate the performance of a 1.3 GHz Intel Centrino processor. We empirically tune the threshold $\delta = 20\%$ to accommodate random network variations in the simulated scenarios. The timeout for *React_Timer* is set as $20(1 - \text{ePDR})$ millisecond, and the accusationTime is set as $250(\text{ePDR} - \text{pPDR})$ second.

We used the SPP high-throughput metric, configured with optimal parameters as recommended in [10]. Data points are averaged over 10 different random environments and over all group members.

Attack Scenarios. We consider the following scenarios:

- *No-Attack*: The attackers do not perform any action in the network. This represents the ideal case where the attackers are identified and completely isolated in the network, and serves as the baseline for evaluating the impact of the attack and the performance of our defense.
- *Drop-Only*: The attackers drop data packets, but participate in the protocol correctly otherwise. The attack has effect only when attackers are selected in the FORWARDING GROUP. We use this scenario to demonstrate that metric manipulation amplifies data dropping attacks.
- *LMM-Drop*: The attackers combine local metric manipulation (LMM) with the data dropping attack. The attackers conduct the LMM attack by re-advertising the same metric they received in JOIN QUERY, which is equivalent to making their link metric of the previous hop equal to 1 (best).
- *GMM-Drop*: The attackers combine global metric manipulation (GMM) with the data dropping attack. The attackers conduct the GMM attack by re-advertising a metric of 1 (best) after receiving a JOIN QUERY.
- *False-Accusation*: The attackers exploit our accusation mechanism by falsely accusing random a honest node at startup for the whole experiment period in order to reduce the PDR. Due to space constraint, we do not present results for attacks that aim to cause large bandwidth overhead through frequent flooding of accusation messages using false accusations. We

can upper bound the frequency of the accusation message flooding from any attacker node to only once a few seconds by imposing a lower bound on the accusation timeout, thus the inflation of overhead is limited.

Metrics. We measure the performance of data delivery using the packet delivery ratio (PDR), defined as $\text{PDR} = n_r/n_s$, where n_r is the average number of packets received by all receivers and n_s is the number of packets sent by the source.

We also measure the strength of the attacks using as metric the PDR Decrease Ratio (PDR-DR), defined as

$$\text{PDR-DR} = \frac{\text{PDR}_{\text{noattack}} - \text{PDR}_{\text{attack}}}{\text{PDR}_{\text{noattack}}},$$

where $\text{PDR}_{\text{attack}}$ and $\text{PDR}_{\text{noattack}}$ represent the PDR when the network is under attack and not under attack, respectively.

The overhead of our defense consists of three components, the control bandwidth overhead due to additional messages and larger message size (e.g. accusation messages, signatures on query messages), the computational overhead due to cryptographic operations, and the additional data packet transmissions caused by our protocol. We measure the control bandwidth overhead per node, defined as the total control overhead divided by the number of nodes. The computational overhead is measured as the number of signatures performed by each node per second. To measure redundant data packet transmissions, we define *data packet transmission efficiency* as the total number of data packets transmitted by all nodes in the network divided by the total number of data packets received by all receivers. Thus, data packet transmission efficiency captures the cost (number of data packet transmissions) per data packet received.

B. Effectiveness of Metric Manipulation Attacks

Fig. 4(a) shows the impact of *Drop-Only* attack on the original ODMRP (not using high-throughput metric). The protocol is quite resilient to attacks, *i.e.*, PDR decreases by only 15% for 20 attackers. This reflects the inherent resiliency of mesh based multicast protocols against packet dropping, as typically a node has multiple paths to receive the same packet.

Fig. 4(b) shows the PDR of the protocol when using high-throughput metric (ODMRP-HT) under different types of attacks. We observe that with the *Drop-Only* attack, the PDR drops quickly to a level below the case when no high throughput metric is used. Thus, simple packet dropping completely nullifies the benefits of high throughput metrics. By manipulating the metrics as in *LMM-Drop* and *GMM-Drop*, the attacker can inflict a much larger decrease in PDR. For example, the PDR




Fig. 5. The effectiveness of S-ODMRP for different attacks

decreases from 72% to only 25% for 10 attackers using *GMM-Drop*, in contrast to 55% for *Drop-Only*. Fig.4(c) compares the impact of the attack in terms of the PDR Decrease Ratio. We see that metric manipulation significantly increases the attack strength. For example, with 10 attackers, the PDR-DR of *GMM-Drop* (68%) is more than double the PDR-DR of *Drop-Only* (32%). Thus, we conclude that metric manipulation attacks pose a severe threat to high-throughput protocols.

