Purdue University

Purdue e-Pubs

2007

BSMR: Byzantine-Resilient Secure Multicast Routing in Multi-hop Wireless Networks

Reza Curtmola

Cristina Nita-Rotaru
Purdue University, crisn@cs.purdue.edu

Curtmola, Reza and Nita-Rotaru, Cristina, "BSMR: Byzantine-Resilient Secure Multicast Routing in Multi-hop Wireless Networks" (2007). *Department of Computer Science Technical Reports.* Paper 1670.
https://docs.lib.purdue.edu/cstech/1670

BSMR: BYZANTINE-RESILIENT SECURE MULTICAST
ROUTING IN MULTI-HOP WIRELESS NETWORKS

Reza Curtmola
Cristina Nita-Rotaru

BSMR: Byzantine-Resilient Secure Multicast Routing in Multi-hop Wireless Networks

Reza Curtmola
Department of Computer Science
The Johns Hopkins University
crix@cs.jhu.edu

Cristina Nita-Rotaru
Department of Computer Science and CERIAS
Purdue University
crisn@cs.purdue.edu

Abstract

Multi-hop wireless networks rely on node cooperation to provide unicast and multicast services. The multi-hop communication offers increased coverage for such services, but also makes them more vulnerable to insider (or Byzantine) attacks coming from compromised nodes that behave arbitrarily to disrupt the network.

In this work we identify vulnerabilities of on-demand multicast routing protocols for multi-hop wireless networks and discuss the challenges encountered in designing mechanisms to defend against them. We propose BSMR, a novel secure multicast routing protocol that withstands insider attacks from colluding adversaries. Our protocol is a software-based solution and does not require additional or specialized hardware. We present simulation results which demonstrate that BSMR effectively mitigates the identified attacks.

I. INTRODUCTION

Multicast routing protocols deliver data from a source to multiple destinations organized in a multicast group. Several protocols were proposed to provide multicast services for multi-hop wireless networks. These protocols rely on node cooperation and use flooding [1], gossip [2], geographical position [3], or dissemination structures such as meshes [4], [5], or trees [6], [7].

A major challenge in designing protocols for wireless networks is ensuring robustness to failures and resilience to attacks. Wireless networks provide a less robust communication than wired networks due to frequent broken links and a higher error rate. Security is also more challenging in multi-hop wireless networks because the open medium is more susceptible to outside attacks and the multi-hop communication makes services more vulnerable to insider attacks coming from compromised nodes. Although an effective mechanism against outside attacks, authentication is not sufficient to protect against insider attacks because an adversary that compromised a node also gained access to the cryptographic keys stored on it. Insider attacks are also known as Byzantine [8] attacks and protocols able to provide service in their presence are referred to as Byzantine resilient protocols.

Previous work focused mainly on the security of unicast services. Several routing protocols [9]–[12] were proposed to cope with outsider attacks. Methods proposed to address insider threats in unicast routing include monitoring [13], multi-path routing [14] and acknowledgment-based feedback [15], [16]. The problem of secure multicast in wireless networks was less studied and only outside attacks were addressed [17].

Security problems related to multicast routing can be classified in routing specific security, such as the management of the routing structure and data forwarding, and application specific security such as data confidentiality and authenticity. Solutions to the latter problem also referred to as secure group communication focus mainly on group key management [18], [19]. In this work we are concerned with multicast routing specific security.

Several differences make the multicast communication model more challenging than its unicast counterpart. First, designing secure multicast protocols for wireless networks requires a more complex trust model, as nodes which are members of the multicast group cannot simply organize themselves in a dissemination structure without the help of other non-member nodes acting as routers.

Second, unlike unicast protocols which establish and maintain routes between two nodes, multicast protocols establish and maintain more complex structures, such as trees or meshes. For example, protocols relying on trees require additional operations such as route activation, tree pruning and tree merging. These actions do not have a counterpart in the unicast case and may expose the routing protocol to new vulnerabilities.

Third, multicast protocols deliver data from one sender to multiple receivers making scalability a major problem when designing attack-resilient protocols. In particular, solutions that offer resiliency against Byzantine attacks for unicast are not scalable in a multicast setting. For example, multi-path routing affects significantly the data dissemination efficiency, while strategies based on end-to-end acknowledgments have high overhead.

In this paper we study vulnerabilities of multicast routing protocols in multi-hop wireless networks and propose a new protocol that provides resilience against Byzantine attacks. Our main contributions are:

- We identify several aspects that make the design of attack-resilient multicast routing protocols more challenging than their unicast counterpart, such as a more complex trust model and underlying routing structure, and scalability. We also discuss potential attacks against such protocols.

- We propose BSMR, an on-demand multicast protocol for multi-hop wireless networks which relies on several mechanisms to mitigate Byzantine attacks. BSMR uses a selective data forwarding detection mechanism that relies on a reliability metric capturing adversarial behavior. Nodes determine the reliability of links by comparing the perceived data rate with the one advertised by the source. Adversarial links are avoided during the route discovery phase. BSMR also prevents attacks that try to prevent or arbitrarily influence route establishment.

- We show through simulations that the impact of several Byzantine attacks (flood rushing, black hole and wormhole) on a previously proposed secure multicast routing protocol is considerable and cannot be ignored. We also demonstrate through simulations that our protocol BSMR mitigates the attacks, while incurring a small overhead.

The remainder of the paper is organized as follows. Section II overviews related work. Section III presents our network and system model. We discuss the attacks against multicast in IV-B and present BSMR in Section V. We present experimental results in Section VI and conclude in Section VII.

II. RELATED WORK

Significant work addresses vulnerabilities of unicast routing protocols in wireless networks. Several secure routing protocols resilient to outside attacks were proposed in the last few years such as Ariadne [11], SEAD [10], ARAN [12], and the work in [9].

Wireless specific attacks such as flood rushing and wormhole were recently identified and studied. RAP [20] 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* [21] which restricts the maximum transmission distance by using either a tight time synchronization or location information, Truelink [22] which uses MAC level acknowledgments to infer if a link exists or not between two nodes, and the technique in [23], which relies on directional antennas.

The problem of insider threats in unicast routing was studied in [13]–[16]. Watchdog [13] relies on a node monitoring its neighbors if they forward packets to other destinations. If a node does not overhear a neighbor forwarding more than a threshold number of packets, it concludes that the neighbor is adversarial. SDT [14] uses multi-path routing to prevent a malicious node from selectively dropping data. ODSBR [15], [16] provides resilience to Byzantine attacks caused by individual or colluding nodes by detecting malicious links based on an acknowledgement-based feedback techniques.

Most of the work addressing application security issues related to multicast in wireless networks focused on the problem of group key management in order to ensure data confidentiality and authenticity [24]–

[28]. Work studying multicast routing specific security problems in wireless networks is scarce with the notable exception of the authentication framework by Roy *et al.* [17]. The framework allows MAODV to withstand several external attacks targeted against the creation and maintenance of the multicast tree. However, it does not provide resilience against Byzantine attacks.

Multicast routing specific security was also studied in overlay networks [29]–[31]. Solutions proposed exploit overlay specific properties such as: existence of network connectivity between each pair of nodes which allows nodes to directly probe non-neighboring nodes, and highly redundant connectivity, guaranteeing that many disjoint paths exist. None of these properties hold in multi-hop wireless networks.

III. NETWORK AND SYSTEM MODEL

A. Network Model

We consider a multi-hop wireless network where nodes participate in the data forwarding process for other nodes. We assume that the wireless channel is symmetric. All nodes have the same transmitting power and consequently the same transmission range. The receiving range of a node is identical to its transmission range.

Nodes are not required to be equipped with additional hardware such as GPS receivers or tightly synchronized clocks. Also, nodes are not required to be tamper resistant: If an attacker compromises a node, it can extract all key material, data or code stored on that node.

B. Multicast Protocol

We assume a tree-based on-demand multicast protocol such as [6]. The protocol maintains bi-directional shared multicast trees connecting multicast sources and receivers. Each multicast group has a corresponding multicast tree. The multicast source is a special node, the group leader, whose role is to eliminate stale routes and coordinate group merges. Route freshness is indicated by a group sequence number updated by the group leader and broadcast periodically in the entire network. Higher group sequence numbers denote fresher routes.

The main operations of the protocol are route discovery, route activation and tree maintenance. During route discovery a node discovers a path to a node that is part of the multicast tree. A requester first broadcasts a route request message that includes the latest known group sequence number. The route request message is flooded in the network using a basic flood suppression mechanism and establishes reverse routes to the source of the request. Upon receiving the route request, a node that is part of the multicast tree and has a group sequence number at least as large as the one in the route request, generates a route reply message and unicasts it on the reverse route. The route reply message includes the last known group sequence number and the number of hops to the node that originated the route reply.

During route activation, the requester selects the freshest and shortest route (i.e., with the smallest number of hops to the multicast tree) from the routes returned by the route discovery operation. The requester activates that route by unicasting a multicast activation message.

Three main operations ensure the tree maintenance: tree pruning, broken link repair and tree merging. Tree pruning occurs when a group member that is a leaf in the multicast tree decides to leave the group. To prune itself from the tree, the node sends a message to indicate this to its parent. The pruning message travels up the tree causing leaf nodes that are not members of the multicast group to prune themselves from the tree, until it reaches either a non-leaf node or a group member. A non-leaf group member must continue to act as a router and cannot prune itself from the multicast tree.

A node initiates a link repair procedure when the upstream link in the multicast tree breaks. If the node cannot reconnect to the tree, it means the tree is partitioned. In this case the node runs a special procedure to prune non-member leaf nodes and elect a group leader for the partition. When two partitions of the same tree reconnect, the leader of one of the partitions coordinates the merge of the partitions, suppressing the other leader.

IV. ATTACKS AGAINST MULTICAST ROUTING

A. Adversarial Model

We assume that nodes may exhibit Byzantine behavior, either alone or colluding with other nodes. Examples of such behavior include: not forwarding packets, injecting, modifying or replaying packets. We refer to any arbitrary action by authenticated nodes resulting in disruption of the routing service as Byzantine behavior, and to such an adversary as a Byzantine adversary.

We consider a three-level trust model that captures the interactions between nodes in a wireless multicast setting and defines a node's privileges: a first level includes the source which must be continually available and assumed not to be compromised; a second level consists of the multicast group member nodes, which are allowed to initiate requests for joining multicast groups; and a third level consists of non-member nodes which participate in the routing but are not entitled to initiate group join requests. In order to cope with Byzantine attacks, even group members cannot be fully trusted.

An attacker can disrupt the physical layer by jamming, and MAC protocols such as 802.11 can be disrupted by attacks using the special RTS/CTS packets. This work only considers attacks targeted against the network level. Also, preventing traffic analysis is not the goal of this work, which focuses instead on survivable routing.

B. Attacks in Multicast in Multi-Hop Wireless Networks

An adversary can attack control messages corresponding to the route discovery, route activation and tree management components of the routing protocol, or can attack data messages.

The route discovery phase can be disrupted by outside attackers creating undesired results by injecting, replaying, or modifying control packets. Nodes that are not in the tree can mislead other nodes into believing that they found and are connected to the tree. Nodes can flood the network with bogus requests for joining multicast groups. A Byzantine adversary can prevent a route from being established by dropping the request and/or response, or can influence the route selection by using wireless specific attacks such as wormhole and flood rushing. A Byzantine adversary can also modify the packets carrying the route selection metric such as hop count or node identifiers.

An outside attacker can inject bogus route activation messages, or prevent correct route activation messages to reach all nodes.

Nodes can maliciously report that other links are broken or generate incorrect pruning messages resulting in correct nodes being disconnected from the network or tree partitioning. In the absence of authentication, any node can pretend to be the group leader. Although many routing protocols do not describe how to select a new group leader when needed, we note that the leader election protocol can also be influenced by attackers.

