Efficient Security Mechanisms For Overlay Multicast-based Content Distribution

Sencun Zhu*,Chao Yao†, Donggang Liu‡, Sanjeev Setia† and Sushil Jajodia†
* Department of Computer Science and Engineering and School of Information Sciences and Technology, The Pennsylvania State University, University Park, PA 16801
† Center For Secure Information Systems, George Mason University, Fairfax, VA 22030
‡ Department of Computer Science, North Carolina State University, Raleigh, NC27695
Email:szhu@cse.psu.edu,cyao,setia,jajodia@gmu.edu,dliu@unity.ncsu.edu

***Abstract*— This paper studies the security issues that arise in an overlay multicast architecture where service providers distribute content such as web pages, static and streaming multimedia data, realtime stock quotes, or security updates to a large number of users. In particular, two major security problems of overlay multicast, *network access control* and *group key management*, are addressed. We first present a bandwidth-efficient scheme, called CRBR, that seamlessly integrates network access control and group key management. Performance analysis and simulation results show that our scheme incurs much smaller communication overhead than two other well-known schemes. Next we propose a DoS-resilient key distribution scheme, called k-RIP, that delivers updated keys to a large fraction of nodes with high probability even if an attacker can *selectively* compromise nodes in the multicast data delivery hierarchy. k-RIP does not rely on knowledge of overlay topology, and can therefore scale up to very large overlay networks. An important application of k-RIP is distributing critical messages (e.g., keys, new virus signatures, or certificate revocation lists) to a large number of nodes that are organized into trees, meshes, or other types of graphs.**

Keywords: Overlay multicast, network access control, group key management, DoS-resilient

I. Introduction

We consider the security issues that arise in an overlay multicast architecture where service providers distribute content such as web pages, static and streaming multimedia data, realtime stock quotes, or security updates (e.g., new virus signatures, certificate revocation lists) [21].

Overlay multicast, also called end system multicast or application-level multicast [12], was proposed as a new group communication mechanism in place of IP multicast whose deployment has been very slow due to both technical and operational concerns. Recently several studies [3], [12], [11], [18], [38] have investigated research problems in overlay multicast such as algorithms for tree or mesh construction, routing, reliability, and resource allocation. However, security issues in overlay multicast have received relatively little attention so far. Previous work on overlay network security either investigates the impact of selfish cheating nodes on the performance of overlay multicast trees [22], or investigates schemes that improve the fault-tolerance or denial-of-service(DoS) resilience of overlay networks by introducing path redundancy [5], [31], [35], [36].

Contributions: We concentrate on two major security problems of overlay multicast: *network access control* and *group key management*. In IP multicast, network access control and group key management were considered as two *independent* issues and they were studied *separately*, one in the network layer [16] and the other in the application layer [7], [23], [27], [28], [32], [33], [37]. In this paper, we propose a bandwidth-efficient scheme called CRBR that seamlessly integrates network access control with group key management. CRBR exploits the special property of overlay multicast that a node is both a group member and a router. We show through analysis and simulation that CRBR greatly outperforms other two representative group rekeying schemes: LKH [32] and SDR [23] when they are applied in overlay multicast. Moreover, using a queueing model, we show the impact of node presence dynamics (i.e.,coming online/going offline) on the performance of group rekeying protocols.

We also propose a simple but effective DoS-resilient key distribution scheme, called k-RIP (stands for k Random Injection Points), that delivers updated keys to a large fraction of nodes via an overlay network. Specifically, in addition to propagating one copy of updated keys using a multicast tree rooted at the source node, our scheme injects k additional copies of updated keys into the multicast tree through k randomly selected nodes in the network. These selected nodes propagate the message to both their child nodes (if any) and parent nodes, thereby spreading the message over the multicast tree. Our simulation and analysis results show that k-RIP can greatly increase the probability that nodes receive messages even if an attacker can *selectively* compromise nodes in the multicast tree, and it can reduce message propagation latency compared to a scheme in which only one copy of the message is injected via the root node. Unlike previously proposed schemes [5], [31], [35], [36], k-RIP does not rely on a knowledge of the overlay topology. Thus, it is scalable to very large overlay networks. An important application of k-RIP is distributing small-size but critical messages (e.g., keys, new virus signatures) to a large number of nodes that are organized into trees, meshes, or other types of graphs.

Organization: The remainder of this paper is organized as follows. Section II discusses some related work on group key

management, network attacks and countermeasures. Section III describes the system model and our design goal. In Section IV, we present our scheme CRBR for providing both network access control and group key management, followed by its security and performance analysis. In Section V we describe our k-RIP key distribution scheme. Finally, Section VI concludes this paper.

II. Related Work

We introduce the related work in four categories: *group key management*, *membership control*, *network attacks and countermeasures*, and *resilient overlay multicast*.

A. Group Key Management

Group key management has been extensively studied in the context of secure multicast in IP multicast. The previous group rekeying schemes can be categorized into stateful and stateless protocols. The stateful class of protocols includes several protocols based upon the use of logical key trees, e.g., LKH [32], [33], OFT [7], [9], ELK [27]. In these protocols, the key server uses key encryption keys that were transmitted to members during previous rekeying operations to encrypt the keys that are transmitted in the current rekeying operation. Thus, a member must have received all the key encryption keys of interest in all the previous rekey operations to decipher the current group key. Adding redundancy in key distribution [30], [37] does not fully address the issue in the case of burst packet loss or nodes going online/offline frequently. *Stateless* group rekeying protocols [20], [23], [29] form the second class of rekey protocols. In these protocols, a legitimate user only needs to receive the keys of interest in the current rekey operation to decode the current group key. The stateless feature makes these protocol very attractive for applications in which members go offline very frequently. However, these protocols usually have much higher communication overhead than the stateful protocols. Our scheme also provides the stateless property, but it incurs significantly smaller communication overhead than the other schemes. Moreover, it also provides network access control.

B. Membership Control

Gothic [17] is a group access control architecture for secure IP multicast and IP anycast. It also includes group access control aware group key management which reduces the overhead of LKH, assuming that the key server has the global knowledge about the topological locations of every member in the multicast tree and that an attacker does not eavesdrop messages on the networks other than its local network. Our scheme is designed specifically for overlay multicast, and it does not make the above assumptions.