C. Effectiveness of the Defense

In Fig. 5 we show the effectiveness of our defense (S-ODMRP) against different types of attacks, compared to the insecure ODMRP-HT protocol. S-ODMRP suffers only a small PDR decrease relative to the baseline *No-Attack* case. For example, a total of 20 attackers causes a PDR drop of only 12%, considerably smaller than the case without defense, which shows a PDR decrease by as much as 55% in the *GMM-Drop* attack. To rule out random factors, we performed a paired t-test [35] on the results showing that S-ODMRP improves the PDR for all attack types, with P-value less than 2.2×10^{-16} . For 10 attackers, S-ODMRP improves the PDR of ODMRP-HT for *Drop-Only*, *LMM-Drop* and *GMM-Drop* by at least 4.5%, 16.7%, 33%, with 95% confidence level. Thus, our defense is very effective against all the attacks. The small PDR decrease for S-ODMRP can be attributed to two main factors. First, common to all reactive schemes, attackers can cause some initial damage, before action is taken against them. Second, as the number of attackers increases, some receivers become completely isolated and are not able to receive data.

Fig. 5 also shows an interesting phenomenon: The PDR decrease for S-ODMRP is similar for all attacks, despite the varying strength of the attacks. This outcome reflects the design of our defense mechanism in which accusations last proportional to the discrepancy between ePDR and pPDR: Attacks that cause a small discrepancy (e.g., *Drop-Only*) are forgiven sooner and can be executed again, while attacks that cause a large discrepancy (e.g., *GMM-Drop*) result in a more severe punishment and can be executed less frequently.

D. Defense Resiliency to Attacks

Attackers may attempt to exploit the accusation mechanism in S-ODMRP. Fig. 6 shows that S-ODMRP is very resilient against the *False-Accusation* attack, in which attackers falsely accuse one of their neighbors. This comes from the controlled nature of accusations, which only allows an attacker to accuse one honest node at a time. Also, as described in Sec. IV-B,




Fig. 6. Impact of the *False-Accusation* attack on S-ODMRP

falsely accused nodes that advertise a good metric may continue to forward data.

E. Overhead of S-ODMRP

Fig. 7(a) and 7(b) show the control bandwidth and computational overhead for S-ODMRP. We observe that for all attack configurations, the bandwidth and computational overhead are maintained at a stable low level of around 0.95 kbps and 0.9 signatures per node per second. To understand the source of the overhead better, we analyzed different components of the overhead. The result shows that the overhead due to reacting to attackers (such as creation and dissemination of ACCUSATION and RECOVERY messages) is negligible, since the attackers, once detected, are accused for a relatively long period of time, thus the frequency of attack reaction is low. The bulk of the overhead comes from the periodic network-wide flooding of authenticated JOIN QUERY packets. Since query flooding is common in all scenarios, we obtain a similar level of overhead across different scenarios. The reason for the slight overhead decrease for an increasing number of attackers for the *False-Accusation* attack is that JOIN QUERY from the falsely accused honest nodes are ignored by their neighbors, resulting a smaller number of transmissions of JOIN QUERY packets.

In Fig. 7(c), we notice that S-ODMRP under various attacks even improves slightly the *data transmission efficiency* of ODMRP-HT with no attacks. This apparent anomaly can be explained because in S-ODMRP nodes that are further away from the source are more likely to be affected by attacks and these are the nodes that require more transmissions to receive data packets.

VI. RELATED WORK

Work studying multicast routing specific security problems in wireless networks is scarce with the notable exception of the authentication framework by Roy *et al.* [24] and BSMR [25] which focus on outsider and insider attacks for the well-known tree-based MAODV multicast protocol.




