

Securing MANET Multicast Using DIPLOMA

Mansoor Alicherry Angelos D. Keromytis

Department of Computer Science, Columbia University

Abstract. Multicast traffic, such as live audio/video streaming, is an important application for Mobile Ad Hoc Networks (MANETs), including those used by militaries and disaster recovery teams. The open nature of multicast, where any receiver can join a multicast group, and any sender can send to a multicast group, makes it an easy vehicle for launching Denial of Service (DoS) attacks in resource-constrained MANETs. In this paper, we extend our previously introduced DIPLOMA architecture to secure multicast traffic. DIPLOMA is a *deny-by-default* distributed policy enforcement architecture that can protect the end-host services and network bandwidth. DIPLOMA uses capabilities to provide a unified solution for sender and receiver access control to the multicast groups, as well as to limit the bandwidth usage of the multicast group. We have extended common multicast protocols, including ODMRP and PIM-SM, to incorporate DIPLOMA. We have implemented multicast DIPLOMA in Linux, without requiring any changes to existing applications and the routing substrate. We conducted an experimental evaluation of the system in the Orbit MANET testbed. The results show that the architecture incurs limited overhead in throughput, packet loss, and packet inter-arrival times. We also show that the system protects network bandwidth and the end-hosts in the presence of attackers.

1 Introduction

Multicast enables delivery of information from one source to many destinations efficiently, without the source to unicasting to individual destinations. In multicast, nodes send packets over a link only once. They create copies of the packet, and send to multiple links when the packets need to go on multiple links to reach destinations. Multicasting is used for content distribution applications, like audio and video streaming.

Mobile ad-hoc networks are increasingly used in tactical military and civil rapid-deployment networks, including emergency rescue operations and disaster relief network, due to their flexibility in deployment. Audio and video content distribution is an important application on these networks, making support for multicast an absolute necessity. Multicast also improves the efficiency of wireless links in MANETs, due to the broadcast nature of the medium.

The set of nodes receiving the messages that are addressed to a common *multicast address* form a *multicast group*. Traditionally, there are three properties of multicast group [3]:

1. All the members receive all the packets send to the multicast group.
2. Any node can join the multicast group.

3. Any node can send packet to the multicast group.

All these properties have security implications, and there are solutions proposed for them in the context of the (wired) Internet. Most of these solutions differentiate the routers from the receiver nodes (or multicast group members), as it is the case in wired networks. The routers are secure and well behaved. These solutions are not suitable for MANETs, since the nodes play the dual role of receivers (and senders) of the traffic and routers for forwarding other node's traffic. Furthermore, exploiting these properties increase the resource usage, making multicast an easy tool for launching denial of service attacks on resource constrained MANETs. In this paper, we propose extensions to DIPLOMA architecture, which stands for **DI**stributed **PO**licy **enFO**rce**ME**nt **AR**chitecture, to provide multicast security in MANETs.

DIPLOMA is a deny-by-default architecture [2] that enforces trust relationships and traffic accountability between mobile nodes through a distributed policy enforcement scheme for MANETs. In that architecture, capabilities propagate both access control rules and traffic-shaping parameters that should govern a node's traffic. In the deny-by-default, model nodes can only access the services and hosts they are authorized for by the capabilities given to them. The enforcement of the capability is done in a distributed manner by all the nodes in the path from the source to the destination. Compromised or malicious nodes cannot exceed their authority and expose the whole network to an adversary. Upon detection, we can prevent a compromised node from further attacking the network simply by revoking its capabilities. Moreover, that architecture helps mitigate the impact of denial of service (DoS) attacks because excess or unauthorized packets are dropped closer to the attack source. Thus, we avoid unnecessary data processing and forwarding at the target node and the network itself.

Multicast security protocols for wired networks have treated receiver access control and sender access control as two separate problems [10]. Receiver access control is provided using a group policy management system and a group member authorization system [3]. Sender access control can be provided using source specific multicast (SSM), in which only single source can transmit to a multicast group. A MANET node's IP address can change when moving between networks, and requires explicit sender access control. Furthermore, because of the broadcast nature of the medium, it is much easier to do IP address spoofing in MANETs.

In this paper, we provide a unified solution for both receiver access control and sender access control for MANETs by extending DIPLOMA to secure multicast traffic. We define capabilities for use with multicast traffic. There are separate capabilities defined for sending and receiving multicast traffic. A node will not be able to send, or join the multicast group without possessing these capabilities. These capabilities also provide bandwidth constraints for the multicast sessions, preventing resource hogging by the multicast group members. The nodes in MANET enforce the access control and bandwidth constraints of the capability in a distributed manner. We propose modifications to multicast protocols to incorporate capabilities and show the modifications for two popular

multicast routing protocols On Demand Multicast Routing Protocol (ODMRP) and Protocol Independent Multicasting Spare Mode (PIM-SM).

We implement the multicast DIPLOMA on Linux. Our implementation does not require any changes to existing multicast applications or the PIM-SM multicast daemon. However, the applications see the benefit in terms of receiving only the authorized traffic, and being able to send the allocated bandwidth even in the presence of rogue nodes that are trying to conduct a DoS attack.

We implement our system in the Orbit Lab testbed. We conduct extensive experiments to evaluate the performance and effectiveness of our system. We show that multicast DIPLOMA incurs minimal overhead in terms of throughput, packet loss and inter-arrival times. We also study the effect on video streaming in our system. Finally, we show that multicast DIPLOMA is effective against attackers. Note that we do not address confidentiality of the multicast messages. *Group key encryption* is used to encrypt the multicast traffic using symmetric keys. Group key management is used for efficient re-keying for dynamic group memberships [3].

We describe the DIPLOMA architecture in Section 2, the threat model in Section 3, its extension to multicast in Section 4, and the implementation in Section 5. We describe our experimental methodology and results in Section 6. Section 7 discusses related work.