Attacks against data messages consist of eavesdropping, modifying, replaying, injecting data, or selectively forwarding data after being selected on a route. A special form of packet delivery disruption is a denial of service attack, in which the attacker overwhelms the computational, sending or receiving capabilities of a node. In general, data source authentication, integrity and encryption can solve the first attacks and are usually considered application specific security. Defending against selective data forwarding and denial of service cannot be done exclusively by using cryptographic mechanisms.

V. SECURE MULTICAST ROUTING PROTOCOL

A. BSMR Overview

Our protocol ensures that multicast data is delivered from the source to the members of the multicast group, even in the presence of Byzantine attackers, as long as the group members are reachable through non-adversarial paths and a non-adversarial path exists between a new member and a node in the multicast tree.

To eliminate a large class of outside attacks we use an authentication framework that ensures only authorized nodes can perform certain operations (e.g., only tree nodes can perform tree operations and only nodes that possess valid group certificates can connect to the corresponding multicast tree).

BSMR mitigates inside attacks that try to prevent a node from establishing a route to the multicast tree by flooding both route request and route reply such that if an adversarial-free route exists, then a route is established.

BSMR ensures resilience to selective data forwarding attacks by using a reliability metric that captures adversarial behavior. The metric consists of a list of link weights in which high weights correspond to low reliability. Each node in the network maintains its own *weight list* and includes it in each route request to ensure that a new route to the tree avoids adversarial links.

A link's reliability is determined based on the number of packets successfully delivered on that link over time. Tree nodes monitor the rate of receiving data packets and compare it with the transmission rate indicated by the source in the form of an MRATE message. If the perceived transmission rate falls below the rate indicated in the MRATE message by more than a threshold, an honest node that is a direct descendant of an adversarial node updates its weight list by penalizing the link to its parent and then tries to discover a new route to the tree.

We note that a strategy based on end-to-end acknowledgments, although shown effective in unicast [14], [16], is not scalable: As the size of the multicast group increases, ACK implosion occurs at the source, which may cause a drastic decrease in data delivery [32]. Moreover, solutions that address the problem of feedback implosion in multicast protocols (e.g., feedback aggregation or a combination of ACK/NACK messages [33]) were designed to operate under non-adversarial conditions; It is questionable if they will work in adversarial networks.

Without loss of generality, we limit our description to one multicast group. Below we describe the previously mentioned authentication framework, the *route discovery*, the *route activation*, *multicast tree maintenance* and the *selective data forwarding detection* mechanisms.

B. Authentication Framework

In order to protect from external attacks against the creation and maintenance of the multicast tree BSMR uses a framework similar with the one in [17]. The framework prevents unauthorized nodes to be part of the network, of a multicast group, or of a multicast tree. These forms of authentication correspond to the trust model described in Section IV-A. Each node authorized to join the network has a pair of public/private keys and a *node certificate* that binds its public key to its IP address. Each node authorized to join a multicast group has an additional *group certificate* that binds its public key and IP address to the IP address of the multicast group.

Nodes in the multicast tree are authenticated using a *tree token*, which is periodically refreshed and disseminated by the group leader in the multicast tree with the help of *pairwise shared keys* established between every direct tree neighbors. Thus, only nodes that are currently on the tree will have a valid tree token. To allow any node in the network to check that a tree node possesses a valid tree token, the group leader periodically broadcasts in the entire network a tree token authenticator $f(tree\ token)$, where f is a collision resistant one-way function. Nodes can check the validity of a given tree token by applying the function f to it and comparing the result with the latest received tree token authenticator.

To prevent tree nodes from claiming to be at a smaller hop distance from the group leader than they actually are, we use a technique based on a one-way hash chain. The last element of this hash chain, referred to as hop count anchor, is broadcast periodically by the group leader.

We assume that nodes have a method to determine the source authenticity of the received data (e.g., TESLA [34]). This allows a node to correctly determine the rate at which it receives multicast data.

C. Route Discovery

BSMR's route discovery allows a node that wants to join a multicast group to find a route to the multicast tree. The protocol follows the typical route request/route reply procedure used by on-demand routing protocols with several differences. To prevent outsiders from interfering, all route discovery messages are authenticated using the public key corresponding to the network certificate. Only group authenticated nodes can initiate route requests and the group certificate is required in each request. Tree nodes use the tree token to prove their current tree status.

Several mechanisms are used to address internal attackers: (a) both route request and route reply are flooded in order to ensure that, if an adversarial-free path exists, it will be found; (b) the path selection relies on the weights list carried in the response flood and allows the requester to select a non-adversarial path; (c) the propagation of weights and path accumulation is performed using an onion-like signing to prevent forwarding nodes from modifying the path carried in the response.

The requesting node broadcasts a route request (RREQ) message that includes the node identifier and its weight list, the multicast group identifier, the last known group sequence number, and a request sequence number. The RREQ message is flooded in the network until it reaches a tree node that has a group sequence number at least as great as that in the RREQ. Only new requests are processed by intermediate nodes.

When a tree node receives for the first time a RREQ from a requester and the node's group sequence number is at least as great as that contained in the RREQ, it initiates a response. The node broadcasts a route reply (RREP) message that includes that node identifier, its recorded group sequence number, the requester's identifier, a response sequence number, the group identifier and the weight list from the request message. To prove its current tree node status, the node also includes in the response the current tree token, encrypted with the requester's public key. The RREP message is flooded in the network until it reaches the requester, using the following *weighted flood suppression* mechanism. Tree nodes with a group sequence number at least as great as that in the RREP ignore RREP messages. Otherwise, a node computes the total path weight by summing the weight of all the links on the specified path from the multicast tree to itself. If the total weight is less than any previously forwarded matching response (same requester, multicast group and response sequence number), and all the signatures accumulated on the reply are valid, the node appends its identifier to the end of the message, signs the entire message and rebroadcasts it. As the RREP message propagates across the network, nodes establish the *forward route* by setting pointers to the node from which the RREP was received. Although several tree nodes may initiate the response flood, the weighted flood suppression mechanism insures the communication overhead is equivalent to only one flood.

When the requester receives a response, it performs the same computation as an intermediate node during the response propagation. The requester updates its information upon receipt of a valid response that contains a better path according to our reliability metric.

D. Multicast Route Activation

The requester signs and unicasts on the selected route a multicast activation (MACT) message that includes its identifier, the group identifier and the sequence number used in the route request phase. An intermediate node on the route checks if the signature on MACT is valid and if MACT contains the same sequence number as the one in the original RREQ message. The node then adds to its list of tree neighbors the previous node and the next node on the route as downstream and upstream neighbors, respectively, and sends the MACT message along the forward route.

The requester and the nodes that received the MACT message could be prevented from being grafted to the tree by an adversarial node, selected on the forward route, which drops the MACT message. To mitigate the attack, these nodes will start a *waiting to connect timer* (*WTC_Timer*) upon whose expiration nodes isolate a faulty link and initiate Route Discovery (Event 3 of Sec. V-F). The timers are set to expire

after a value proportional to a node's hop distance to the tree, in the hope that the nodes closer to the tree will succeed in avoiding the adversarial node and will manage to connect to the tree. After a node becomes aware of its expected receiving data rate, it cancels its *WTC_Timer* and behaves as described in Sec. V-F.

E. Multicast Tree Maintenance

The tree maintenance phase ensures the correct operation of the protocol when confronted with events such as pruning, link breakage and node partitioning. Routing messages exchanged by tree neighbors, such as pruning messages (described in Sec. III-B) are authenticated using the pairwise keys shared between tree neighbors. If a malicious node prunes itself even if it has a subtree below it, the honest nodes in this subtree will reconnect to the tree following the procedure described in Sec. V-F. The link repair procedure is initiated by nodes that detect a broken link and is similar with Route Discovery.

The group leader periodically broadcast in the entire network a signed Group Hello message that contains the current group sequence number, the tree token authenticator and the hop count anchor (described in Sec. V-B).

F. Selective Data Forwarding Detection

The source periodically signs and sends in the tree a multicast rate (MRATE) message that contains its data transmission rate ρ_0. As this message propagates in the multicast tree, nodes may add their perceived transmission rate to it. The information in the MRATE message allows nodes to detect if tree ancestors perform selective data forwarding attacks. Depending on whether their perceived rate is within acceptable limits of the rate in the MRATE message, nodes alternate between two states. The initial state of a node is *Disconnected*; After it joins the multicast group and becomes aware of its expected receiving data rate, the node switches to the *Connected* state. Upon detecting a selective data forwarding attack, the node switches back to the *Disconnected* state.

A network operating normally exhibits some amount of natural "loss", which may cause the rate perceived by a node to be smaller than the rate perceived by its tree parent. This natural rate decrease is cumulative as data travels further away from the source. We define a threshold δ as the upper bound for the tolerable loss rate on a single link. If a node's perceived rate is smaller than the last recorded rate in MRATE by more than δ, the node concatenates its identifier and its rate to MRATE and signs the entire message before forwarding it. These added rates serve as proofs that nodes which previously forwarded the MRATE message did not perceive losses much larger than natural losses.

In order to prevent a malicious node from introducing a rate decrease significantly larger than δ, we use another threshold $\Delta > \delta$. Upon receiving an MRATE message, each node first checks if the difference between the last rate in MRATE and the node's perceived rate is greater than Δ. If so, this indicates that there exists at least an adversarial node in between this node and the node that added the last rate to MRATE. The first honest node that notices a difference larger than Δ incriminates the link to its tree parent as faulty (by using an multiplicative weight increase scheme) and assumes responsibility for finding a new route to the tree. The nodes in the subtree below this node will notice there is a "gap" greater than Δ between the rates included in MRATE; They will defer taking any action to isolate the faulty link for an amount of time proportional to the distance from the node that already started the repair procedure, in the hope that the nodes closer to the faulty link will succeed in isolating it. Upon detecting that the expected data packet rate has been restored, nodes cancel the repair procedure.

Figures 1, 2 and 3 describe how a *Connected* node reacts to the following events, respectively: (1) receipt of an MRATE message, (2) timeout of the *MRATE_Timer*, and (3) timeout of the *WTC_Timer*. ρ_{node} denotes the rate at which the node receives packets from its tree parent.

Tree nodes expect to periodically receive MRATE messages, otherwise the *MRATE_Timer* will expire. Note that each tree node stores the latest received MRATE message and uses it to re-initiate the propagation

Fig. 1: receipt of MRATE $= (\rho_0, (id_1, \rho_1), \ldots, (id_k, \rho_k))$

```
 1. if this is the first MRATE message received then
 2.     switch to Connected state
 3.     cancel WTC_Timer
 4. store MRATE message and cancel MRATE_Timer
 5. if state = Connected and WTC_Timer ≠ PENDING then
 6.     if MRATE contains a "gap" larger than Δ then
 7.         start WTC_Timer timer
 8.         forward MRATE
 9.         return
10. else if WTC_Timer = PENDING then
11.     if MRATE contains a "gap" larger than Δ then
12.         MRATE = cat_and_sign(MRATE, (id_node, ρ_node))
13.         forward MRATE
14.         return
15.     else
16.         cancel WTC_Timer
17.         switch to Connected state
18. if ρ_k − ρ_node > Δ then
19.     MRATE = cat_and_sign(MRATE, (id_node, ρ_node))
20.     if WTC_Timer = PENDING then
21.         cancel WTC_timer
22.     switch to Disconnected state
23.     increase weight of the link to the parent
24.     initiate Route Discovery
25. else if ρ_k − ρ_node > δ then
26.     MRATE = cat_and_sign(MRATE, (id_node, ρ_node))
27. forward MRATE message
28. start MRATE_Timer
```

Fig. 2: timeout of MRATE_Timer

```
 1. if state = Connected and WTC_Timer ≠ PENDING then
 2.     retrieve stored MRATE = (ρ_0, (id_1, ρ_1), ..., (id_k, ρ_k))
 3.     if ρ_k − ρ_node > Δ then
 4.         MRATE = cat_and_sign(MRATE, (id_node, ρ_node))
 5.         switch to Disconnected state
 6.         increase weight of the link to the parent
 7.         initiate Route Discovery
 8.     else if ρ_k − ρ_node > δ then
 9.         MRATE = cat_and_sign(MRATE, (id_node, ρ_node))
10.     forward MRATE message
```

Fig. 3: timeout of WTC_Timer

```
 1. switch to Disconnected state
 2. increase weight of the link to the parent
 3. initiate Route Discovery
```

of MRATE if *MRATE_Timer* expires. Also, when *MRATE_Timer* expires a node compares its perceived rate with the expected rate from the stored MRATE message.