Fig. 7. The overhead of S-ODMRP

Significant work focused on the security of unicast wireless routing protocols. Several secure routing protocols resilient to outside attacks were proposed in the last few years such as Ariadne [18], SEAD [17], ARAN [19], and the work in [16].

Wireless specific attacks such as flood rushing and wormhole were identified and studied. RAP [28] prevents the rushing attack by waiting for several flood requests and then randomly selecting one to forward, rather than always forwarding only the first one. Techniques to defend against wormhole attacks include *Packet Leashes* [29] which restricts the maximum transmission distance by using time or location information, Truelink [30] which uses MAC level acknowledgments to infer if a link exists or not between two nodes, and the work in [31], which relies on directional antennas.

The problem of insider threats in unicast was studied in [18], [20]–[23]. Watchdog [20] detects if a node is adversarial by using an approach where a node monitors its neighbors if they forward packets to other destinations. SDT [21] and Ariadne [18] use multi-path routing to prevent a malicious node from selectively dropping data. ODSBR [22], [23] provides resilience to colluding Byzantine attacks by detecting malicious links based on an acknowledgment-based feedback technique.

VII. CONCLUSION

We considered the security implication of using high throughput metrics in multicast protocols in wireless mesh networks. In particular, we identified metric manipulation attacks that can inflict significant damage on the network. The attacks not only have a direct impact on the multicast service, but also raise additional challenges in defending against them due to their metric poisoning effect. We overcome the challenges with our novel defense scheme that combines measurement-based attack detection and accusation-based reaction. Our defense also copes with transient network variations and malicious attempts to attack the network indirectly by exploiting the defense itself. We demonstrate through experiments that our defense is effective against the identified attacks, resilient to malicious exploitations, and imposes a small overhead.

REFERENCES

- [1] Y. B. Ko and N. H. Vaidya, "Flooding-based geocasting protocols for mobile ad hoc networks," *Mob. Netw. Appl.*, vol. 7, no. 6, 2002.
- [2] R. Chandra, V. Ramasubramanian, and K. Birman, "Anonymous gossip: Improving multicast reliability in mobile ad-hoc networks," in *Proc. of ICDCS*, 2001.
- [3] Y.-B. Ko and N. H. Vaidya, "GeoTORA: a protocol for geocasting in mobile ad hoc networks," in *Proc. of ICNP*. IEEE, 2000, p. 240.
- [4] E. L. Madruga and J. J. Garcia-Luna-Aceves, "Scalable multicasting: the core-assisted mesh protocol," *Mob. Netw. Appl.*, vol. 6, no. 2, 2001.