2 System Architecture

2.1 DIPLOMA Overview

In our architecture, one or more pre-defined nodes act as a *group controller* (GC), which is trusted by all the group nodes. A GC has authority to assign resources to the nodes in MANET. This resource allocation is represented as a credential (capability) called *policy token*, and it can be used to express the services and the bandwidth a node is allowed to access. They are cryptographically signed by the GC, which can be verified any node in the MANET.

When a node (initiator) requests a service from another MANET node (responder) using the policy token assigned to the initiator, the responder can provide a capability back to the initiator. This is called a *network capability*, and it is generated based on the resource policy assigned to the responder and its dynamic conditions (*e.g.*, level of utilization).

Figure 1 gives a brief overview of DIPLOMA. All nodes in the path between an initiator to a responder (*i.e.*, nodes relaying the packets) enforce and abide by the resource allocation encoded by the GC in the policy token and the responder in the network capability. The enforcement involves both access control and bandwidth allocation. A responder accepts packets (except for the first) from an initiator only if the initiator is authorized to send, in the form of a valid network capability. It accepts the first packet only if the initiator’s policy token is included. An intermediate node will forward the packets from a node only if they have an associated policy token or network capability, and if they do not violate the conditions contained therein. Possession of a network capability does not imply resource reservation; they are the maximum limits a node can use.

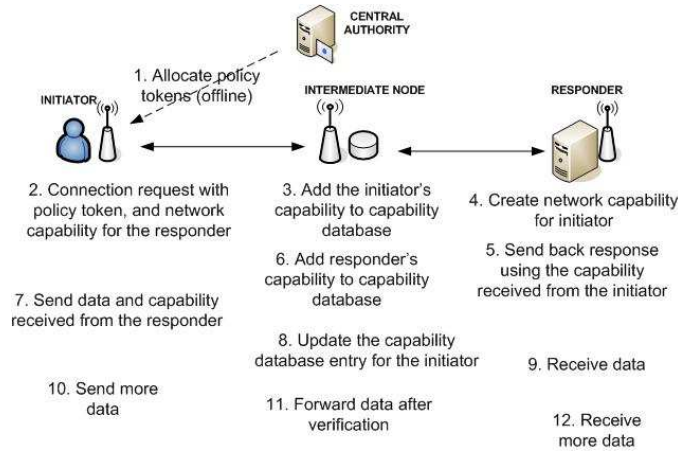


Fig. 1. System overview

Available resources are allocated by the intermediate nodes in a fair manner, in proportion to the allocations defined in the policy token and network capability.

The capability need not be contained in all packets. The first packet carries the capability, along with a transaction identifier (TXI) and a public key. Subsequent packets contain only the TXI and a packet signature based on that public key. Intermediate nodes cache policy tokens and network capabilities in a *capability database*, treating them as soft state. A capability database entry contains the source and the destination addresses, TXI, the capability, public key for the packet signature and packet statistics. Capability retransmissions update the soft state of intermediate nodes when the route changes due to node mobility. The soft state after a route change is also updated using an on-demand query for the capability database entry from the upstream nodes.

2.2 Multicast Capability

DIPLOMA use multicast capabilities for access control and bandwidth limitations. They have same syntactic structure as unicast capabilities [2].

```

serial: 1307467
owner: unit01.nj.army.mil (public key)
destination: 225.1.1.8
service: video
bandwidth: 512kbps
expiration: 2010-12-31 23:59:59
flags: MCAST RW
issuer: captain.nj.army.mil
signature: sig-rsa 23455656769340646678

```

The above represents a policy token assigned by node captain.nj.army.mil to unit01. This is a multicast capability, since the destination address is a multicast address. Unit01 can multicast video traffic up to 512 kbps to the group 225.1.1.8.

There are two types of multicast capabilities: **Multicast Send Capability (MSC)** and **Multicast Receive Capability (MRC)**. The flags in the capability indicate the type of the multicast capability. The nodes possessing a MSC can send the traffic to the multicast group, limited by the bandwidth allocation on the capability. They can also join the multicast group and receive the traffic from the group. The nodes possessing a MRC can join the multicast group only to receive data; They do not have authority to send data to the group.

The group controllers allocate MSCs, and hence they are of type policy tokens. MRCs can be either a policy token or a network capability. The group controller or a sender that has authority in the form of a policy allocates them.

3 Threat Model

Our goal is to protect network resources and the multicast traffic from denial of service attacks, and to enforce access control rules in the absence of a fixed topology. Thus, we want a receiver node to be able to access only the multicast services it is entitled to, and to limit the amount of traffic that can be sent to any multicast group by the authorized senders. To preserve bandwidth and power, we need to filter any unauthorized traffic early on.

We assume MANET environments where an adversary may be an existing node that has been compromised (insider) or a malicious external node that might want to participate in the MANET. In addition, there may be multiple cooperating adversaries; and compromised nodes may not be detected as such immediately, or ever (depending on their actions).

The resources needed to access a service are allocated by the *group controller(s)* (GCs) of the MANET. Group controllers are nodes responsible for maintaining the group membership for a set of MANET nodes, and *a priori* authorize communications within the group. This means that GCs do not participate in the actual communications, nor do they need to be consulted by nodes in real time; in fact, if they distribute the appropriate policies ahead of time, they need not even be members of the MANET. In most cases, the GC may be reachable through a high-energy-consumption, high-latency, low-bandwidth long-range link (*e.g.*, a satellite connection); interactions in such an environment should be kept to a minimum, and only for exceptional circumstances (*e.g.*, for revoking access for compromised nodes).

Without compromising a GC, an external node can participate in a MANET only by stealing the authorization credentials that are bound to the identity of a legitimate node. Because we envision GCs as being primarily offline or, at best, intermittently reachable (with respect to the MANET), we are not addressing the issue of compromised controllers in this paper.