VI. EXPERIMENTAL RESULTS

To the best of our knowledge, the only security mechanism proposed on-demand multicast protocols is the authentication framework by Roy et al. [17], to which we refer as A-MAODV. Although A-MAODV withstands several external attacks targeted against the creation and maintenance of the multicast tree, it does not provide additional resilience against Byzantine attacks. In this section, we study the effect of several Byzantine attacks on the performance of A-MAODV and we simulate the same attacks against BSMR in order to show its effectiveness in mitigating the attacks.

Implementation. We implemented BSMR using the ns2 simulator [35], starting from an MAODV implementation [36]. We assumed the protocol uses RSA [37] with 1024-bit keys for public key operations,

AES [38] with 128-bit keys for symmetric encryptions and HMAC [39] with SHA1 as the message authentication code.

The values used for δ and Δ were 10% and 20% of the source's rate, respectively. We developed a protocol-independent Byzantine attack simulation module for ns2.

A. Experimental Methodology

To capture a protocol's effectiveness in delivering data to the multicast group, we used as a performance metric the packet delivery ratio (PDR), defined as:

$$PDR = \frac{P_r}{P_s \cdot N}$$

where P_r is the number of data packets received by multicast group members, P_s is the number of data packets sent by the source and N is the size of the group.

Because external attacks can be prevented using the authentication framework described in Sec. V-B, in this paper we focus on the following three Byzantine attacks:

− **black hole attack:** This is a selective data forwarding attack, in which adversaries only forward routing control packets, while dropping all data packets.

− **wormhole attack:** Two colluding adversaries cooperate by tunneling packets between each other in order to create a shortcut (or wormhole) in the network. The adversaries use the low cost appearance of the wormhole in order to increase the probability of being selected on paths; Once selected on a path, they attempt to disrupt data delivery by executing a black hole attack.

− **flood rushing attack:** The attack exploits the flood duplicate suppression technique used by many wireless routing protocols. By "rushing" an authenticated flood through the network before the flood traveling through a legitimate route, a Byzantine adversary ends up controlling many routes. The attack can be implemented by simply ignoring the small randomized delays which are normally required to reduce the number of collisions. Flood rushing can be used to increase the effectiveness of a black hole or wormhole attack.

In order to quantify the impact of adversarial positioning, we consider the following scenarios:

− **random placement:** adversaries are placed randomly in the simulation area;

− **strategic placement:** adversarial placement is as follows:

- **black hole attack:** adversaries are placed strategically around the multicast source, equidistant on a circle with radius of 200 meters.
- **wormhole attack:** Given k adversaries, one adversary is placed near the source, at coordinates (650,650). The other $k - 1$ adversaries are placed throughout the simulation area so that the areas covered by their transmission range overlap as little as possible. Each one of the $k - 1$ adversaries is connected via a wormhole tunnel to the adversary placed near the source, thus creating $k - 1$ wormholes.

In Fig. 4, we illustrate the strategic adversarial placement for the black hole attack (in the presence of 6 adversaries) and for the wormhole attack (in the presence of 7 adversaries).

To study the influence of whether the adversaries explicitly join the multicast group and the order of joining, we consider two scenarios:

− **NJOIN:** adversarial nodes do not join the group;

− **JOIN:** adversarial nodes explicitly join the group before any of the honest members join. The adversaries are considered group members in the formula for PDR.

We chose these test case scenarios in order to study the impact of the attacks under a light set of conditions (adversaries are placed randomly, or they do not explicitly join the multicast group) and under a more extreme set of conditions (adversaries are placed strategically, or they join the group before honest nodes do).
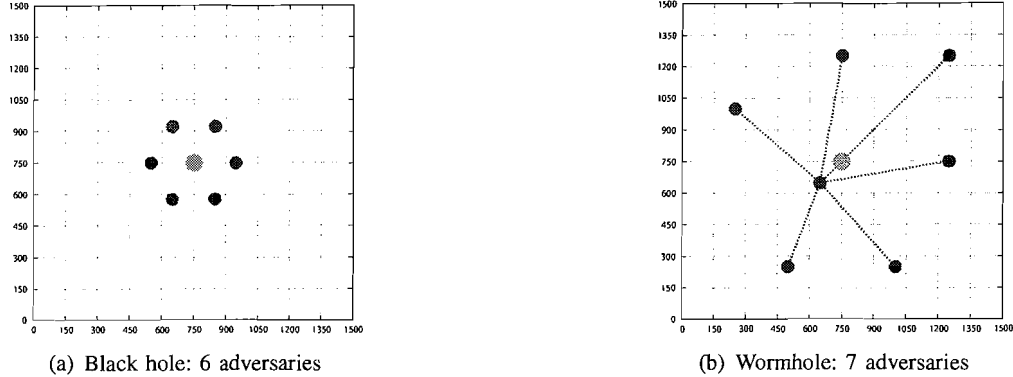
(a) Black hole: 6 adversaries   (b) Wormhole: 7 adversaries

Fig. 4: Examples of strategic adversarial placement for the black hole and wormhole attacks

B. Simulation Setup

We performed simulations using the ns2 network simulator [35]. Nodes were set to use 802.11 radios with 2 Mbps bandwidth and 250 meters nominal range. The simulated time was 600 seconds. We randomly placed 100 nodes within a 1500 by 1500 meter area and the multicast source in the center of the area at coordinates (750,750). We experimented with different values for group size (10, 30 and 50), for number of adversaries (between 16% and 60% of the group size) and for "max" speed (0, 2 and 5 m/s).

Group members join the group sequentially in the beginning of the simulation, each one at an interval of 3 seconds. Then the source transmits multicast data for 600 seconds at a rate of 5 packets per second, each packet of 256 bytes (resulting in loads between 100-500 Kbps across all receivers). The members stay in the group until the end of the simulation. Adversaries added to the network replace honest nodes, thus modeling the capture of honest nodes.

We used a random way-point mobility model, but we incorporated changes to address concerns raised in [40] about the validity of the standard random way-point model. In particular, nodes select a speed uniformly between 10% and 90% of the given maximum speed to achieve a more steady mobility pattern and ensure that the average speed does not drop drastically over the course of the simulation. In addition, 300 seconds of mobility are generated before the start of the simulation so that nodes are already in motion. This allows the average speed and node distribution to stabilize before the simulation starts.

Each data point in the figures is averaged over 30 different random environments and over all group members.

We evaluate the PDR as a function of the number of adversaries, for different group sizes and levels of mobility. Each graph of sections VI-C and VI-D illustrates the effect of the attacks with and without flood rushing.

C. Attack Resilience: The black hole attack

    a) Impact of Adversarial Placement: Figures 5 and 6 show the results for random and strategic adversarial placement, respectively. For random placement we see that, for the same group size, the PDR of A-MAODV decreases as the number of adversaries increases. For the same number of adversaries, it also decreases as we increase the group size. However, random adversarial placement causes the number of group members in the subtree below an adversary to be low; Thus a relatively large number of adversaries is needed to cause a significant disruption (e.g., 30 adversaries for a group of size 50 can cause the PDR to drop below 50%). In the presence of flood rushing, the PDR decreases further because adversaries actively try to get selected themselves as part of the tree. The impact of flood rushing decreases as the group size and the nodal speed increase.

We notice that, for A-MAODV, increasing the nodal speed does not have a negative effect on the PDR; On the contrary, at higher speeds we even see a slight increase in PDR. The effect of link breaks due to
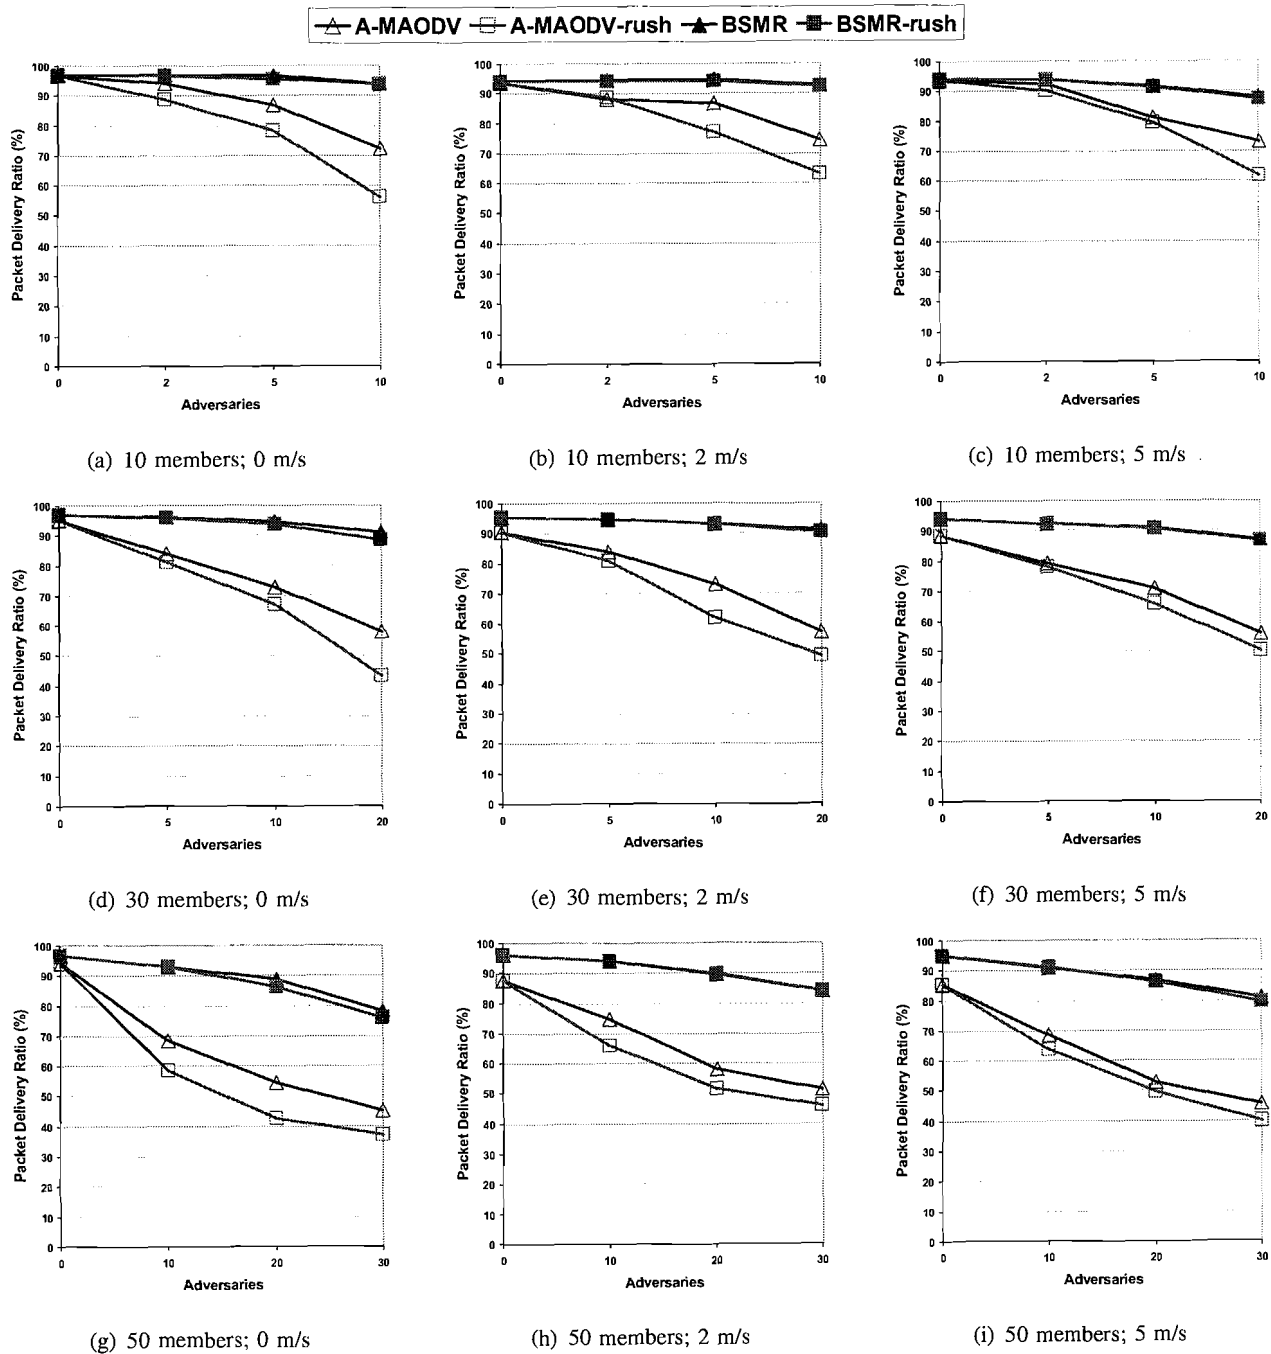
Fig. 5: Black hole attack and flood rushing combined with black hole: **Random placement (NJOIN)**