C. Network Attacks and Countermeasures

Mathy et al [22] studied the impact of selfish nodes cheating about their distance measurements in application-level multicast overlay tree. Badishi et al [6] proposed a gossip-based multicast protocol called Drum, which combines multiple techniques such as push, pull, random port selections, and resource bounds, for mitigating DoS attacks in secure gossip-based multicast. Wright et al [35] presented k-redundant depender graphs for distributing public-key certificate revocation lists (CRLs), which provides every node in the graph with k disjoint paths to the root of the graph, thus guaranteeing delivery even when up to $k-1$ paths between them have failed. Song et al [31] improved the scalability of the above scheme by presenting expander graphs for constructing robust overlay networks that have constant degree. Their scheme provides a lower bound on the probability of each node receiving a revocation list when a certain fraction of nodes may fail. Yang et al [36] proposed to augment tree-like hierarchy with hierarchical overlay networks, which is actually also a type of graphs, to achieve DoS resilience.

All these schemes provide stronger fault-tolerance or DoS resilience at the cost of higher (re)construction complexity to maintain their security property, especially when nodes join or leave the tree frequently. Moreover, these schemes are subject to selective attacks in which an attacker can prevent a large number of nodes from receiving messages by compromising (or becoming) the nodes close to the root. Our random injection points scheme directly works with the existing overlay multicast schemes without changing trees into graphs, and it is especially suitable for distributing small-size but critical messages. Our scheme is robust to selective attacks; therefore, we believe that the combination of our scheme with the other DoS-resilient schemes will make a distribution system more robust to DoS attacks.

D. Resilient Overlay Multicast

Banerjee et al [5] introduced a probabilistic forwarding scheme for overlay multicast. In their scheme, every node forwards received packets to a randomly selected set of nodes, assuming that every node has global knowledge of overlay topology or it can discover other nodes on the fly. Our k-RIP scheme also uses randomness, but the randomness is used by the key server to inject packets into the overlay network, not used by the regular nodes to forward packets. The main reason we do not employ their scheme directly is because of scalability consideration. For large-scale and dynamic overlay networks, the overhead for discovering other nodes on the fly or maintaining global topological knowledge would be very large. In k-RIP neither the key server nor the nodes need knowledge on overlay topology. The key server infers the presence of the selected nodes only through some heuristics; the nodes do not know anything about the nodes other than their parent and child nodes. This allows k-RIP to scale up to arbitrarily large overlay networks.

III. System Model and Design Goal

This section describes our system model and design goal. For ease of presentation, we use the terms “join” and “leave” to denote the actions of a subscriber coming online and going offline, respectively, whereas use “add” and “revoke” to denote the actions of becoming a member and cancelling the membership status of a subscriber, respectively.

A. System Model

There are potentially a large number of application scenarios of overlay multicast, which are characterized by different parameters. For example, *group size* could vary from tens to millions, *number of data sources* from one to many, *membership dynamics* from static to frequent subscribing/unsubscribing, *node presence* from always online to frequently going offline, and *types of traffic* from real-time information to delay-insensitive bulk data. It seems unlikely that a single system model can describe all these scenarios. Therefore, we focus on a specific application scenario, which we believe is (or will be) very representative. Security solutions for this scenario can be applied to many other scenarios as well.

We consider a commercial application of overlay multicast, in which a service provider distributes data (e.g., live content or streaming media) to a large number of subscribers (also called member nodes hereafter) simultaneously. For simplicity, we assume that online nodes are self-organized into an overlay multicast delivery tree rooted at the distribution server of a service provider, although our security schemes work for various distribution infrastructures, such as trees, meshes, or other types of graphs. The algorithms for constructing and maintaining overlay multicast trees [3], [12], [18] are out of the scope of our work.

The population of the system could be up to hundreds of thousands or even millions of nodes. We assume that a node may join or leave a multicast group very frequently and at any time. For example, a user may subscribe to multiple service providers for different programs. She may switch between multiple channels to find an interesting program to join. A user may also leave a channel immediately after she has received the data of interest to her.

In this model, the service provider has three types of servers playing different roles. A key server (or many key servers for scalability) provides subscription services to users. Before a user is able to join the group for the first time, it needs to subscribe to the key server (e.g., through a website). After successful subscription based on certain policies or rules (e.g., agreeing to pay service fee), a user is provided with a service credential that allows it to join the multicast delivery tree later. A user must also contact the key server to cancel its membership later if it wants. The key server also manages the update of data encryption keys (DEKs). When it changes its DEK, it sends a new DEK to the data server, which encrypts the future messages with the new key. The key server also sends to the distribution server its (updated) network access control policy or access control list indicating which nodes are currently authorized to join the group. The data server is mainly engaged in processing the data to be distributed, e.g., computing encryptions and digital signatures. It transmits the prepared data to the distribution server for distribution.

B. Design Goal

Security requirements of overlay multicast are similar to those of other networks. Some of the general security properties are *authentication*, *confidentiality*, *network access control*, *availability*, *anonymity*, and *fairness*.

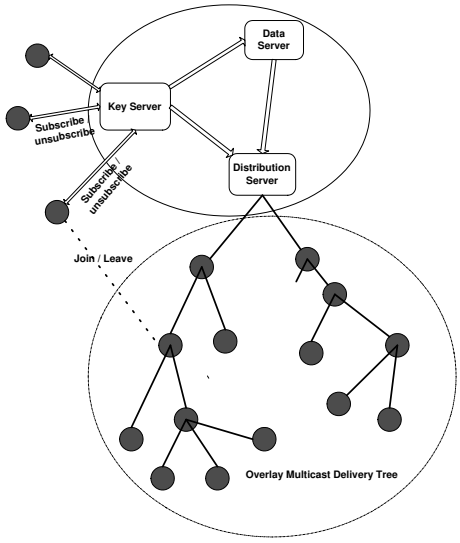

Fig. 1. A system model

In this paper, however, we focus on two of these security issues in the context of overlay multicast. First, we want to provide data confidentiality and network access control. Data confidentiality ensures that only authorized nodes can understand the multicast data. It must be provided because an unauthorized user may attempt to receive multicast data by eavesdropping on the communication links of authorized nodes or even of Internet routers. Network access control is also critical because it ensures that only authorized nodes can join the overlay multicast tree; otherwise, the resources of a legitimate node are consumed for forwarding data to unauthorized nodes. Second, we want to provide a DoS-resilient key distribution scheme that delivers keys to existing member nodes with a high probability even if some selectively compromised nodes drop the keys they are supposed to forward.

We assume the existence of an appropriate multicast source authentication scheme (e.g., digital signature or other schemes [25], [26], [31], [24], [34] that amortize the cost of a digital signature over multiple packets) by which the data server and the key server can authenticate their data or keys to member nodes. We leave the other topics such as incentive for forwarding, service fairness, and insider attacks to our future study.