If a node is compromised, an adversary can only access the services and bandwidth that node is authorized to access. If other MANET nodes are adhering to our architecture, a compromised node does not have the ability to disrupt or interfere with end-to-end service connectivity and other nodes beyond its local radio communication radius. The nodes providing services will receive only the

traffic that the compromised node is authorized to transmit, unless the adversary is in the local communication radius.

4 DIPLOMA for Multicast Protocols

Unlike the unicast implementation of DIPLOMA, multicast implementation depends on the underlying multicast protocol used. This is because a multicast forwarding node does not know about the receiver nodes to enforce the multicast receive capability, without interfacing with the multicast routing protocol. Hence, our implementation influence the protocol by snooping and filtering the packets, even though it does not directly modify the multicast protocol processing modules. DIPLOMA may also modify the packet immediately before the packet is sent to the physical interface and immediately after it is received on the interface.

There are two types of multicast routing protocols. The first type is flooding based protocols, where the multicast tree is created for the entire topology based on flooding. Later, part of the tree that does not have any receivers is pruned by explicit prune or status discovery messages. An example of this type of protocol is Protocol Independent Multicasting in Dense Mode (PIM-DM). This type of protocol is useful when most of the nodes in the network are members of the group. The second type, which is more predominant, creates a tree (or mesh) based on the membership. A branch in the tree is created only if there a node in that branch that wants to receive the multicast traffic from the group. There is no wasted data bandwidth in this protocol, even though efficiency of the bandwidth usage depends on the type of the tree construction. Examples of this type of protocol include Protocol Independent Multicasting in Sparse Mode (PIM-SM), MAODV, ODMRP *etc.* In this paper, we focus on implementing the DIPLOMA on this type of protocols.

The receivers are required to send explicit messages to join the multicast tree. This message may traverse multiple intermediate nodes to reach the tree or the node in charge of constructing the tree. Depending on the protocol, the intermediate node may directly forward this message, or send a different message to the same effect to the upstream node. We call these messages collectively as **Join-Tree** messages. In PIM-SM protocol, Join-Tree messages constitute IGMP *membership report* message, as well as *Join/Prune* message. In ODMRP protocol, it is the *Join Reply* message serving this role. In DIPLOMA, we make use of Join-Tree messages to send the MRCs. The nodes drop the Join-Tree messages that do not contain the valid MRCs. When there are multiple downstream receivers, the forwarding node needs to send only one of the MRCs upstream.

Join-Tree messages forwarded by a node may contain the MRC of its downstream node, instead of its own. This happens when the node is just a forwarding node but not a member of the multicast group. To avoid MRC reuse by rogue forwarding nodes for future multicast sessions, the receivers add an expiration time to the MRC in Join-Tree messages. Receivers sign the (capability, timestamp) tuple with their public key. We call that message *time stamped MRC*.

Many multicast protocols have explicit messages initiated by the sender to form the tree. For example, ODMRP has *Join Query* message send by the sender to initiate the tree creation. Not all protocols have this mechanism. For example, PIM-SM does not require the sender to join the multicast group for sending multicast packets. Hence, we do not rely on any of the protocol messages to send the MSC. Instead, we send the MSC when data traffic starts flowing, like in the unicast case. This has an advantage of treating multicast and unicast data the same way, independent of the underlying protocol. To provide added security, we also send the MSC on the protocols that require explicit tree create message from the sender.

An intermediate node forwards a multicast data packet only if both of the following conditions are satisfied:

1. The data packet has an associated MSC from the sender in the node's capability database, and the data packet is conformant to the capability in the form of valid packet signature and the bandwidth constraints.
2. The node has a valid multicast receive capability from one of the receivers in the downstream path. The intermediate node forwards the packet on an interface only if it has a time stamped MRC for a receiver that is reachable on that interface.

A receiving node may leave the multicast tree in two ways depending on the multicast protocol. Some protocols support explicit leave messages. Since it may not be always possible to send a leave message (*e.g.*, the receiver node crashed), the protocols also has periodic membership query. When the receiver node receives a query, it sends some form of a Tree-Join message. In a DIPLOMA enabled systems, the receive node also sends a time stamped MRC in those messages. Then the intermediate (forwarding) node forwards one of the time stamped MRC to the upstream node in its Tree-Join message. When a node does not receive any time stamped MRCs from the downstream nodes on an interface, that interface is pruned from the multicast tree (or mesh).

4.1 Security Analysis

We now discuss how our architecture relates to the threat model described in Section 3.

Since the capabilities are signed by a GC and are verifiable by all nodes, adversaries cannot generate their own valid capabilities. Adversaries can create valid capabilities only if the GC is compromised. Since the individual packets are signed, an adversary cannot use a transaction id that does not belong to it to transmit packets.

A compromised or malicious node that does not enforce the capability protocol can only have impact within its communication radius. Packets generated without the capability or with a snooped transaction id by a malicious node will be dropped by the neighboring nodes due to invalid signatures. A compromised node can only access the services it is authorized to. Packets of nodes trying to

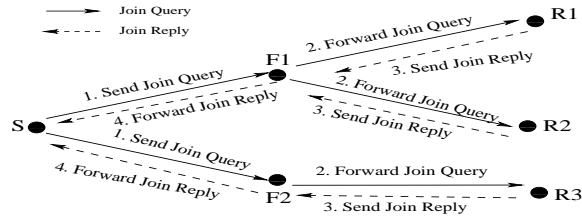


Fig. 2. ODMRP Protocol

use more bandwidth than is allocated to them will be rejected. A malicious node frequently doing this can be detected and isolated.

A multicast receiver can only join the multicast groups for which it possesses a MRC. Similarly, a multicast sender can send traffic only to the groups for which it possesses a MSC. Furthermore, this send traffic is limited by the bandwidth constraints of the MSC. Only the links which are part of the multicast tree or mesh actually carry the multicast traffic. Since the packets are signed, any injection of packets into a data stream is easily detectable by the nodes in the path.