mobility is compensated by the fact that group members get a chance to reconnect to the multicast tree in a different position, possibly connected to the source through an adversarial-free path. For the same reason, the effect of flood rushing is diminished as the nodal speed increases.

BSMR is almost unaffected by the black hole attack (see Fig. 5). The PDR drops by less than 10% even in the presence of 20 adversaries. In addition, the influence of flood rushing is unnoticeable. This shows the effectiveness against flood rushing of BSMR's strategy, which includes the processing of all response flood duplicates and the metric capturing past behavior of adversarial nodes. Mobility causes a slight PDR decrease, which is natural because higher speeds will cause more link breaks.
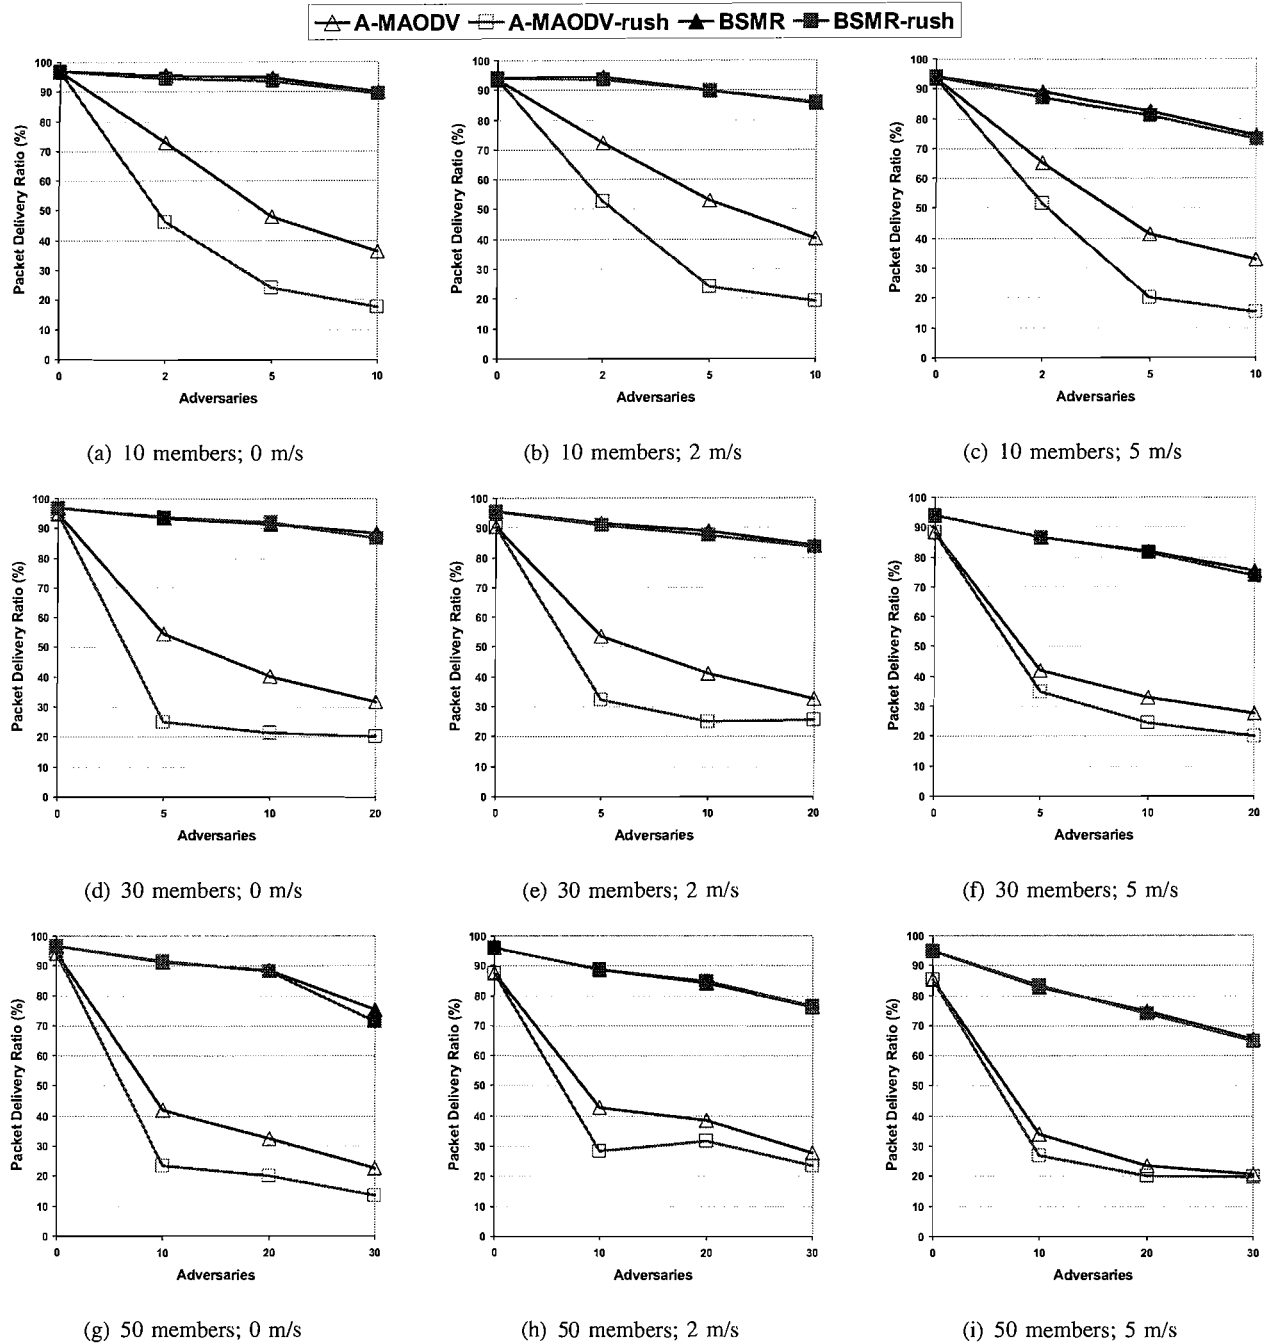
Fig. 6: Black hole attack and flood rushing combined with black hole: **Strategic placement (NJOIN)**

In the previous experiments, the adversaries were randomly placed. Fig. 6 shows the results when adversaries are strategically positioned as described in Section VI-A. For A-MAODV we notice a drastic drop in the PDR. For example, at 0 m/s, when the group size is 30, only 5 adversaries (representing 16% of the group size) are able to reduce the PDR to 25% by executing the black hole attack with flood rushing. This is a direct consequence of the fact that an adversary is now selected in the tree closer to the root and the subtree below it may potentially contain many group members. For the same reason, the negative effect of the flood rushing attack is now emphasized when compared to the random placement case. We conclude that strategic adversarial positioning has a crippling effect on the performance of A-MAODV.
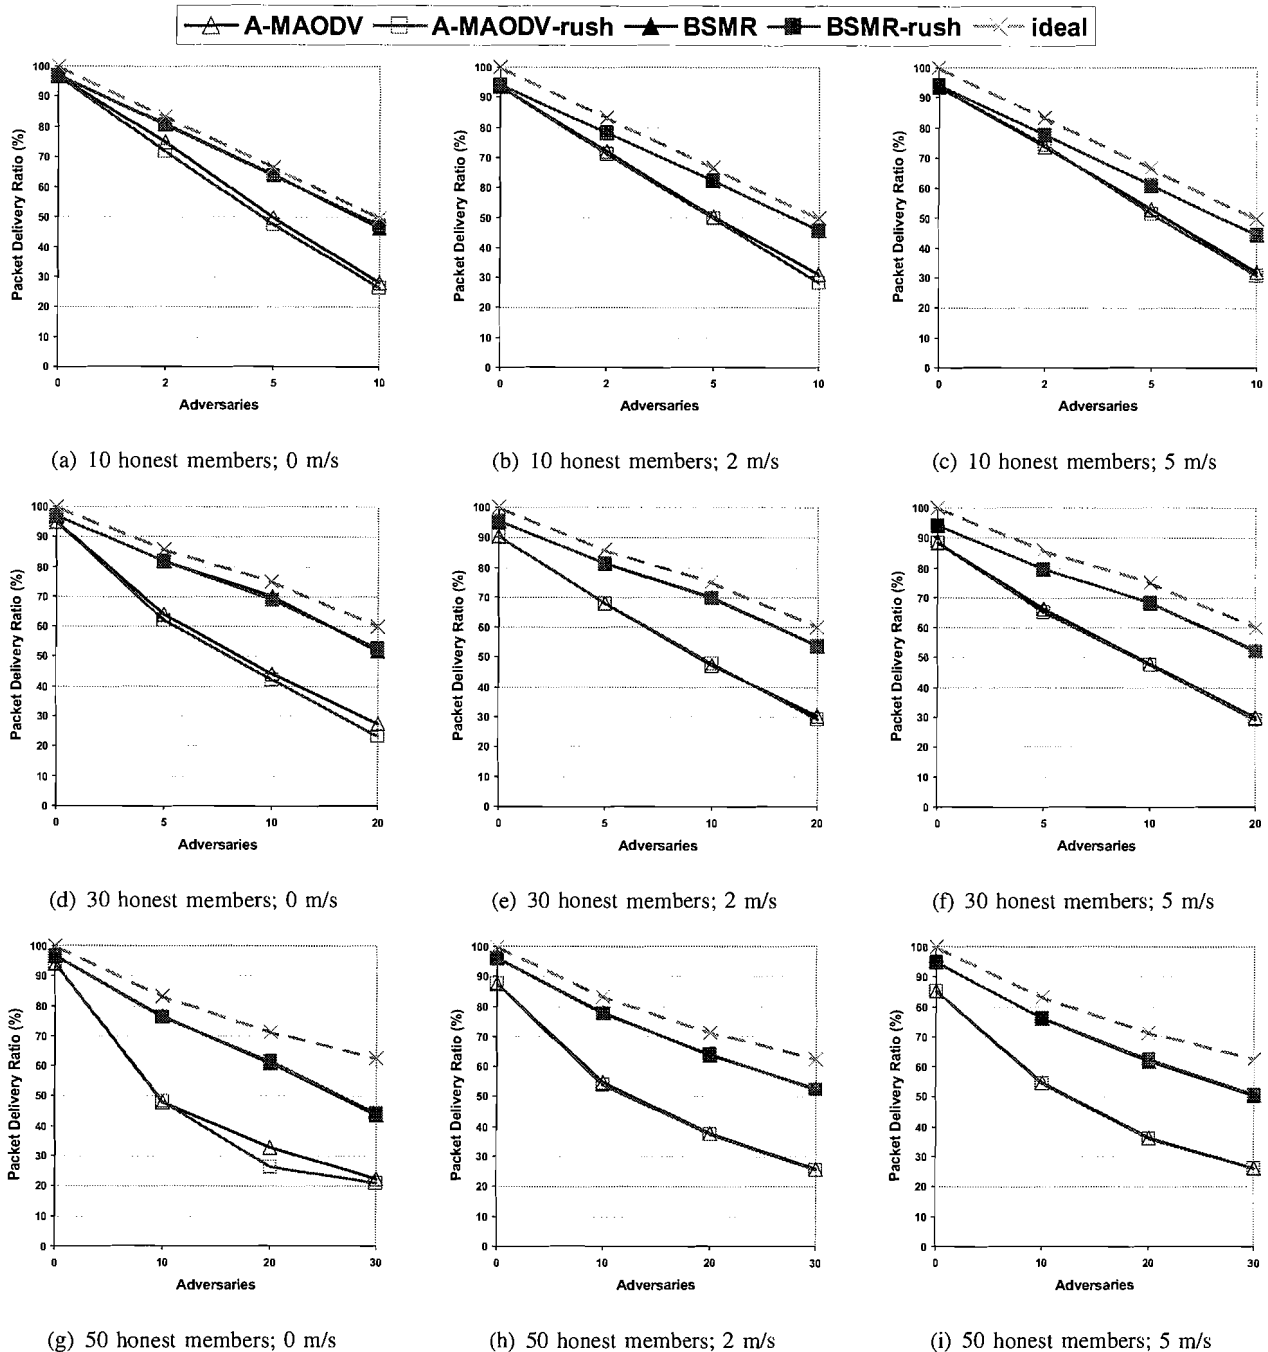
Fig. 7: Black hole attack and flood rushing combined with black hole: **Random placement** (**JOIN**)