IV. A Certificate Revocation Based Group Rekeying Scheme (CRBR)

An important, while challenging, issue for providing multicast data confidentiality is group key management. To enforce both backward confidentiality (i.e., a new user should not be able to decipher the data distributed before its subscription) and forward confidentiality (i.e., a revoked user should not be able to decipher the future data) [33], it is required to distribute a new group data encryption key (DEK) to all authorized members in a *secure*, *reliable*, and *timely* fashion when group membership changes. This is referred to as group rekeying.

Unicast-based group rekeying, in which the key server sends a new DEK to every individual node, has the communication

complexity of $O(N)$ keys. Recently proposed group rekeying schemes [7], [23], [27], [32], [33] use logical key trees to reduce the complexity of a group rekeying operation from $O(N)$ to $O(logN)$. Further, it has been proposed that groups be re-keyed periodically instead of on the basis of every membership change [28], [37]. Periodic or batched rekeying can reduce both the processing and communication overhead at the key server, and improve the scalability and performance of key management protocols. Note that in all these schemes, the key server includes keys for *all* the member nodes when distributing its rekeying message, and every member receives the entire message although it is only interested in a small fraction of the content.

Network access control, which is another critical security service, was studied *independently* with group key management in IP multicast. It is usually enforced by Internet edge routers [16]. Specifically, each router maintains inclusion or exclusion access control lists (ACLs) for all supported multicast groups, and every member presents its authorization certificate or token to its edge router to join a group. Network access control is hard to implement in IP multicast because routers are required to authenticate packets, to establish trust relationship with individual group controllers, and to keep their ACLs up-to-date.

A. Scheme Overview

We exploit the property of an overlay network that nodes are both hosts and routers when designing group rekeying and network access control schemes. In IP Multicast, all group members are end hosts, and they have no responsibility for forwarding keying materials to other group members. In contrast, for group communication in an overlay network, the nodes in the delivery tree also act as routers. As such, the key server only has to deliver a new DEK securely to a small number of nodes, which are its immediate children, and these children then forward the new DEK securely to their own child nodes. In this way, a group key is propagated to all the online member nodes in a hop-by-hop fashion. The amortized transmission cost per node is one key, independent of the group size.

For the above scheme to work, a basic requirement is the existence of a secure channel between every pair of neighboring nodes. We employ conventional public key techniques for establishing pairwise keys between two nodes. Symmetric-key techniques might also work, but have several limitations. For example, during node registration, the key server could assign to every node N pairwise keys, each of which is shared uniquely with another node. However, this scheme is not scalable with group size for distributing and storing these keys, and it is also not flexible for adding new members. Although the technique of probabilistic key pre-deployment [14], [19] can be used to reduce the storage requirement, it only provides probabilistic or threshold security in the sense that a certain number of colluding members can greatly jeopardize the secure links shared between other members in the system.

The use of public key techniques can additionally provide network access control because public key techniques such as digital signatures support strong source authentication. In overlay multicast, because nodes are both routers and hosts, network access control will be achieved as long as every node authenticates every other node that contacts it for joining the network.

In our system model, it is very natural that data access control (through encryption) and network access control (through authentication) be integrated. A node should have both privileges if it is authorized, and it should not have either of them if it is revoked. This motivates us to update group keys and invalidate the public keys (or certificates) of revoked nodes simultaneously. Moreover, since periodic group rekeying is much more scalable than individual rekeying [28], [37] and certificate revocation information is also distributed periodically [8], group rekeying and the distribution of certificate revocation information can be performed with the same time interval. An appropriate rekeying interval is application dependent and requires a trade-off between security and performance. The selection of rekeying interval is out of the scope of this paper.

B. Scheme Specifications

This subsection describes our scheme. We first show the major steps, followed by more details.

- **Node Registration:** The key server issues every member a public-key certificate when the member subscribes.
- **Security Update Generation:** The key server generates a new certificate revocation list (CRL) and a new DEK K_g for every group rekeying. It further computes a digital signature over the CRL, K_g, and a timestamp. Denote SU as a *security update* that includes the CRL, the timestamp, and the above digital signature (note that K_g is not included in SU). The key server sends SU and K_g to the distribution server.
- **Security Update Distribution:** For the ease of presentation, here we use a traditional, non-DOS-resilient scheme for the distribution of security updates. This is referred to as base scheme. We will present a DOS-resilient scheme in Section V. In the base scheme, the distribution server forwards SU and K_g to each of its child nodes. SU is sent in cleartext, whereas K_g is encrypted with a pairwise key shared between two nodes in every link. A node establishes a pairwise key with another node and then propagates K_g to it only if the CRL indicates that node is still a legitimate member. Also note that every node can verify the authenticity of the received group key K_g by verifying the signature. After successfully verifying the message, these child nodes forward the message to their own child nodes. Recursively, the security update and K_g are propagated to all on-line nodes in a hop-by-hop fashion.
- **Local Recovery:** A node that has missed one or several security update information because it was off-line can authenticate itself to any one of the online nodes to obtain the up-to-date security update and the group key when it joins the network, because that node knows if the joining node is legitimate or not based on the CRL it possesses.

1) Certificate Management: The key server issues every node a unique public-key certificate if the node is authorized. The node is also given the public key of the key server, which allows the node to verify the certificates (hence their public keys) presented by other nodes.

The key issue in using digital certificate is certificate management. In general, there are mainly two challenges. The first challenge is for a node to verify a received certificate in an efficient and timely fashion in the presence of complex CA hierarchy. In our applications, fortunately, there is no CA hierarchy since every user receives a certificate from the same key server. The second challenge is certificate revocation. On one hand, it is more efficient and economic for a CA to issue certificates with long validity periods, because there is considerable computational and communication overhead (e.g., computing a digital signature) involved in issuing a certificate. On the other hand, it is more likely that a certificate with a long validity period needs to be revoked explicitly before it is expired, because in some applications users may unsubscribe from the program at any time. To address this problem, a conventional approach is for the key server to issue a certificate revocation list (CRL) that contains the information of all revoked certificates. The issue is: how can the key server make the revocation information available to other nodes in an efficient and timely manner. There are mainly two techniques: pull or push. In a pull-based scheme, a node contacts a dedicated directory to verify a received certificate from another node. Since every node needs to authenticate itself to multiple nodes, including its parent, children, and multiple other nodes it has to contact during its join process, the pull-based scheme makes the directory a performance bottleneck.