4.2 DIPLOMA on ODMRP

Figure 2 gives a high level overview of the On Demand Multicast Routing Protocol (ODMRP) [13] protocol. There is a sender node S that wants to multicast data into a group. Three receiver nodes $R1$, $R2$ and $R3$ are part of the multicast group. Two nodes $F1$ and $F2$ are in the path from the node S to the receivers. We call those nodes as intermediate nodes. When the node S has data to multicast, it broadcasts a *Join Query* message to the neighboring nodes to discover a multicast tree. This message is received by the intermediate nodes $F1$ and $F2$, which in turn broadcasts to their neighbors. The nodes $R1$ and $R2$ receives the *Join Query* from the node $F1$, and the node $R3$ receives the *Join Query* from the node $F2$. The receiver nodes send a *Join Reply* message back to the nodes from which it received the *Join Query* message (i.e. the upstream nodes $F1$ and $F2$). Once the nodes $F1$ and $F2$ receive the *Join Reply* messages they become part of the *forwarding group* and forwards the *Join Reply* messages to node S .

In DIPLOMA systems that are running over ODMRP protocol, *Join Query* messages are modified to contain the MSC of the sender, and the transaction id and the key that will be used by the sender for subsequent communication. The intermediate nodes store this capability information temporarily and forward the *Join Query* message to its neighbors. On receiving this *Join Query*, a receiver node in the multicast group responds with a *Join Reply* message. This *Join Reply* message is modified to contain the receiver's time stamped MRC that authorizes the node to be part of the multicast group. On receiving a *Join Reply*, the intermediate node becomes part of the *Forwarding Group*. The intermediate node installs the saved MSC in its capability database. The intermediate node then forwards the *Join Reply* to its upstream node (i.e. towards the sender). It is possible for the intermediate node to receive *Join Replies* from multiple receivers

with different MRCs. The intermediate node needs to forward only one of them to its upstream node. Then the forwarding node starts forwarding the multicast data traffic to the downstream nodes. Similar to the unicast case, the forwarding nodes enforce the capability for all the multicast packets.

Whenever the time stamped MRCs expire, a forwarding node stops forwarding any multicast packet received by the node. ODMRP is a stateless protocol that does not have any multicast leave or prune messages. Instead, the tree is valid only for certain duration. The tree is completely dissolved when that timer expires. Furthermore, the receiver nodes can respond with Join-Reply messages only when it receives Join-Request message from a sender. There is no mechanism for a new receiver to add itself to an existing multicast tree. The sender maintains the multicast tree, and adds new receivers by periodically sending the Join-Query message. The DIPLOMA keeps the time stamped MRC up to date through this periodic tree maintenance protocol. Whenever a receiver gets a new Join-Query message, it creates a new time stamped MRC to respond back in the Join-Reply. To maintain continuous multicast data session, it is important for the period in which a new Join-Query is generated to be less than the validity duration of the time stamped MRC.

4.3 DIPLOMA on PIM-SM

Protocol Independent Multicast - Sparse Mode (PIM-SM) is a popular multicast routing protocol that is independent of the underlying unicast protocol. This protocol works in conjunction with the Internet Group Membership Protocol (IGMP). The protocol explicitly creates a tree from the sender to the receivers.

In PIM-SM, one of the router is designated as a *Rendezvous Point (RP)* for a multicast group. All the other routers need to join the group through RP. Whenever a node wants to join a multicast group, it conveys the message through an IGMP *membership report* message. A *designated router (DR)* for the node sends periodic PIM Join/Prune messages towards the RP for the multicast group. Each router along the path to RP updates the packet forwarding state (routing entries) and sends the Join/Prune message towards the RP.

Whenever a node wants to send the traffic to the multicast group, its DR encapsulates the data in PIM *Register* messages and unicasts it to the RP. The RP decapsulates the message and sends the data towards the receivers in the multicast tree. If the data-rate from the sender is high, then the RP sends a source specific Join/Prune message towards the sender. This extends the tree to the sender, and the sender can directly send multicast messages to the tree without encapsulating the messages. If the data rate warrants it, any DR can join source specific shortest path tree by sending a Join/Prune message towards the sender, and prune the shared tree towards the RP.

We can enable DIPLOMA in multicast systems running PIM-SM by including the multicast capabilities in the IGMP and PIM messages. Whenever a receiver sends an IGMP membership report message, its timestamped MRC is included. DIPLOMA systems reject any membership report without the capability. A DR includes one of the time stamped capabilities of the downstream

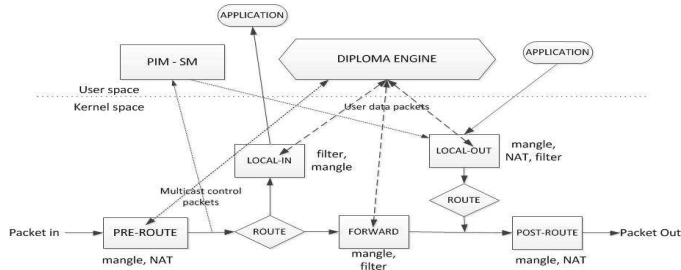


Fig. 3. DIPLOMA Implementation

nodes in the Join/Prune messages it sends towards the RP or the source node. When a router receives a prune message, the corresponding time stamped MRC is removed from its tables. The node stops forwarding the packets, when it does not have any valid time stamped MRC from the downstream receivers.

The multicast packets are sent similar to unicast case. Before sending the packet, the sender multicasts it's MSC in a capability request packet with the capability, transaction identifier and the key for the packet signatures. This packet goes to the RP as a regular multicast packet or a Register packet; the RP in turn sends the packet to the multicast group (after decapsulation for the Register packets). All the nodes in the multicast tree add the capability to their capability database. If it is a register packet, then the nodes in the path between the sender and the RP will also extract the capability, transaction id and the key for the signature from the capability request, and install in their database. Any subsequent data packet multicast by the sender contains the transaction id and the packet signature. The signature is verified and the bandwidth is enforced by all the nodes in the multicast tree, and by the nodes between the sender and the RP in the case of the Register packets.