On the contrary, the effect of strategic adversarial positioning on BSMR is minor (Fig. 6). Like for random placement, the PDR drops by less than 10% even in the presence of 20 adversaries, at low nodal speeds. When more adversaries are present, we see a slightly larger PDR decrease because there are less available honest nodes left in the network to serve as intermediaries for the group members. The resilience of BSMR to attacks that otherwise have a devastating effect on A-MAODV validates the effectiveness of BSMR's approach.

b) Impact of Explicit Join and Join Order: To analyze the impact of explicit join of adversaries to the multicast group (**JOIN**), when compared to the **NJOIN** case, we look again at the cases where adversaries
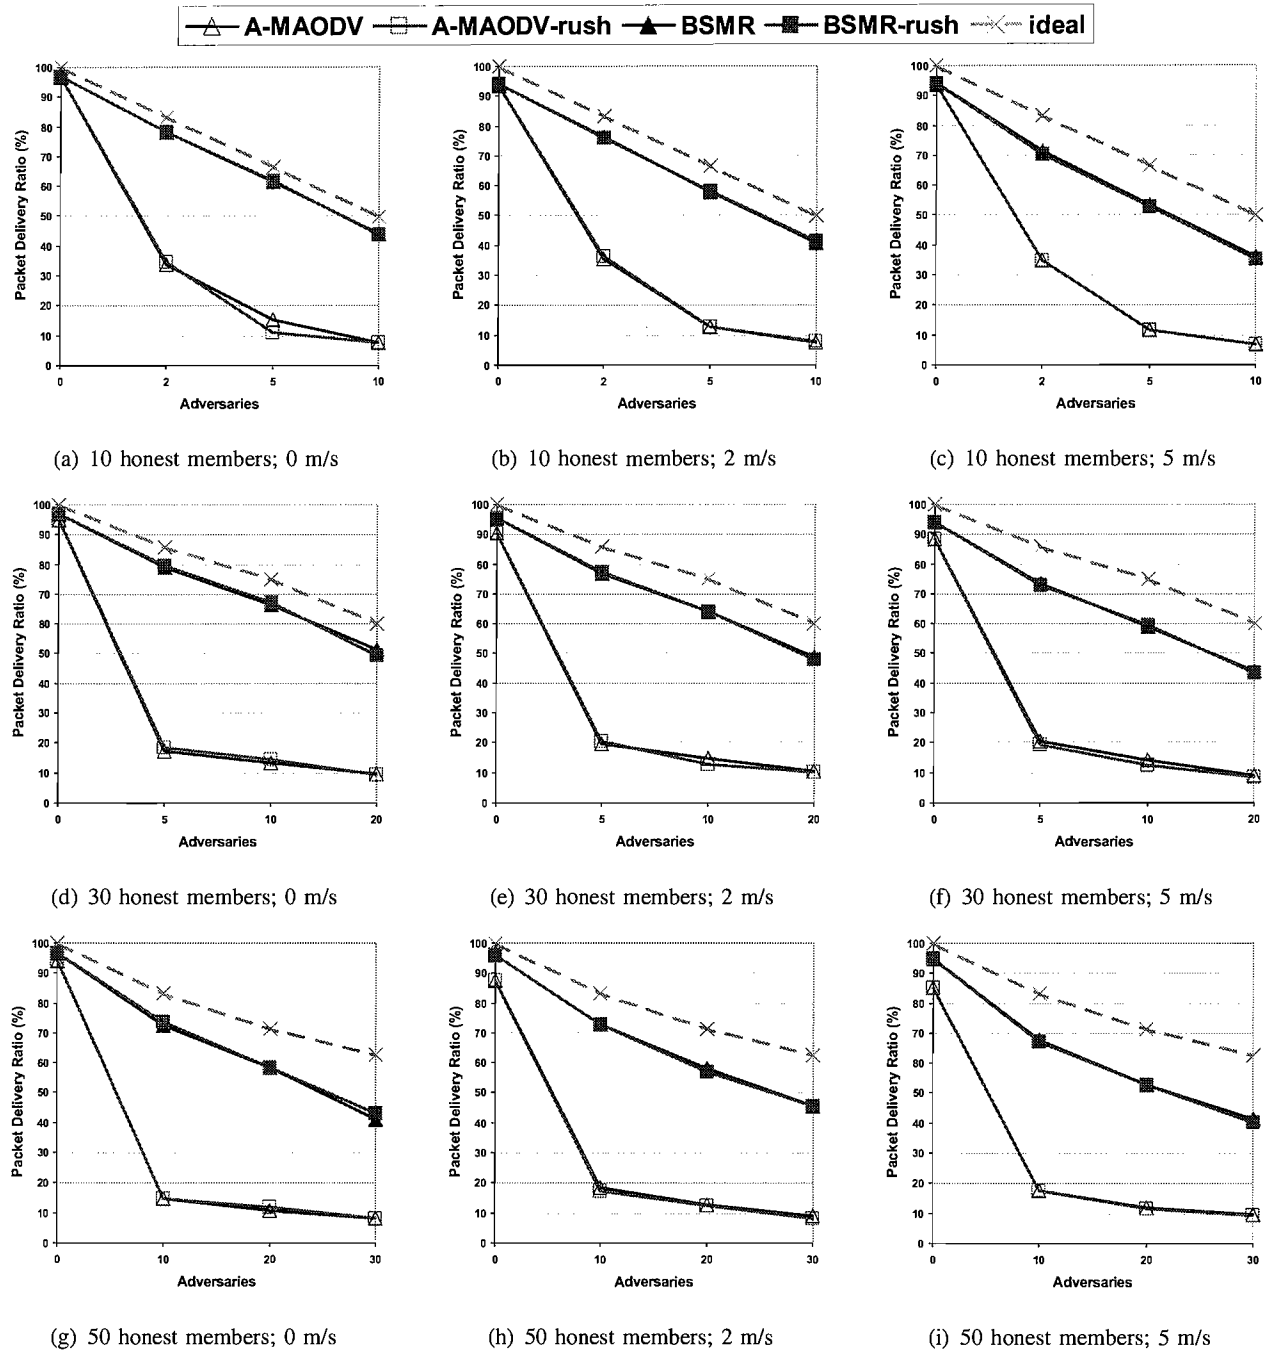
Fig. 8: Black hole attack and flood rushing combined with black hole: **Strategic placement (JOIN)**

are randomly and strategically placed (Figures 7 and 8, respectively). Each figure also shows the ideal PDR (labeled **ideal**), which would be obtained if every one of the honest group members would receive all the packets sent by the source. The effectiveness of an attack should be interpreted as the *difference* between the **ideal** line and a protocol's PDR line. An attack is effective if this difference increases as the number of adversarial nodes is increased; On the other hand, a protocol resilient to attack appears as a line that remains parallel to the ideal line.

For random adversarial placement in Fig. 7, just like in the NJOIN case, we also see a decrease in

PDR as the number of adversaries increases. In addition, we see a major difference from the NJOIN case: When the adversaries explicitly join the group before the honest nodes join, the impact of flood rushing is minimal because the adversaries are already part of the group and rushing control packets does not provide any additional benefit. On the contrary, in this case flood rushing may actually improve the PDR because, by rushing control packets, adversaries may help legitimate nodes to find routes faster. Also, in MAODV a joining node activates a route only after waiting an interval of time, during which it may receive several route reply messages. At the end of this interval, it selects the best route according to a hop count metric. Thus, even if a adversarial tree node managed to rush its route reply, it may get overwritten by a legitimate route reply with a better hop metric.

A drastic drop in the PDR is observed for A-MAODV when adversaries are placed strategically (see Fig. 8). We conclude that strategic positioning has a crippling effect on the performance of A-MAODV even more when adversaries explicitly join the multicast group. For both random and strategic adversarial placement, BSMR is barely affected by the attacks: In most cases the PDR line remains almost parallel to the ideal line, which shows little degradation occurs as the number of adversaries increases. The impact of the attacks on BSMR increases slightly when a large number of adversaries have joined the group, because there are less available honest nodes left in the network to serve as intermediaries for honest group members. We conclude that BSMR's strategy is effective in the **JOIN** case as well.

D. Attack Resilience: The wormhole attack

   c) Impact of Adversarial Placement: Figures 9 and 10 show the results for random and strategic adversarial placement, respectively. Unlike in the experiments for the black hole attack, where the maximum number of adversaries was varied proportionally with the group size, for the wormhole attack only up to 12 adversaries (forming 6 wormholes) and up to 7 adversaries (forming 6 wormholes) were used for the random and strategic placement cases, respectively, regardless of the group size.

When adversaries are placed randomly, the delivery ratio of A-MAODV is affected by an increase in the number of adversaries, as expected. In general, we notice that the wormhole attack causes more damage than the black hole attack because less adversarial nodes are needed to decrease the delivery ratio by the same amount. As an example, for speed of 0 m/s and group size of 30 members, only 10 nodes forming wormholes can reduce the PDR to 60%; On the other hand, 20 adversaries executing the black hole attack are needed to reduce the PDR by the same amount. Flood rushing has a noticeable impact especially for small group sizes and for low mobility levels. Increasing the nodal speed reduces slightly the impact of the wormhole and flood rushing attacks. The explanation is similar with that for the black hole attack: Mobility causes link breaks and honest nodes reattach themselves in different positions in the multicast tree, possibly connected to the source through non-adversarial paths.

The wormhole attack has a minor effect on the delivery ratio of BSMR and the addition of flood rushing has no effect. The PDR drops by less than 10% for all simulated scenarios, even at higher levels of mobility and remains above 90% for most simulated scenarios. This confirms the effectiveness of BSMR's strategy against the wormhole attack.