We adopt a push-based approach in which the key server distributes its CRLs periodically. To minimize the communication overhead, we employ the following techniques.

- The key server assigns a unique integer to every node as the identifier of the certificate of the node. The integer starts from '0' and is incremented by one for a new node. Note that the key server does not assign the ids of revoked nodes to any new nodes. Also, the certificate of a user does not contain any personal information about the user. Instead, the key server has a database that records the personal information of a user and its certificate.
- The CRL is a bit string of size Z, where Z is the number of nodes that have so far subscribed to the system. Every certificate is mapped to a bit in the bit string and the id of the certificate is the index of the bit in the bit string. A bit value of '0' indicates that a corresponding certificate is invalid and '1' indicates valid. When a node is revoked from the system, its corresponding bit in the bit string is set '0'.

2) Node Joining: When a node u joins the multicast delivery tree after its subscription process, it follows the existing overlay multicast routing protocol, except that it authenticates to all the nodes it contacts with and also verifies any messages from those nodes. For example, nodes u and its parent node v can authenticate to each other and establish a pairwise key K_{uv} based on their certificates and their public/private keys. Node v also checks if the bit indexed by node u's id in its CRL is '1'. Then node v sends node u the current K_g encrypted by K_{uv} and SU. Node u can verify the authenticity of K_g based on SU. Note that a pairwise key is not merely for delivering K_g. In all the overlay multicast routing protocols [3], [12], [11], [18], [38], two neighboring nodes in the multicast tree exchange KEEPALIVE messages periodically. They can use their pairwise key for authenticating these KEEPALIVE messages to each other.

C. Security Analysis

In our scheme, no unauthorized nodes can get the group key K_g because a member node only forwards K_g to other member nodes. Nor can a compromised node inject a false group key into the network because the group key is signed by the key server. An unauthorized node cannot join the multicast tree either. However, a compromised node might share its private key with other unauthorized nodes so that those nodes can join the network. This attack could be detected if the key server has the global knowledge of the online nodes in the multicast tree. Another approach, which works for users with static IP addresses, is to include the IP address of a node in its certificate. To receive data, a node cannot lie about its IP address, thus this approach thwarts those unauthorized nodes from using the certificate of an authorized user.

Our scheme also provides weak anonymity in the sense that the certificate of a user only has an integer field to uniquely identify the user. A node cannot figure out the identities of other users in the system; however, it can know the IP addresses of other nodes it is communicating with.

D. Performance Evaluation

This subsection compares the performance of our scheme CRBR with two well-known group rekeying schemes: LKH [32] and SDR [23], supposing they are employed in overlay multicast. Unlike CRBR, neither of LKH and SDR provides network access control. Nevertheless, in the comparison we consider the cost of distributing the CRL for CRBR but not for LKH or SDR, thus biasing the comparison in favor of the latter two schemes.

The metric of interest is key server bandwidth overhead. Because security updates are propagated in the entire multicast tree, the more the key server distributes, the more network and node resources are consumed. Hence, key server bandwidth overhead is also an indication of the total network bandwidth overhead. We do not consider computational complexity of nodes for group rekeying. In CRBR, two neighboring nodes establishes a pairwise key only once, and the pairwise key is then used for encrypting group keys or authenticating KEEPALIVE messages. The cost of establishing a pairwise key is therefore amortized to multiple messages. Moreover, a small number of public-key operations such as public-key decryptions or digital signatures are not a concern for nodes. In a PC with a Celeron 450 MHz processor, a 1024-bit RSA digital signature takes only 27 milliseconds [10].

Two scenarios are studied. The first scenario considers the bandwidth overhead of the key server for multicasting keys to

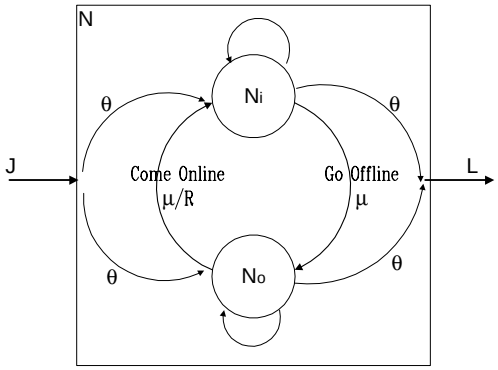

Fig. 2. An analytical model

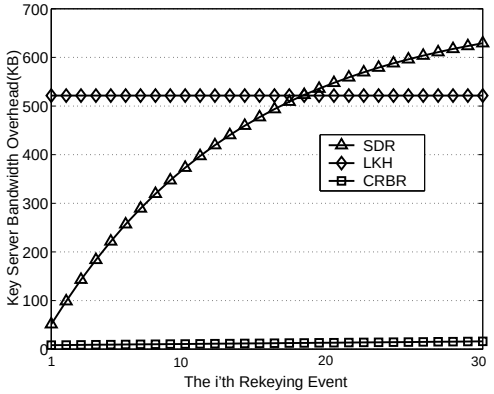

Fig. 3. Rekeying cost of LKH, SDR, and CRBR

online nodes. The second scenario considers the bandwidth overhead of the key server for unicasting the current keys to individual nodes that have missed one or several previous group rekeying operations because they were offline. To study the performance of these schemes quantitatively, below we first present an analytical model for node presence dynamics.

1) The Analysis of Node Presence Dynamics: A member node can be in either of two statuses: presence (online) or absence (offline), and it can switch its status between these two statuses until its membership duration is expired and then revoked from the group. We use the term "presence duration" and "absence duration" to denote a continuous time period a node stays in a group and stays outside a group, respectively. Previous study [1] based on multiple sessions in MBone showed that presence durations in a multicast session follow either an exponential distribution or a Zipf distribution. For simplicity, in this study we assume that the durations of node membership follow an exponential distribution with mean $1/\theta$. We also assume that presence durations of nodes follow another exponential distribution with mean $1/\mu$, and further assume that the absence durations of nodes are exponentially distributed with mean R/μ. Thus R is the ratio of the average time for which a node is absent to the average time for which it is present.

Figure 2 depicts our analytic model. Let the group rekeying interval be T. When the system is in its steady status, during T the number of new subscribers J is equal to the number of revoked subscribers L. Based on queueing theory, the revocation rate of the system is $N \cdot \theta$ where N is the population of the system. Thus, $L = N \cdot \theta \cdot T$.