If a receiver node joins the multicast tree after the transmission of the initial capability request packet by the sender, then it will not be able to validate the multicast data packets. DIPLOMA solves this by two means: Firstly, the sender periodically multicasts the capability request packet. The new receiver node can start accepting the data packets after the periodic multicast. Secondly, the receiver sends a request for the capability towards the sender using a DIPLOMA control (or error) packet. On receiving this request, either an intermediate node or the sender responds with the capability and the public key for the signature.

5 Linux Implementation

We now describe the implementation of Multicast DIPLOMA on Debian Linux system running kernel 2.6.30. For multicast routing, we use *pimd*, a PIM-SM package that comes with the Debian distribution. Since PIM-SM requires a separate unicast routing, we use University of Uppsala's AODV implementation called AODV-UU. Our implementation does not require any changes to the application program, routing module or PIM-SM daemon.

The multicast DIPLOMA is implemented as a user level process, called DIPLOMA engine that interfaces with rest of the Linux packet processing subsystem using netfilter framework. We use netfilter queue to receive, modify, and filter packets in the DIPLOMA engine.

Figure 3 shows how the DIPLOMA engine interfaces with netfilter subsystem. A brief description of Netfilter framework and how the DIPLOMA uses it for handling the unicast traffic can be found in [1]. The dotted lines are the hooks used only for the multicast traffic. The solid lines show the hook for both unicast and the multicast traffic. The reason for requiring additional hooks for the multicast traffic is due to the implementation of PIM-SM in Linux. It uses raw sockets to send and receive traffic; these packets do not go through INPUT and OUTPUT hooks, but traverse PREROUTING and POSTROUTING hooks.

Next, we describe the packet flow for control (*i.e.*, IGMP and PIM packets) and multicast data packets.

5.1 Membership Messages

When the system sends a membership message, in the form of an ICMP message or a PIM Join/Prune message, the DIPLOMA engine receives on the packet on OUTPUT hook. It checks for a valid MRC for the message in its database. The valid capability may be either its own capability, or a capability it received from a downstream node. It adds the capability in the packet and sends an ACCEPT verdict on the hook.

When the system receives a membership message on the PREROUTING hook, it validates the packet. A valid packet needs to contain a valid MRC. The node saves the MRC in its tables for subsequent request to the upstream node. The capability is removed from the packet and an accept verdict is given. The PIM-SM daemon receives this packet over the RAW socket. The engine drops any membership message without a valid capability.

5.2 Capability establishment

When a sender needs to multicast data, it creates a transaction identifier for use with subsequent packets to identify the session. It also creates an RSA key for signing the data packets of that session. The sender sends the transaction id, public key and the MSC authorizing the sender to send the multicast packet as a DIPLOMA control message. The DIPLOMA engine sends this message when it first sees a packet from the sender for a multicast group. The application program sending the multicast data need not be aware of this step.

When a multicast member node or a forwarding node receives this message, it validates the capability and stores the transaction id, the public key and the MSC in its capability database. These nodes validate the subsequent data packets coming from the sender against the capability and verify the packet signatures. For updating the new receivers or new intermediate node after a route change, the sender multicasts the capability establishment packet periodically. A receiver node can also request the sender on a unicast message to send the capability establishment packet, when it does not have that information due to late joining or a route change.

5.3 Multicast Data Packets

All the multicast data packets need to contain an associated capability. The DIPLOMA engine at the sender modifies the outgoing packets in the OUTPUT hook by including a capability header, which contains the transaction identifier and the packet signature. The packets sent to a multicast group are treated together as a block for the signature computation [1]. A packet block contains maximum of *block size (P)* packets that are sent within the interval *block timeout (T)*. The packet signatures for a block consist of RSA signature for the first packet and SHA-1 hashes for the remaining packets. The RSA signature is verifiable with the key sent in the capability establishment phase. The SHA-1 hashes are integrity protected by including them in the first packet.

The engine at the intermediate node receives the multicast packet on FORWARD hook. The engine validates the packet against the capability using the transaction identifier. The validation including checking if there is a valid MSC in its database associated with the transaction identifier, if the packet has valid signature, and if the packet conforms to the bandwidth constraints of the capability. If the packet is valid, then the engine gives an ACCEPT verdict for forwarding the packet.

If the packet is destined to the node as a receiver on the multicast group, DIPLOMA receives the packet on the INPUT hook. The engine validates the packet as above, removes the capability header from the packet and gives an ACCEPT verdict, causing the kernel to deliver the packet to the application.

6 Experimental evaluation

In this section, we evaluate the effectiveness of multicast DIPLOMA. First, we compare the throughput, packet loss and inter arrival times of the systems with and without multicast DIPLOMA using periodic traffic. We also study these parameters using real video streaming traces. Finally, we study the effectiveness of DIPLOMA in containing the attacker nodes.

6.1 Testbed

We implemented the multicast DIPLOMA engine as described in Section 5 in Linux systems running Debian Linux with kernel 2.6.30. We use AODV-UU for routing unicast traffic, modified to handle multiple interfaces. For multicast routing, we use PIM-SM implementation called *pimd* that is available with Debian Linux distribution. We run the resulting system on multiple nodes in the Orbit lab¹ wireless testbed. Orbit is an indoor wireless testbed consisting of 400 nodes arranged as a 20x20 grid on a physical area of (20m x 20m). Each node contains 1-GHz VIA C3 processor, 512 MB RAM, a 20 GB hard disk, two wireless mini-PCI 802.11 a/b/g interfaces, and two 100BaseT Ethernet ports.