- [5] S. J. Lee, W. Su, and M. Gerla, "On-demand multicast routing protocol in multihop wireless mobile networks," *Mob. Netw. Appl.*, 2002.
- [6] E. M. Royer and C. E. Perkins, "Multicast ad-hoc on-demand distance vector (MAODV) routing," in *Internet Draft*, July 2000.
- [7] J. G. Jetcheva and D. B. Johnson, "Adaptive demand-driven multicast routing in multi-hop wireless ad hoc networks," in *Proc. of MobiHoc*, 2001, pp. 33–44.
- [8] H. Lundgren, E. Nordstrom, and C. Tschudin, "Coping with communication gray zones in IEEE 802.11b based ad hoc networks," in *Proc. of WOWMOM '02*. ACM Press, 2002, pp. 49–55.
- [9] D. S. J. D. Couto, D. Aguayo, J. C. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," in *Proc. of MOBICOM '03*. ACM, 2003, pp. 134–146.
- [10] S. Roy, D. Koutsonikolas, S. Das, and C. Hu, "High-throughput multicast routing metrics in wireless mesh networks," in *Proc. of ICDCS '06*, 2006.
- [11] A. Chen, D. Lee, G. Chandrasekaran, and P. Sinha, "HIMAC: High throughput MAC layer multicasting in wireless networks," in *Proc. of Mobile Adhoc and Sensor Systems (MASS '06)*, October 2006.
- [12] B. Awerbuch, D. Holmer, and H. Rubens, "The medium time metric: High throughput route selection in multirate ad hoc wireless networks," *MÖNET, Spec. Iss. on Internet Wireless Access: 802.11 and Beyond*, 2005.
- [13] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou, "A multi-radio unification protocol for IEEE 802.11 wireless networks," in *Proc. of BroadNets '04*, 2004.
- [14] S. Keshav, "A control-theoretic approach to flow control," *Proc. of the Conference on Communications Architecture and Protocols*, 1993.
- [15] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *Proc. of MOBICOM '04*. ACM, 2004.
- [16] P. Papadimitratos and Z. Haas, "Secure routing for mobile ad hoc networks," in *Proc. of CNDS*, January 2002, pp. 27–31.
- [17] Y.-C. Hu, D. B. Johnson, and A. Perrig, "SEAD: Secure efficient distance vector routing for mobile wireless ad hoc networks," in *Proc. of WMCSA*, June 2002.
- [18] Y.-C. Hu, A. Perrig, and D. B. Johnson, "Ariadne: A secure on-demand routing protocol for ad hoc networks," in *Proc. of MOBICOM*, 2002.
- [19] K. Sanzgiri, B. Dahill, B. N. Levine, C. Shields, and E. Belding-Royer, "A secure routing protocol for ad hoc networks," in *Proc. of ICNP*, 2002.
- [20] S. Marti, T. Giuli, K. Lai, and M. Baker, "Mitigating routing misbehavior in mobile ad hoc networks," in *Proc. of MOBICOM*, August 2000.
- [21] P. Papadimitratos and Z. Haas, "Secure data transmission in mobile ad hoc networks," in *Proc. of WiSe*, 2003, pp. 41–50.
- [22] B. Awerbuch, D. Holmer, C. Nita-Rotaru, and H. Rubens, "An on-demand secure routing protocol resilient to byzantine failures," in *Proc. of WiSe '02*. ACM Press, 2002.
- [23] B. Awerbuch, R. Curtmola, D. Holmer, C. Nita-Rotaru, and H. Rubens, "On the survivability of routing protocols in ad hoc wireless networks," in *Proc. of SecureComm '05*. IEEE, 2005.
- [24] S. Roy, V. G. Addada, S. Setia, and S. Jajodia, "Securing MAODV: Attacks and countermeasures," in *Proc. of SECON '05*. IEEE, 2005.
- [25] R. Curtmola and C. Nita-Rotaru, "BSMR: Byzantine-resilient secure multicast routing in multi-hop wireless networks," in *Proc. of IEEE SECON '07*, June 2007.
- [26] R. Draves, J. Padhye, and B. Zill, "Comparison of routing metrics for static multi-hop wireless networks," in *Proc. of SIGCOMM '04*, 2004.
- [27] S. Roy, D. Koutsonikolas, S. Das, and C. Hu, "High-throughput multicast routing metrics in wireless mesh networks," *Elsevier Ad Hoc Ntwks*, 2007.
- [28] Y.-C. Hu, A. Perrig, and D. B. Johnson, "Rushing attacks and defense in wireless ad hoc network routing protocols," in *Proc. of WiSe*, 2003.
- [29] Y.-C. Hu, A. Perrig, and D. B. Johnson, "Packet leashes: A defense against wormhole attacks in wireless ad hoc networks," in *Proc. of INFOCOM*, 2003.
- [30] J. Eriksson, S. V. Krishnamurthy, and M. Faloutsos, "Truelink: A practical countermeasure to the wormhole attack in wireless networks," in *Proc. of ICNP '06*, 2006.
- [31] L. Hu and D. Evans, "Using directional antennas to prevent wormhole attacks," in *Proc. of NDSS*, 2004.
- [32] A. Perrig, R. Canetti, D. Song, and D. Tygar, "Efficient and secure source authentication for multicast," in *Proc. of NDSS*, February 2001.
- [33] M. Poturalski, P. Papadimitratos, and J.-P. Hubaux, "Secure Neighbor Discovery in Wireless Networks: Formal Investigation of Possibility," Tech. Rep., 2007.
- [34] "Global mobile information systems simulation library - glomosim." [Online]. Available: <http://pcl.cs.ucla.edu/projects/glomosisim/>
- [35] D. S. Moore and G. P. McCabe, *Introduction to the Practice of Statistics*. New York: W.H. Freeman, 2003.