In this model, the rate of an online user going offline is μ and the rate for an offline user coming online is μ/R. Let N_i and N_o be the populations of online and offline users just after a rekeying operation, respectively, then $N_i + N_o = N$. Denote S as the number of nodes switching from offline status to online status in T. In the steady status, S is also the number of nodes switching from online to offline in T. For periodic batched rekeying, both node additions and revocations are processed at the end of a group rekeying. Thus, $S = N_i \cdot \mu \cdot T = N_o \cdot \mu/R \cdot T$. We have $N_i = \frac{N}{R+1}$ and

$$S = \frac{N \cdot \mu \cdot T}{R+1}. \quad (1)$$

We will use S in our evaluation study shortly.

2) The Impact of Presence Dynamics on Group Rekeying:

Scenario I: Multicast Cost For LKH, we adopt the analytic result in [37] to compute the communication cost of the LKH-based group rekeying scheme. For SDR, the upper bound of group rekeying cost is $2r$, where r is the accumulated number of nodes that have been revoked from the system so far. We use simulations to show its average cost in different rekeying periods.

The setting for the comparison is as follows. We assume that in the steady status, $N = 65536$. Let the average membership duration be $1/\theta = 30$ days and the group rekeying interval be $T = 1$ day. Let the size of a key be 20 bytes (128-bit AES encryption [2], with a key version field and encoding overhead). Figure 3 depicts the bandwidth overhead of LKH, SDR, and CRBR in the first thirty rekeying events. We can observe that LKH has the same bandwidth overhead during different rekeying operations, whereas in SDR the bandwidth overhead starts at a small value and eventually exceeds that in LKH. The bandwidth overhead in CRBR increases slightly with time, but it is still far smaller than that in other two schemes.

Note that although in this simulation setting the bandwidth overhead seems not a big concern in any one of these three schemes, the bandwidth saving of our scheme over other two schemes is very meaningful for a very large group. For example, when N reaches one million, under the same setting every rekeying cost in LKH is 8.3 MB, in SDR it becomes several megabytes when r reaches hundreds of thousands, whereas in CRBR it is upper bounded by 128 KB. Because of its small size, the rekeying message in CRBR can be distributed using our DoS-resilient k-RIP scheme introduced in Section V.

Scenario II: Unicast Cost When the key server distributes its security update, none of the N_o offline nodes receive it. As we showed earlier, S of these nodes come back online during T. The issue is that they cannot decrypt the data without the current group key. One solution is letting these nodes wait until the next group rekeying operation. However, this does not address the problem completely. When a stateful rekeying

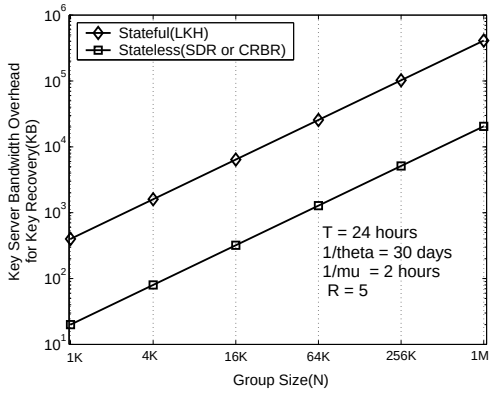

Fig. 4. The bandwidth overhead of key server for unicasting keys to nodes.

protocol such as LKH or OFT is employed, a node cannot decrypt the next rekeying message to get the next group key either. A stateless protocol does allow a node to decode the next group key from the next rekeying message independently. However, a subscriber has to wait for an average time interval of $T/2$ for the next rekeying interval. Therefore, to update their keys, these S nodes have to ask for retransmission. Generally, a node can keep the last security update, but it does not bear the responsibility for keeping all the previous security updates. Therefore, when LKH is employed, these S nodes would need to ask the key server for retransmission, whereas in SDR and CRBR, these nodes can get retransmission from other nodes.

Next we show the overall retransmission cost in these schemes. In SDR or CRBR, only the current group key (no KEKs) is retransmitted to a requesting node. Thus the overall bandwidth overhead is S keys if we do not count other packet overhead. In LKH, a node needs to receive the current group key and some of its KEKs that have been updated. We assume that in LKH, a node that comes online needs to receive on average $h/2$ keys [1], where h is the height of the key tree maintained by the key server. Therefore, during T, the key server needs to retransmit $\frac{Sh}{2}$ keys.

Figure 4 plots the bandwidth overhead for unicast-based key updating in a stateful protocol (LKH) and a stateless protocol (SDR or CRBR). Again we set $T = 1$ day, presence interval $1/\mu = 2$ hours, and $R = 5$. We can observe that the key server bandwidth overhead is nontrivial and it increases linearly with group size. For example, for a group of size of $65,536$, in LKH the key server has to transmit 25.6 MB to help nodes update keys. In SDR or CRBR, the cost is 1.28 MB. When N reaches one million, the cost in LKH becomes greater than 400 MB.

Overall, the analysis of these two scenarios shows that CRBR outperforms LKH and SDR for the applications under consideration. Moreover, our simulation (although not shown) indicates that in many cases we can greatly reduce the size of the CRL by compressing it using a compression program, e.g.,"zip", before performing digital signing. Our id assignment rule renders many of '1's (also '0's) contiguous. This allows us to use a much larger bitmap (e.g., millions or billions users), the size of which after compression has almost no difference with that of a compressed bitmap of smaller size (e.g., less than one million). An advantage of using a large bitmap is to accommodate a potentially very large number of users.

[1]The exact number depends on the group characteristics N and L. Our analysis shows that it is about $\frac{3h}{4}$ keys in our earlier setting if a node has only missed one group rekeying event.

V. A DoS-Resilient Key Distribution Scheme (k-RIP)

This section describes our DoS-Resilient key distribution scheme called k-RIP. The scheme can also be used for distribution of other small-size but critical information (e.g., new virus and worm signatures, CRLs) in overlay multicast group.

In overlay multicast, messages are normally injected into the network from the distribution server, i.e., the root node, and are then forwarded hop-by-hop to all the other nodes in the tree (in Section IV-B we used this as the base scheme for key distribution). If a malicious node in the tree intentionally discards the message it receives from its parent node, its downstream nodes will not receive the message. This attack is specially severe when the malicious node is very close to the root. We note that this attack is also effective to non-tree based delivery infrastructure. Schemes [31], [35], [36] based on more complex graphs are more resilient to the attack in general, but they are still subject to selective attacks in which an attacker selectively compromises several nodes close to the injection point. The injection point is fixed and well-known, thus naturally becoming a point of interest for denial of service attacks. To launch DoS attacks, an attacker could also manage to compromise the nodes close to the injection point.