Since the nodes are within the communication range of each other in Orbit testbed, we use channel hopping to create multi-hop topologies. The traditional methods of MAC address based filtering to create multi-hop topologies are not

¹ <http://www.orbit-lab.org/>

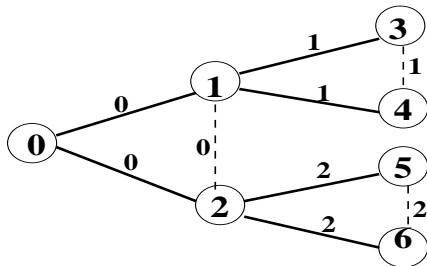


Fig. 4. Tree topology

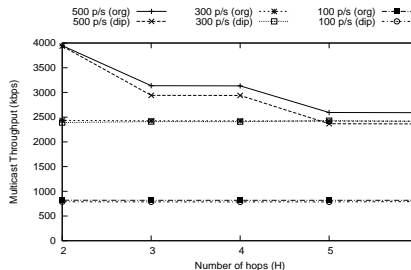


Fig. 5. Throughput for line topology

suitable for studying the security system like DIPLOMA, since an attacker node can cause damage in its communication radius.

Since the DIPLOMA engine is a user-level process, all packets are queued for user-level processing before transmission. To make a fair comparison, we also do a similar queuing of the packet to a user level process on systems not running the DIPLOMA (called *original*). The user level program gives an ACCEPT verdict on all the packets, without any processing.

For measuring the performance of the DIPLOMA, we use two topologies: a line topology and a tree topology. In line topology nodes are allocated channels in such a way that each node can directly communicate only with its neighbors on either side (except for the first and last, which has only one neighbor). In this topology, the first node is the sender of the multicast. All the remaining nodes subscribe to the multicast group. The tree topology is shown on figure 4. The links are labeled with the channel with which the nodes communicate. Here the sender is the node 0 (root), and the multicast receivers are nodes 3,4,5,6 (leaf nodes). In the figure, the solid lines show the multicast tree and the dashed lines shows the links that are not participating in the multicast.

We use the multi-generator tool *mgen* [14] from Naval Research Laboratory to send and receive traffic in our experiments. Each data points in this section represent an average of running six experiments, each experiments sending traffic for 30 seconds each.

6.2 Line Topology

In this set of experiments, we study the performance of DIPLOMA and the original schemes for the line topology. The sender sends periodic traffic of size 1024 bytes at the rate of 100, 300 and 500 packets per second. This corresponds to the rates of 819.2 Kbps, 2.4576 Mbps and 4.096 Mbps respectively.

Figure 5 shows the throughput received by the nodes at different hop lengths for different transmission rates. For the rate of 100 and 300 pkts/sec, both the DIPLOMA and the original schemes receive bandwidth close to the send bandwidth for all the hops. The bandwidth for the DIPLOMA is minimally (0.7% and 3.7% respectively) lower than the original. For the rate of 500 pkts/sec, the received bandwidth reduces as the hop count increases. This is because the available bandwidth decreases as the number of hops increases. The bandwidth

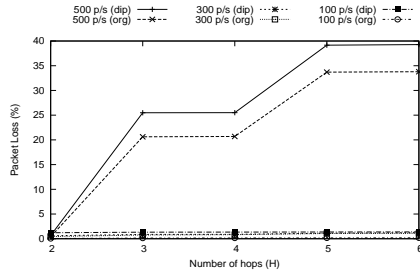


Fig. 6. Packet loss for line topology

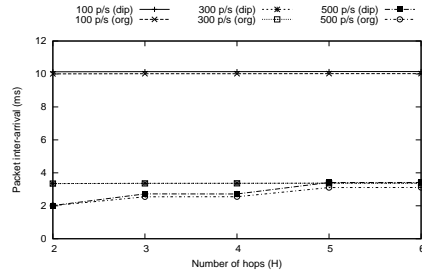


Fig. 7. Packet inter arrival times for line topology

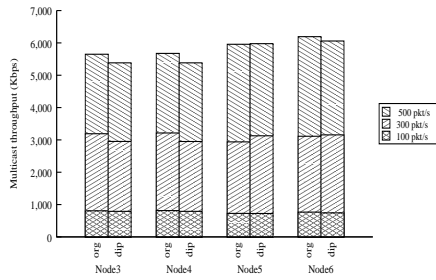


Fig. 8. Throughput for tree topology

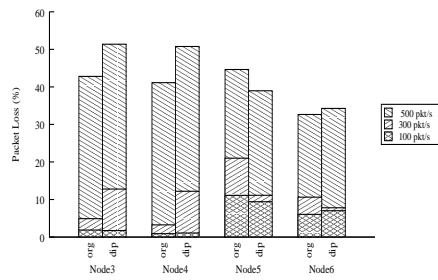


Fig. 9. Packet loss for tree topology

for the DIPLOMA is 6.6% lower than the original. This is due to larger headers and processing required for the DIPLOMA.

Figure 6 shows the packet loss for the same experiment. The packet losses are less than 1% for both the schemes for the rate of 100 & 300 pkts/sec. The packet losses are higher for the rate of 500 kbps, which explains the lower throughput as the hops count increases. The packet loss is about 5% more for DIPLOMA, due to larger headers, which require more bandwidth.

Figure 7 shows the packet inter arrival times for the same experiments. For the rates 100 and 300 pkts/sec, the inter arrival is close to the inverse of their send rate. The inter arrival for diploma is slightly higher than the original, due to larger processing required. For the 500 pkts/sec rate, inter-arrival time increases with hop count, due to correspondingly higher packet loss.