Fig. 10 shows the results when adversaries are strategically placed and form wormholes as described in Section VI-A. Because one end of each wormhole is near the source, the other end of the wormhole will present attractive, short routes to the nodes that are in its transmission range. This has a serious impact on the delivery ratio of A-MAODV: When 7 adversaries are present, forming 6 wormholes, they cover almost completely the entire simulation area; Thus, every honest node that is normally more than three hops away from the source will be tricked into connecting to the multicast tree through one of the adversarial nodes. When flood rushing is employed, even nodes that are normally less than three hops away from the source are also tricked into connecting through adversarial nodes. This has a serious impact on the PDR of A-MAODV, which drops as low as 23% (a drop of over 70%) when flood rushing is present
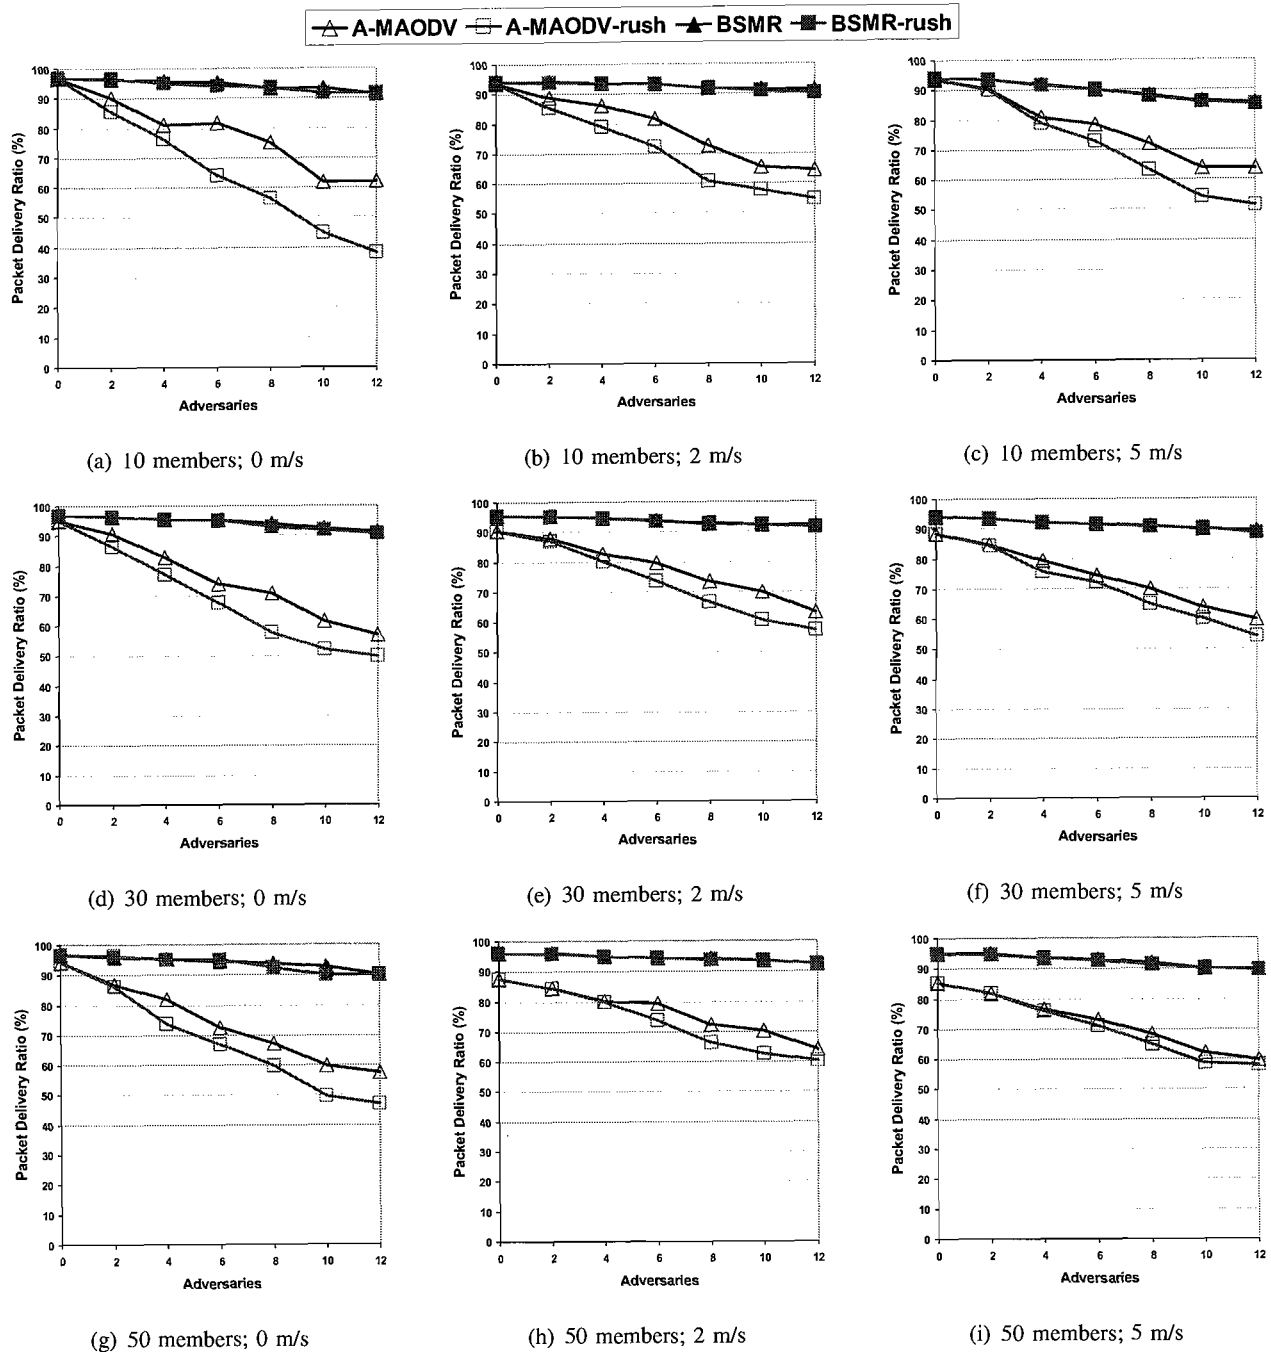
Fig. 9: Wormhole attack and flood rushing combined with wormhole: **Random placement (NJOIN)**

for 0 m/s and group size is 30. We conclude that a small number of strategically placed adversaries can cause considerable damage when A-MAODV is used.

For BSMR, strategic placement results in a slim increase in the delivery ratio drop, when compared to the random placement case. For most simulated scenarios, the PDR drops by less than 10%. The solid performance of BSMR serves as proof for the robustness of its approach.

d) Impact of Explicit Join and Join Order: To analyze the impact of explicit join of adversaries to the multicast group (**JOIN**), when compared to the **NJOIN** case, we look again at the cases where adversaries are randomly and strategically placed (Figures 11 and 12, respectively). Each figure also shows the ideal
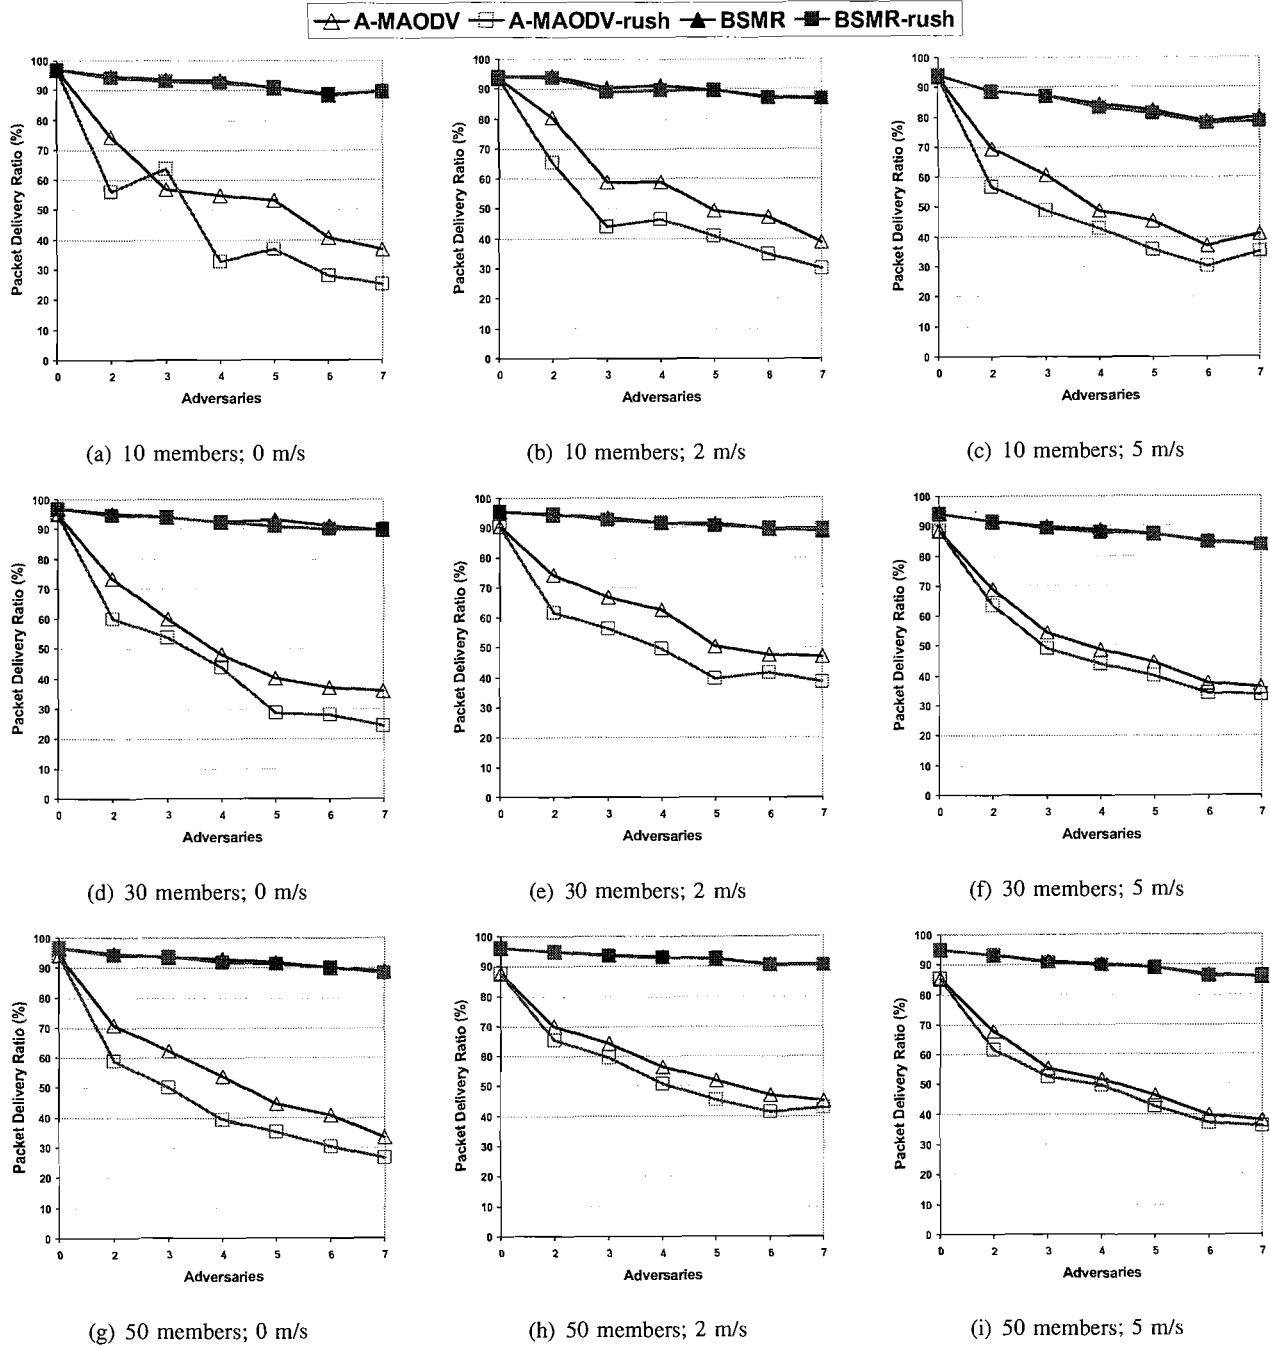
Fig. 10: Wormhole attack and flood rushing combined with wormhole: **Strategic placement (NJOIN)**

PDR (labeled **ideal**), which would be obtained if every one of the honest group members would receive all the packets sent by the source. The effectiveness of an attack should be interpreted as the *difference* between the **ideal** line and a protocol's PDR line. An attack is effective if this difference increases as the number of adversarial nodes is increased; On the other hand, a protocol resilient to attack appears as a line that remains parallel to the ideal line.