Note that we cannot solely rely on detection and retransmission mechanisms to address this attack. If every node that detects message losses asks the key server for retransmission, the key server will become the performance bottleneck. Therefore, it is very important that the majority of the member nodes could receive messages even in the presence of DoS attacks.

A. Scheme Overview

To address the above attacks, we propose that in addition to propagating its message through the root node, the key server also randomly picks k nodes (not including the root) in the tree and sends its message to these k nodes. All these k nodes propagate the message towards their children (if any) as well as their parents if their children or parents have not received the message yet. Thus, if a small number of nodes do not forward the message, other nodes might still be able to get it from their children or parents with high probability.

In this scheme, we can simply use sequence numbers to suppress duplicated messages, thus every node only receives one copy of the message and forwards the message to another node at most once[2]. Moreover, this scheme has the additional benefit of reducing the overall latency for all online nodes to receive the message. On the other hand, this scheme incurs the bandwidth overhead for the key server to transmit k additional

[2]A node may first send to a recipient node a probing packet including only the sequence number of the message, and then decides whether or not to send the entire message based on the feedback from the recipient node.

copies of the message. However, for small-size messages (e.g., tens or hundreds of kilobytes) and a small k (e.g., ≤ 20), in practice this transmission overhead should not be a big concern for the key server.

B. Node Selection

The very first question is which k nodes to select? To answer this question, we need to consider two factors: latency and DoS-resilience. Ideally, we should select k nodes such that the overall latency is minimized and the number of nodes that can receive messages is maximized. It is relatively easy to select k nodes to minimize the overall latency if we have the complete knowledge of the tree topology and the delay in each link of the tree. A brute-force–based algorithm works by calculating the overall latency in every possible combination of k nodes out of the total N_a nodes in the tree and then finding out the set producing the minimum latency. The computational complexity is $\binom{N_a}{k}$ calculations of the overall latency and there is room for optimization. We can also calculate the number of nodes that can receive the CRL for every combination if we know which nodes in the tree are malicious.

In practice, it is hard to achieve the above goal because the key server might not have the precise knowledge of the tree topology due to the presence dynamics of the member nodes. The key server might know which nodes have joined the tree from a rendezvous point (RP) in many routing algorithms [3], [38] because a joining node contacts a RP for information assisting the node to find a position in the tree. However, for scalability the RP does not keep track of the position of a specific node in the tree. Moreover, the RP does not know either if a node is online or offline because a node might leave the group at any time and it does not notify the RP its leaving. Thus, for our scheme to work, a practical issue is to determine which nodes are online.

A Heuristic Selection Algorithm A simple solution works as follows. The key server randomly selects its member nodes to connect to. If a member node is unreachable, it picks another one. The key server repeats this process until it discovers k online nodes. One problem with this scheme is that the key server might not know the IP addresses of its member nodes because nodes might have dynamic IP addresses. This problem can be addressed by letting the RP record the IP addresses of the nodes that have recently contacted it. Because nodes normally do not change their IP addresses during a session, the key server can use these IP addresses directly.

Using the same group characteristics as that used in Section IV-D, we know that for a system that has the registered population of N and the average network size of $N_a = \frac{N}{R+1}$, the key server needs to try an average number of $\frac{kN}{N_a} = k(R+1)$ times to find k online nodes. This shows that the efficiency of this algorithm relies on the node presence dynamics. For a small R, this selection algorithm should work fine. When R is large, we may exploit the following heuristics to increase the hit ratio. The idea is that the key server could make a good guess of online members based on the joining times of the members. Again, we assume that presence durations in a multicast session approximately follow an exponential distribution [1]. Assume that the mean of presence durations is $1/\theta$, which can be calculated if every member node records its every presence period and reports its mean presence time to the RP when joining the tree.

The probability $p_i(t)$ that a member node i is still online t time after it joins is $p_i(t) = e^{-\theta \cdot t}$. $p_i(t)$ decreases with t, indicating that the nodes joining more recently are more likely to be online than those joining earlier. Thus, the RP simply tells the key server the ids of m distinct nodes that joined the tree most recently, such that $\sum_{i=1}^{m} p_i(t_i) \geq \varphi k$, where t_i ($1 \leq i \leq m$) is the time difference between the current time and the joining time of that node. Here $\varphi \geq 1$ is a parameter reflecting the probability that the key server finds k online nodes from m candidates, and it is variable and should be determined by the presence dynamics of an actual application.

We note that there is a potential attack against this selection algorithm if the message (e.g., the CRL in CRBR) is distributed periodically, because multiple malicious nodes may join the tree just before the distribution time point. Based on our selection algorithm, the RP will likely report their ids to the key server. Thus, these malicious nodes are selected as injection points, reducing the effectiveness of our scheme. To mitigate this attack, it is important that the key server randomly picks nodes from the m candidates for presence test, not preferring the nodes that joined more recently.

C. Evaluation of Effectiveness

This subsection reports the effectiveness of our k-RIP scheme in reducing propagation latency and increasing DoS-resilience.

1) Simulation Setting: We first generate a random graph of 10,000 nodes and then construct a tree out of the graph based on the joining algorithm that is also used in [4], [18], [38]. Specifically, every joining node searching from the root downwards along the tree for the (possible) nearest node as its parent, thus geometrically adjacent nodes become neighbors in the tree. The link delay between any two nodes is randomly selected from a uniform distribution between 10 and 200 ms, and the outdegree of a node is a random number between 1 and 5. Our simulation programs were written using the Csim simulation library [13]. We use the method of independent replications for our simulations and all our results have 95% confidence intervals that are within 5% of the reported values.

2) Reducing Latency: We first evaluate the effectiveness of k-RIP in reducing the propagation delay when no nodes are compromised. Figure 5 plots the latency histogram when k nodes are randomly selected from 10,000 nodes. In the base scheme, there is only one injection point, which is the root. Note that for simplicity we calculate the message propagation latency in each link as triple of the link delay to mimic TCP-like unicast schemes; hence, the results should only be interpreted as relative performance. We can see that with the increase of k, more and more nodes can receive a distributed message in shorter time intervals. We have also calculated the average latency in each case (not shown in this figure). When $k = 20$, the average latency is about 15% less than

that in the base scheme. Note that the reduced latency is not very significant over the base scheme because the nodes are randomly selected. Due to the tree structure, most of the nodes are selected from leaves or close to leaves.