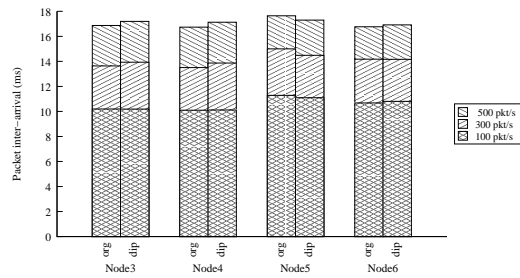


Fig. 10. Packet inter arrival times for tree topology

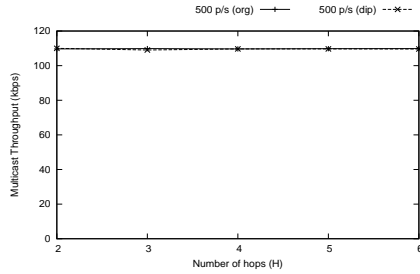


Fig. 11. Streaming video throughput for line topology

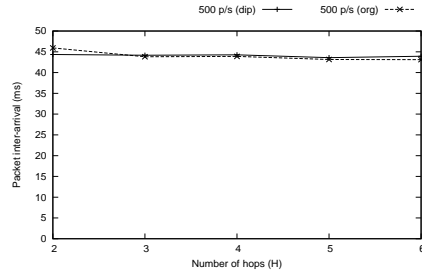


Fig. 12. Streaming video packet inter-arrival times for line topology

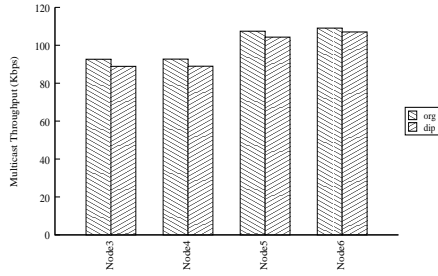


Fig. 13. Streaming video throughput for the tree topology

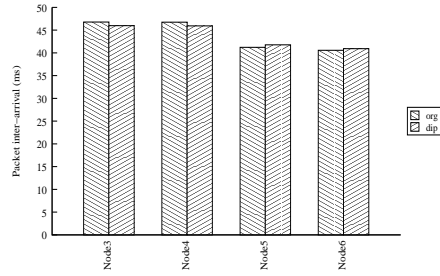


Fig. 14. Streaming video packet inter-arrival times for the tree topology

6.3 Tree topology

We study the throughput, packet loss and inter arrival times for the tree topology given in Figure 4. Here the root node 0 is the sender and the leaf nodes 3,4,5,6 are the receivers.

Figure 8 shows the throughput at the nodes for both the schemes. Though the nodes are at same distance from the root, they receive different bandwidths. This may be because of the channel conditions and the packet scheduling. For some nodes, the DIPLOMA receives higher bandwidth than the original. The sum of the bandwidth received by all the four receivers is slightly higher for the original scheme compared to the DIPLOMA scheme. This total bandwidth is 2.1%, 2.1% and 3.6% higher respectively for the rates 100, 300 and 500 pkts/sec for original compared to DIPLOMA.

Figure 8 shows the packet loss at the nodes for both the schemes. Unlike the line topology, there was some packet losses (6% to 9%) for the rates of 100 and 300 pkts/sec for both the schemes on some of the nodes. This may also be due to channel conditions.

Figure 10 shows the packet inter arrival times for both the schemes. Here also for some receivers, the inter arrival times were shorter for the DIPLOMA. However, on average, the inter arrival times for the DIPLOMA was slightly higher than the original, due to larger processing delays and the extra headers in DIPLOMA.

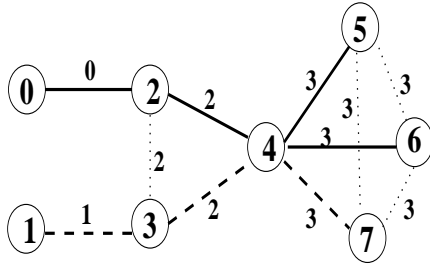


Fig. 15. Attack topology

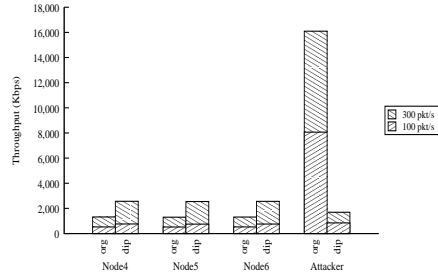


Fig. 16. Throughput in presence of unicast attacker

6.4 Streaming video

In this set of experiments, we study the performance of streaming video. The experiments were conducted by creating a trace of streaming video using *evalvid* [12], and sending that packets based on that trace using *mgen*.

Figures 11 and 12 shows the throughput and the inter arrival times for the streaming video for the line topology. The results show that both the DIPLOMA and the original schemes receive the full bandwidth the video, and the packets are received at constant inter-arrival times.

Figures 13 and 14 shows the results for the tree topology. There was a small loss in two of the nodes for both the schemes. This behavior is similar to the results for the periodic traffic.

6.5 Attacker resiliency

Now we study the effectiveness of multicast DIPLOMA in containing attackers. We use the topology given in figure 15. The solid lines show the multicast tree and the dashed lines show the unicast path. The labels on the links show the channels. In the experiments below, the nodes 0 and 1 are the senders. These nodes have only its neighboring nodes 2 and 3 respectively in its communication radius. Hence only nodes 2 or 3 cannot be protected by DIPLOMA, when these nodes misbehave at Physical or MAC layer.

We study how DIPLOMA can protect multicast sessions when there is a DoS attacker sending high-rate traffic. Node 1 (attacker) sends periodic traffic of size 1024 bytes at the rate of 1000 packets per second (*i.e.*, rate of 8.19 Mbps) to node 7. The allocated bandwidth for the attacker was 1 Mbps. At the same time, node 0 multicasts to receivers nodes 4, 5 and 6 a periodic traffic of size 1024 at rates 100 pkts/s (*i.e.*, 819.2 Kbps) or 300 pkts/s (*i.e.*, 2.45 Mbps).