In general, for both random and strategic adversarial placement, the effect of flood rushing on both A-MAODV and BSMR is unnoticeable. However, the reasons behind the immunity to flood rushing are different: BSMR prevents flood rushing attacks by using a reliability metric and by processing all flood
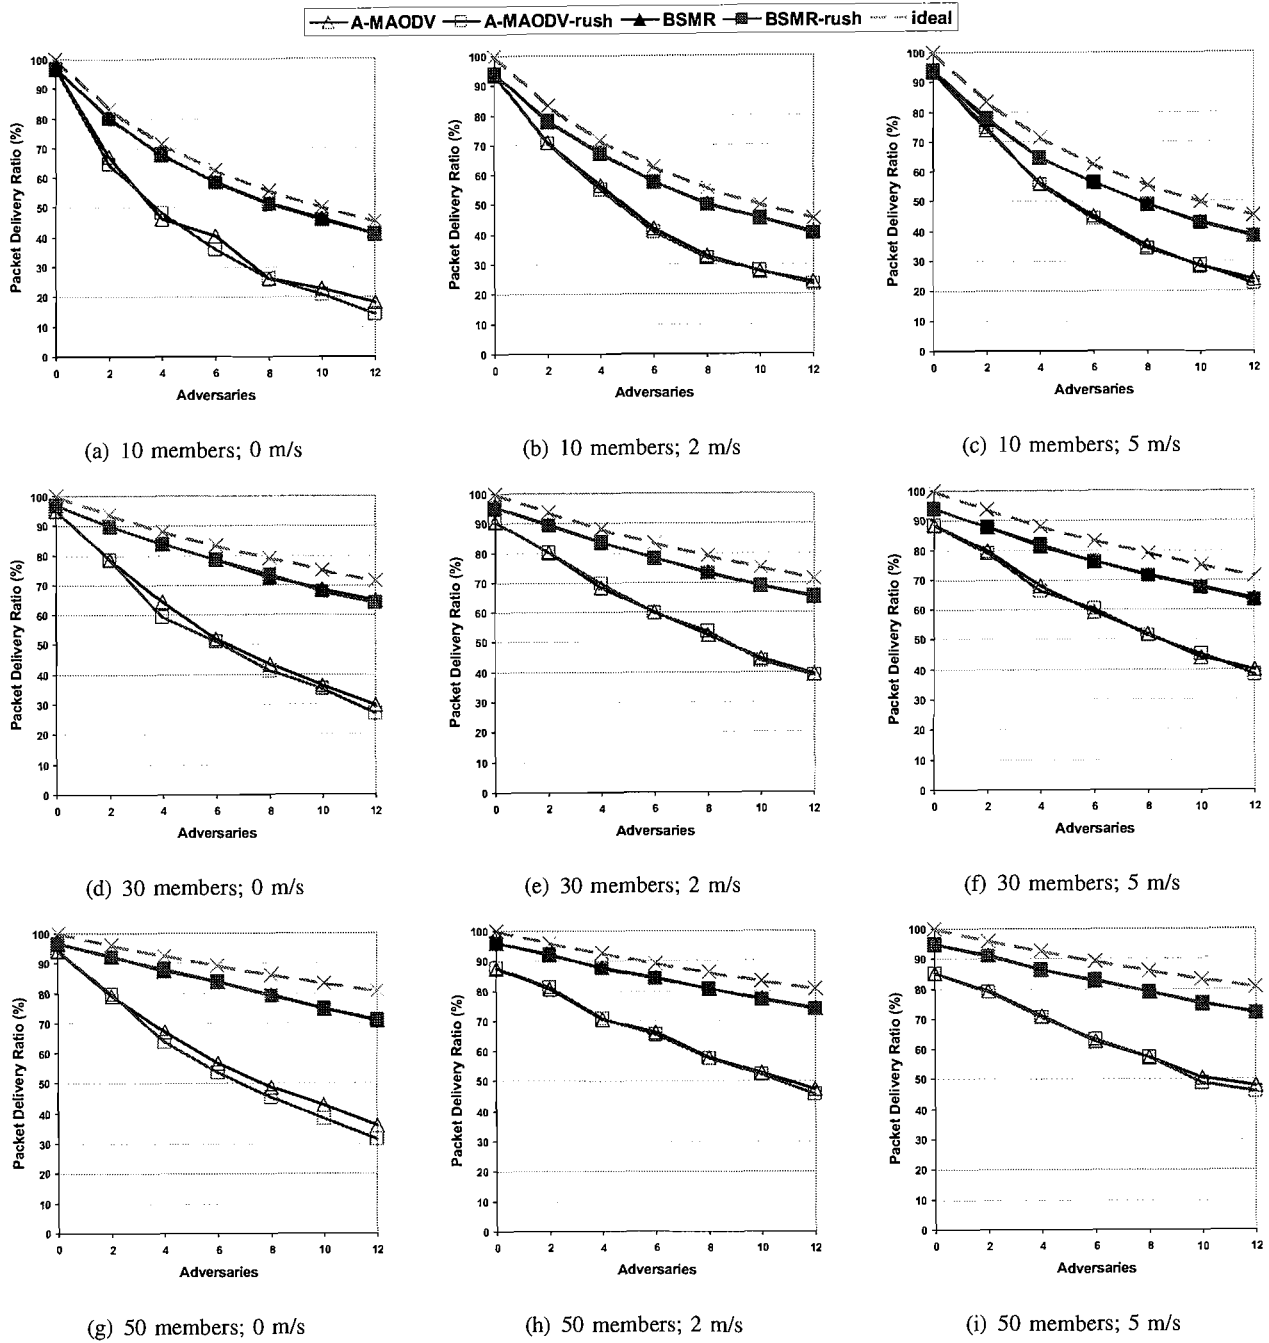
Fig. 11: Wormhole attack and flood rushing combined with wormhole: **Random placement (JOIN)**

duplicates, forwarding the ones with a better metric. As already explained in Sec. VI-C, A-MAODV's immunity to flood rushing is simply an artefact that the adversarial nodes join the multicast tree before the honest nodes do.

When adversaries are placed randomly (Fig. 11), the impact of the wormhole attack on A-MAODV is significant and increases as the number of wormholes increases. BSMR's delivery ratio line remains almost parallel with the ideal line, which means that little degradation occurs as more adversaries join the multicast group.

The devastating impact of the wormhole attack on A-MAODV becomes obvious when adversaries are
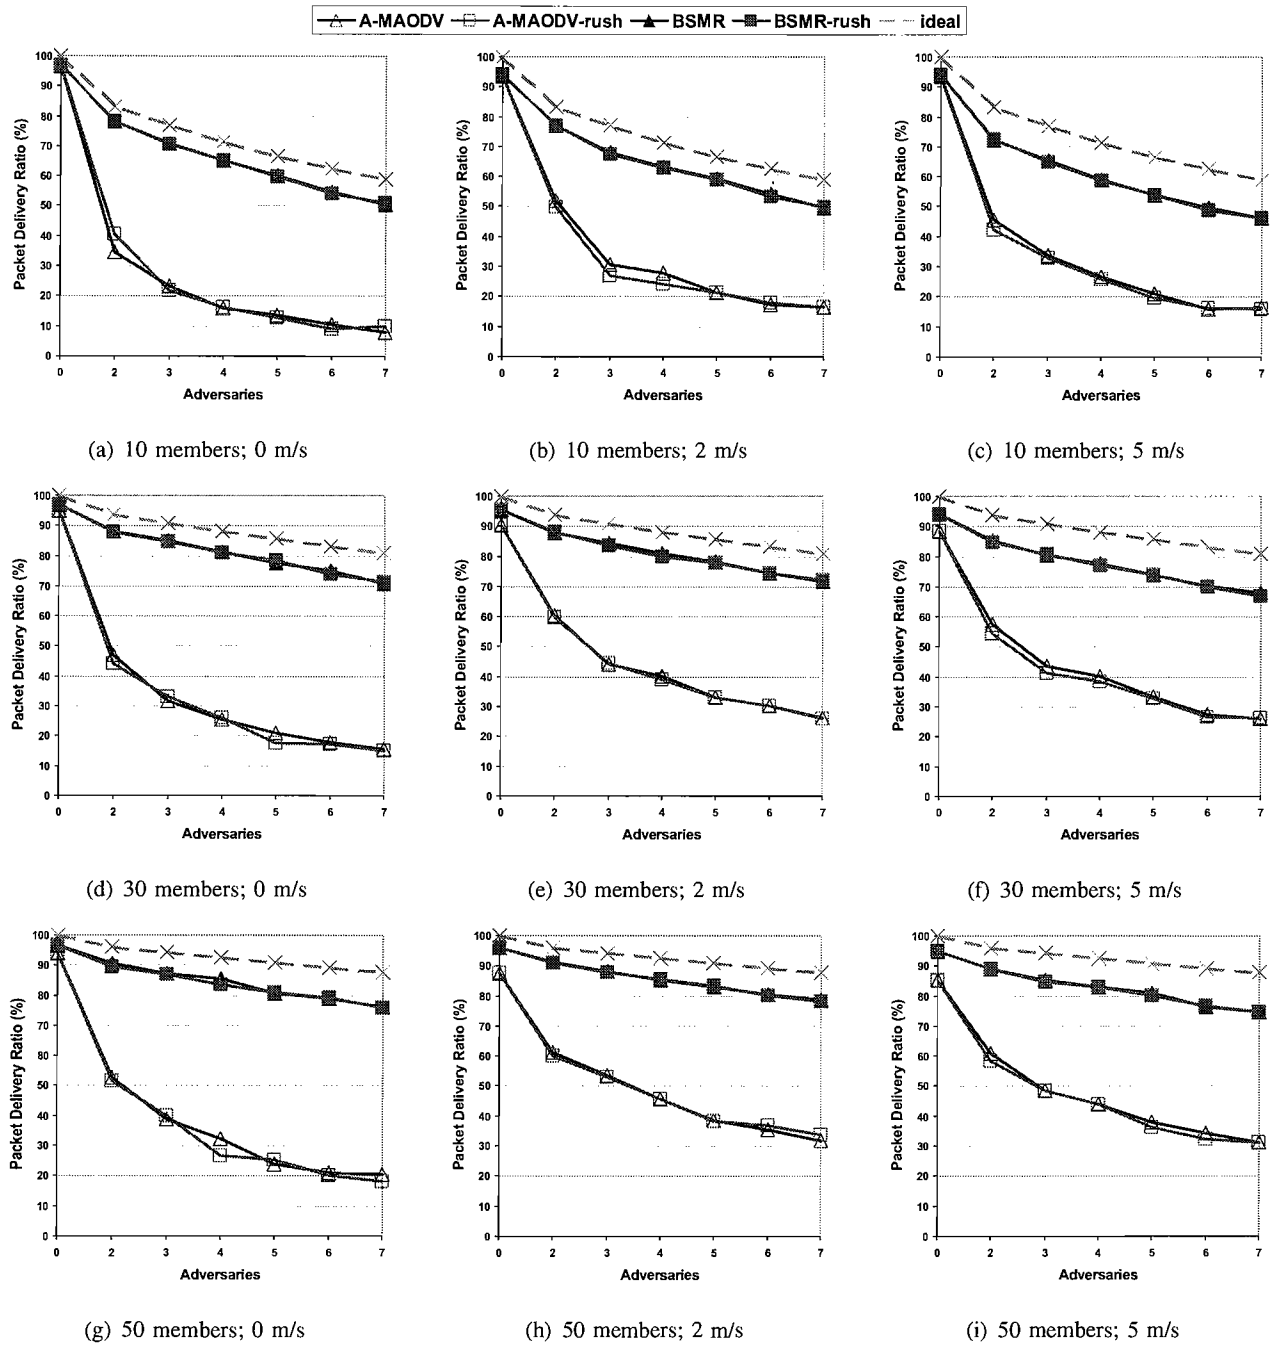
Fig. 12: Wormhole attack and flood rushing combined with wormhole: **Strategic placement (JOIN)**