We investigate the effectiveness of the heuristic algorithm by randomly choosing k ($k \leq 20$) injection points from $m = 50$ most recently joined nodes, which shows that it is not as effective as choosing k nodes from all N_a nodes. When $k = 20$, the average latency is only about 11% less than in the base scheme. This is because the recently joined nodes are mostly leaf nodes in the tree.

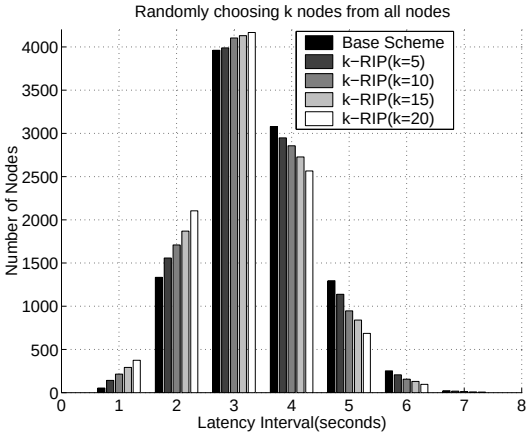

Fig. 5. A histogram showing the number of nodes receiving the distributed message in each time interval.

3) Increasing DoS-Resilience: We first show the analytical model, followed by simulations.

a) Performance Analysis: Figure 6 depicts an example multicast tree that has degree of $d = 2$ and group size of N_a nodes (excluding the root node that is the distribution server). The solid nodes are good nodes and the empty one is a compromised node that drops messages going through it. Let h be the height (in hops) of the compromised node from the root and s be the number of nodes in the subtree rooted at the compromised node. Then we have $s = \frac{N_a}{d^h} - \frac{1-(\frac{1}{d})^{h-1}}{d-1}$.

Let z be the number of good nodes that can receive messages. Let base scheme be the one in which only a single copy of a message is injected via the root node. It is easy to see that in base scheme $z = N_a - s$. In our k-RIP scheme, besides the root node, k randomly picked nodes also inject the message into the tree simultaneously. If a good node in the subtree rooted at a child node of the compromised node is selected , all the nodes in this subtree will receive the message. If one good node is selected from each subtree rooted at each child node of the compromised nodes, all the good nodes in the network will receive the message. More generally, denote $p(i)$ as the probability that at least one node is selected from each of i ($0 \leq i \leq d$) subtrees rooted at i child nodes (but no nodes are selected from the other $(d - i)$ subtrees) of the compromised node. Further denote $x = (s - 1)/d$ and $y = N_a - s + 1$, as shown in Figure. 6. Basic probability and combinatorics

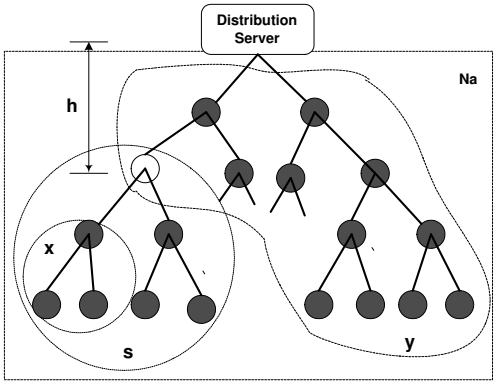

Fig. 6. An example tree of size N_a excluding the root. The empty node is a compromised node whose tree size is s.

arguments can be used to derive $p(i)$:

$$p(i) = \begin{cases} \binom{y}{k} / \binom{N_a}{k} & (i = 0) \\ \binom{y+ix}{k} / \binom{N_a}{k} - \sum_{j=0}^{i-1} \binom{i}{j} p(j) & (0 < i \leq d) \end{cases}$$

The expected value of z is:

$$E[z] = \sum_{i=0}^{d} \binom{d}{i} p(i)(y - 1 + ix). \quad (2)$$

This analytic result has been validated by simulations.

b) A Single Compromised Node: Figure 7 illustrates the effectiveness of our k-RIP scheme compared to the base scheme, based on eqn. 2. We observe that in the base scheme, a compromised node with the height $h = 1$ could prevent half of existing nodes from receiving the message, whereas our scheme allows a much larger fraction of nodes to receive the message. For example, when $k = 3$, about 80% nodes can receive the message. More nodes get the message when k increases. For example, when $k = 10$, about 97% nodes could receive it. We also observe that if a compromised node is far from the root of the tree, it will not be able to affect many nodes even in the base scheme. This is because less nodes are in the tree rooted at the compromised node. In addition, the figure indicates that in the base scheme, to cause denial of service to more nodes, an attacker should manage to become as close to the root as possible. However, when our scheme is deployed, that is not necessarily the best strategy for the attacker. For example, when $k = 10$, becoming a node with $h = 2$ or $h = 3$ gives the attacker a little more advantage than becoming a node with $h = 1$.

Next we study the effectiveness of our heuristic selection algorithm that selects k nodes from mostly recently joined m nodes (denoted as k-RIP-h, dashed lines in Figure 8), compared with the base scheme in which we randomly select k nodes from N_a nodes (denoted as k-RIP, solid lines). We set the outdegree of a node $d = 3$ and the size of the candidate set $m = 50$. Figure 8 indicates that except when $h = 1$, the effectiveness of k-RIP is only slightly affected when choosing

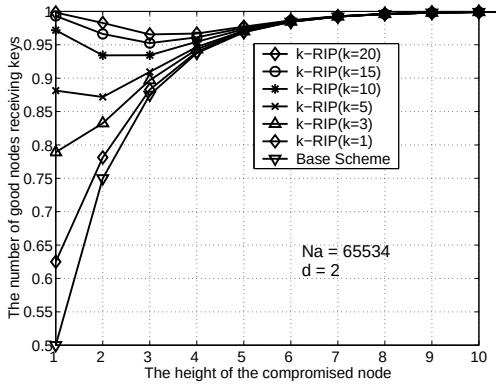

Fig. 7. The fraction of nodes that can receive messages as a function of the location of one single compromised node in the multicast tree and k.