Figure 16 shows the throughput at the three multicast receivers and the unicast receiver (attack traffic). In DIPLOMA, the attacker is able to achieve a bandwidth of 844 Kbps, which is the allocated bandwidth (minus the overhead). The multicast receivers receive close to their send bandwidth. The multicast receivers receive on average 749 Kbps and 1.80 Mbps respectively for 100 and 300 pkt/s traffic. For the original scheme, the attacker is taking up most of the

bandwidth, at 8.04 Mbps. The multicast traffic receives only a fraction of its send bandwidth. The multicast receivers receive on average only 517 kbps and 788 kbps respectively for 100 pkt/s and 300 pkt/s traffic.

7 Related Work

The concept of capabilities was used in operating system for securing resources [17]. Follow-on work investigated the controlled exposure of resources at the network layer using the concept of “visas” for packets [7], which is similar to network capabilities. More recently, network capabilities were proposed to prevent DoS in wired networks [4]. We extend the concept to MANET and use it for both access control rules and traffic shaping parameters.

A survey of security issues and solutions for multicast in wired networks is presented in [3]. They classify the issues and solution based on the three properties described in Section 1. The solutions are specific to wired networks and not directly applicable to MANETs, which have no specialized router nodes.

A number of solutions have been proposed for multicast receiver access control [10, 9, 5]. These solutions have trusted routers or query centralized servers, thus neither is suitable for MANETs. These protocols do not also have limitations on the amount of service accessed. DIPLOMA provides a unified solution to both receiver and sender access control, and supports bandwidth constraints.

There are a number of multicast routing protocols proposed for MANETs. A survey of these protocols is present in [6]. There has been work dealing with security issues of these protocols. A discussion of possible attacks on MAODV (Multicast-extended AODV) routing can be found in [15]. The authors also propose an authentication framework to protect an MAODV network against these attacks. Tactical MAODV [16] extends MAODV through the integration of the security services necessary for the tactical deployment of MANETs, such as forward and backward secrecy and data confidentiality. In [8] authors extend their multicast MANET protocol MMARP with digital signatures using a public key scheme. They use Cryptographically Generated Addresses (CGA) to keep attackers from impersonating other nodes. [11] introduces a protocol for secure communication in multicast groups with a pair of multicast trees for each multicast group; one for security information and the other for data traffic.

8 Conclusions and Future Work

We presented multicast DIPLOMA, an architecture for securing multicast traffic in MANETs. DIPLOMA is a deny-by-default, distributed policy enforcement architecture based on network capabilities. It prevents unauthorized senders from sending packets to a multicast group, and unauthorized receivers from joining the multicast group, protecting the end-host resources and the network bandwidth. We showed that popular multicast protocols such as ODMRP and PIM-SM can be modified to incorporate DIPLOMA. We implemented the DIPLOMA in Linux running PIM-SM without any changes to applications and routing. We evaluated the system on the Orbit MANET testbed. We showed that the impact of the

scheme is minimal on throughput, packet loss, and packet inter-arrival times. We also showed that DIPLOMA allocates resources in a fair manner even in the presence of attackers, protecting legitimate traffic.

Acknowledgements

This work was supported in part by the National Science Foundation through Grant CNS-07-14277 and by ONR through Grant N00014-09-10757. Any opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the NSF, ONR, or the US Government.

References

1. M. Alicherry and A. D. Keromytis. DIPLOMA: Distributed Policy Enforcement Architecture for MANETs. *International Conference on Network and System Security*, September 2010.
2. M. Alicherry, A. D. Keromytis, and A. Stavrou. Deny-by-Default Distributed Security Policy Enforcement in Mobile Ad Hoc Networks. *SecureComm*, September 2009.
3. P. J. M. Ammar. Security issues and solutions in multicast content distribution: A survey. *IEEE Network*, 17, 2003.
4. T. Anderson, T. Roscoe, and D. Wetherall. Preventing Internet Denial-of-Service with Capabilities. *Proc. of Hotnets-II*, 2003.
5. A. Ballardie and J. Crowcroft. Multicast-Specific Security Threats and Countermeasures. *SNDSS*, 1995.
6. C. M. Cordeiro, H. Gossain, and D. Agrawal. Multicast over Wireless Mobile Ad Hoc Networks: Present and Future Directions. *IEEE Network*, 17, 2003.
7. D. Estrin, J. C. Mogul, and G. Tsudik. Visa protocols for controlling interorganizational datagram flow. *IEEE JSAC*, May 1989.
8. F. J. Galera, P. M. Ruiz, A. F. Gomez-Skarmeta, and A. Kassler. Security Extensions to MMARP Through Cryptographically Generated Addresses. *Lecture Notes on Informatics*, 2005.
9. T. Hardjono and B. Cain. Key Establishment for IGMP Authentication in IP Multicast. *IEEE ECUMN*, 2000.
10. P. Judge and M. Ammar. Gothic: A Group Access Control Architecture for Secure Multicast and Anycast. *INFOCOM*, 2002.
11. T. Kaya, G. Lin, G. Noubir, and A. Yilmaz. Secure Multicast Groups on Ad Hoc Networks. *ACM Workshop on Security of Ad Hoc and Sensor Networks*, 2003.
12. J. Klaue. EvalVid - A Video Quality Evaluation Tool-set. <http://www.tkn.tu-berlin.de/research/evalvid/>.
13. S.-J. Lee, M. Gerla, and C.-C. Chiang. On-Demand Multicast Routing Protocol. October 1999.
14. Naval Research Laboratory. Multi Generator (MGEN). <http://cs.itd.nrl.navy.mil/work/mgen/>.
15. S. Roy, V. G. Addada, S. Setia, and S. Jajodia. Securing MAODV: Attacks and Countermeasures. *IEEE Intl. Conf. SECON*, 2005.
16. D. Slezak, T. Kim, A. C. Chang, T. Vasilakos, M. Li, and K. Sakurai. Security in Tactical MANET Deployments. *Comm. and NetInt. Conf., FGCN/ACN*, 2009.
17. E. Wobber, M. Abadi, M. Burrows, and B. Lampson. Authentication in the Taos Operating System. *ACM Trans. on Computer Systems*, 12, February 1994.