strategically placed (Fig. 12). When 7 adversaries are present, the difference between the ideal line and A-MAODV's line increases by more than 60% at 0 m/s, for both group sizes of 30 and 50 (causing the PDR to drop to 15-20%). BSMR is affected slightly more than when adversaries are randomly placed, but in all simulated scenarios the difference between the ideal line and BSMR's line does not grow more than 8%, even in the presence of 7 adversaries. This validates the effectiveness of BSMR's strategy in the **JOIN** case a well.
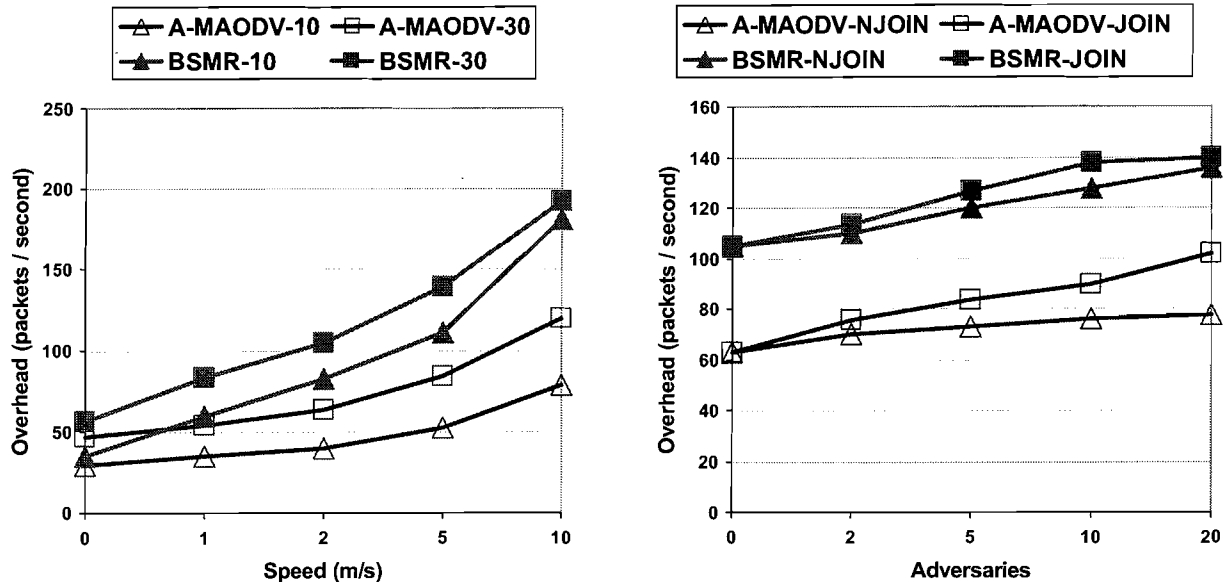
(a) Non-adversarial scenario (group size ∈ {10, 30})   (b) Attack Scenario (group size = 30, 2m/s, strategic placement)

Fig. 13: BSMR overhead

E. Protocol Overhead

We first compare the overhead incurred by A-MAODV and BSMR in a non-adversarial scenario (Fig. 13(a)), for group sizes of 10 and 30, under different levels of mobility. BSMR has higher overhead because both route request and route reply are broadcast, and because of the extra MRATE packets broadcast periodically. BSMR's overhead due to double flooding in the route discovery phase becomes more noticeable especially at higher levels of mobility, when link breaks occur more often.

We then compare the overhead under adversarial conditions. In Fig. 13(b) we focus on one of the strongest studied attack scenarios: Black hole with strategic adversarial placement. For the NJOIN case, BSMR's additional overhead compared to A-MAODV grows slowly as the number of adversaries increases (from 40 more packets/sec. for 0 adversaries to 55 more packets/sec. for 20 adversaries). For the JOIN case, the additional overhead does not grow as we increase the number of adversaries. This indicates that as the number of adversaries increases, BSMR incurs little extra overhead over the non-adversarial case. This is not surprising: the bulk of the additional overhead is caused by the initial route discovery phase which leads to the creation of the multicast tree with avoidance of adversaries; Afterwards, BSMR's additional overhead consists only of periodical MRATE packets and of occasional route discovery in case a link breaks due to mobility.

VII. CONCLUSION

In this paper we have discussed several aspects that make designing attack-resilient multicast routing protocols for multi-hop wireless networks more challenging when compared to their unicast counterpart. A more complex trust model and underlying structure for the routing protocol make solutions tailored for unicast settings not applicable for multicast protocols. In the absence of defense mechanisms, Byzantine attacks can prevent multicast protocols to achieve their design goals.

We have proposed BSMR, a routing protocol which relies on novel general mechanisms to mitigate Byzantine attacks. BSMR identifies and avoids adversarial links based on a reliability metric associated with each link and capturing adversarial behavior. Our experimental results show that BSMR's strategy is

effective against strong insider attacks such as black holes and flood rushing. We believe that this strategy can also be effective against wormhole attacks and defer the experimental validation for future work.

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, pp. 471–480, 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 Computer Society, 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, pp. 151–165, 2001.

[5] S. J. Lee, W. Su, and M. Gerla, "On-demand multicast routing protocol in multihop wireless mobile networks," *Mob. Netw. Appl.*, vol. 7, no. 6, pp. 441–453, 2002.

[6] E. Royer and C. 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] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," in *Advances in Ultra-Dependable Distributed Systems*. IEEE Computer Society Press, 1995. [Online]. Available: citeseer.nj.nec.com/lamport82byzantine.html

[9] P. Papadimitratos and Z. Haas, "Secure routing for mobile ad hoc networks," in *Proc. of CNDS*, January 2002, pp. 27–31.

[10] 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.

[11] Y.-C. Hu, A. Perrig, and D. B. Johnson, "Ariadne: A secure on-demand routing protocol for ad hoc networks," in *Proc. of MOBICOM*, September 2002.

[12] 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*, November 2002.

[13] S. Marti, T. Giuli, K. Lai, and M. Baker, "Mitigating routing misbehavior in mobile ad hoc networks," in *Proc. of MOBICOM*, August 2000.

[14] P. Papadimitratos and Z. Haas, "Secure data transmission in mobile ad hoc networks," in *Proc. of WiSe*, 2003, pp. 41–50.

[15] 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.

[16] 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.

[17] S. Roy, V. G. Addada, S. Setia, and S. Jajodia, "Securing MAODV: Attacks and countermeasures," in *Proc. 2nd IEEE Int'l. Conf. SECON*. IEEE, 2005.

[18] C. Zhang, B. DeCleene, J. Kurose, and D. Towsley, "Comparison of inter-area rekeying algorithms for secure wireless group communications," *Perform. Eval.*, vol. 49, no. 1-4, 2002.

[19] K. H. Rhee, Y. H. Park, and G. Tsudik, "An architecture for key management in hierarchical mobile ad hoc networks," *Journal of Communication and Networks*, vol. 6, no. 2, June 2004.

[20] Y.-C. Hu, A. Perrig, and D. B. Johnson, "Rushing attacks and defense in wireless ad hoc network routing protocols," in *Proc. of WiSe*. ACM, 2003.

[21] ——, "Packet leashes: A defense against wormhole attacks in wireless ad hoc networks," in *Proc. of INFOCOM*, April 2003.

[22] J. Eriksson, S. Krishnamurthy, and M. Faloutsos, "Truelink: A practical countermeasure to the wormhole attack in wireless networks," in *Proc. of ICNP'06*. IEEE, 2006.

[23] L. Hu and D. Evans, "Using directional antennas to prevent wormhole attacks," in *Proc. of NDSS*, 2004.

[24] D. Bruschi and E. Rosti, "Secure multicast in wireless networks of mobile hosts: protocols and issues," *Mobile Networks and Applications*, vol. 7, no. 6, pp. 503–511, 2002.

[25] T. Kaya, G. Lin, G. Noubir, and A. Yilmaz, "Secure multicast groups on ad hoc networks," in *Proc. of SASN'03*. ACM Press, 2003, pp. 94–102.

[26] S. Zhu, S. Setia, S. Xu, and S. Jajodia, "Gkmpan: An efficient group rekeying scheme for secure multicast in ad-hoc networks," in *Proc. of Mobiquitos'04*. IEEE, 2004, pp. 42–51.

[27] L. Lazos and R. Poovendran, "Power proximity based key management for secure multicast in ad hoc networks," 2005, aCM Journal on Wireless Networks (WINET).

[28] R. Balachandran, B. Ramamurthy, X. Zou, and N. Vinodchandran, "CRTDH: an efficient key agreement scheme for secure group communications in wireless ad hoc networks," in *Proc. of ICC 2005*, vol. 2, 2005, pp. 1123–1127.

[29] S. Banerjee, S. Lee, B. Bhattacharjee, and A. Srinivasan, "Resilient multicast using overlays," *SIGMETRICS Perform. Eval. Rev.*, vol. 31, no. 1, pp. 102–113, 2003.

[30] V. Pappas, B. Zhang, A. Terzis, and L. Zhang, "Fault-tolerant data delivery for multicast overlay networks," in *Proc. of ICDCS '04*, 2004, pp. 670–679.

[31] L. Xie and S. Zhu, "Message dropping attacks in overlay networks: Attack detection and attacker identification," in *Proc. SecureComm '06*. IEEE and Create-NET, 2006.

[32] R. Curtmola and C. Nita-Rotaru, "Secure multicast routing in wireless networks," to be published in ACM MC^{2}R, 2006, www.cs.jhu.edu/crix/publications/mc2r.pdf.

[33] K. Tang, K. Obraczka, S.-J. Lee, and M. Gerla, "A reliable, congestion-controlled multicast transport protocol in multimedia multi-hop networks," in *Proc. of IEEE WPMC'02*, 2002.

[34] A. Perrig, R. Canetti, D. Song, and J. D. Tygar, "Efficient and secure source authentication for multicast," in *Proc. of NDSS'01*, 2001.

[35] "The network simulator - ns2," http://www.isi.edu/nsnam/ns/. [Online]. Available: http://www.isi.edu/nsnam/ns/

[36] Y. Zhu and T. Kunz, "MAODV implementation for NS-2.26," Carleton University, Technical Report SCE-04-01.

[37] R. L. Rivest, A. Shamir, and L. M. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978.

[38] *Advanced Encryption Standard (AES)*. National Institute for Standards and Technology (NIST), 2001, no. FIPS 197.

[39] *The Keyed-Hash Message Authentication Code (HMAC)*. NIST, 2002, no. FIPS 198.

[40] J. Yoon, M. Liu, and B. Noble, "Random waypoint considered harmful," in *Proc. of INFOCOM '03*, 2003.