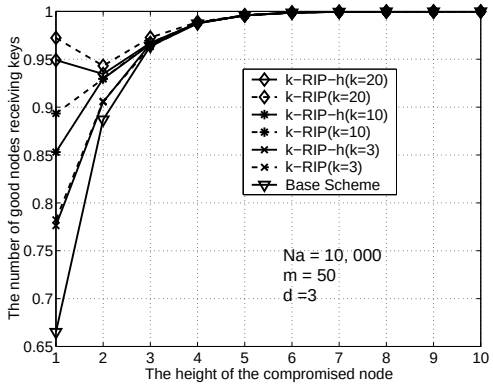

Fig. 8. The effectiveness of our heuristic selection algorithm (through simulation, 95% C.I. is within 5% interval of a reported value.)

k ($k \leq 20$) nodes from 50 most recently joined nodes instead of choosing from all $N_a = 10,000$ nodes.

c) Multiple Compromised Nodes: Figure 9 shows the effectiveness of our scheme when all the d^h nodes in height h are compromised, instead of a single compromised node in the previous case. When $h = 1, 2, 3$, in the base scheme almost no good nodes receive keys. The effectiveness of our scheme decreases with h. For example, in the case of $h = 3$ (i.e., all the 8 nodes in level 3 are compromised), when $k = 20$ only about 72% nodes receive keys. Note that if the key server has the global knowledge of the multicast tree, higher robustness can be achieved with smaller communication cost. For example, in the case of $h = 3$, there are totally 16 subtrees of these 8 compromised nodes. By selecting one node from each of the 16 subtrees, all good nodes can receive keys.

VI. Conclusions and Future Work

We have presented a bandwidth efficient scheme that integrates network access control and group key management. Performance analysis and simulation study show that our scheme incurs much smaller communication overhead than two other well-known schemes. Our idea of hop-by-hop secure forwarding of group keys may also be applied to multi-hop ad hoc and sensor networks, but we will need to study the new challenges that arise due to resource constraints of ad hoc and sensor nodes.

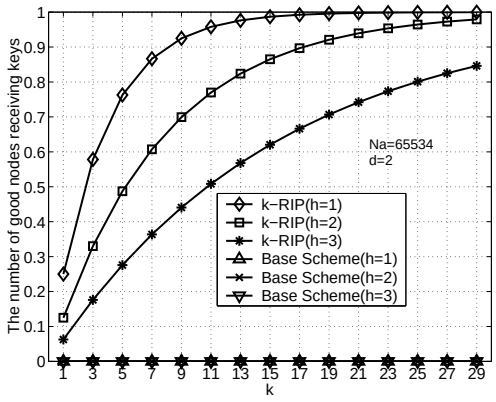

Fig. 9. The fraction of nodes that can receive messages when all the nodes in the height h are compromised.

We also proposed a DoS-resilient information distribution scheme that delivers small-size but critical messages (e.g., keys) to a large fraction of nodes with high probability even if an attacker can *selectively* compromise nodes in the multicast data delivery hierarchy. The scheme has only considered to distribute one message and use a fixed k. If multiple messages are to be distributed, we may randomly sample a fraction of nodes to obtain their receiving statuses of the previous messages and then decides an appropriate k for the next message. We will study efficient and effective sampling schemes for this purpose.

References

[1] K. Almeroth and M. Ammar, Multicast Group Behavior in the Internet's Multicast Backbone (Mbone), IEEE Communications, June 1997.

[2] Advanced Encryption Standard (AES). http://csrc.nist.gov/CryptoToolkit/aes/

[3] S. Banerjee, B. Bhattacharjee, C. Kommareddy. Scalable Application Layer Multicast. In Proc. of ACM Sigcomm, 2002.

[4] S. Banerjee, C. Kommareddy, K. Kar, B. Bhattacharjee, and S. Khuller. Construction of an Efficient Overlay Multicast Infrastructure for Real-time Applications In Proc. of IEEE Infocom 2003, San Francisco, April 2003.

[5] S. Banerjee, S. Lee, B. Bhattacharjee, A. Srinivasan. Resilient Multicast Using Overlays. In Proc. of ACM Sigmetrics 2003.

[6] G. Badishi, I. Keidar, and A. Sasson. Exposing and Eliminating Vulnerabilities to Denial of Service Attacks in Secure Gossip-Based Multicast. In Proc. of Dependable Systems and Networks (DSN), 2004.

[7] D. Balenson, D. McGrew, and A. Sherman. Key Management for Large Dynamic Groups: One-Way Function Trees and Amortized Initialization. IETF Internet draft (work in progress), August 2000.

[8] CCITT Recommendation X.509: The Directory-Authentication Framework. 1988.

[9] R. Canetti, J. Garay, G. Itkis, D. Micciancio, M. Naor, B. Pinkas. Multicast Security: A Taxonomy and Some Efficient Constructions. In Proc. of IEEE INFOCOM'99, March 1999

[10] http://packetstormsecurity.nl/crypt/LIBS/cryptolib/ benchmarks.html.

[11] Y. Chu, S. Rao, S. Seshan, and H. Zhang. Enabling conferencing applications on the internet using an overlay multicast architecture. In Proc. of ACM SIGCOMM 2001.

[12] Y. Chu, S. Rao, and H. Zhang. A case for endsystem multicast. In Proc. of ACM Sigmetrics '00, 2000.

[13] http://www.mesquite.com/.

[14] W. Du, J. Deng, Y. Han, and P. Varshney. A Pairwise Key Predistribution Scheme for Wireless Sensor Networks. In Proc. of 10th ACM Conference on Computer and Communications Security (CCS), Washington DC, October 27-31, 2003.

[15] M. Handley. An examination of MBONE performance. Technical Report ISI/RR- 97-450, USC/ISI, 1997.
[16] H. He, T. Hardjono, and B. Cain. Simple Multicast Receiver Access Control. draft-irtf-gsec-smrac-00.txt, Nov. 2001.
[17] P. Judge and M. Ammar. GOTHIC: A Group Access Control Architecture for Secure Multicast and Anycast. In Proc. of IEEE Infocom 2002.
[18] J. Jannotti, D. Gifford, K. Johnson, M. Kaashoek, and J. O'Toole. Overcast: Reliable, Multicasting with an Overlay Network. In Proc. of 4th USENIX OSDI Symposium, 2000.
[19] D. Liu and P. Ning. Establishing Pairwise Keys in Distributed Sensor Networks. In Proc. of the 10th ACM Conference on Computer and Communications Security (CCS '03), Washington D.C., October, 2003.
[20] D. Liu, P. Ning, and K. Sun. Efficient Self-Healing Group Key Distribution with Revocation Capability. In Proc. of the 10th ACM CCS, 2003.
[21] J. Li, P. Reiher, and G. Popek. Resilient Self-Organizing Overlay Networks for Security Update Delivery. IEEE Journal on Selected Areas in Communications, Vol.22., January 2004.
[22] L. Mathy, N. Blundell, V. Roca, and A. Elsayed. Impact of Simple Cheating in Application-Level Multicast. In Proc. of IEEE Infocom 2004.
[23] D. Naor, M. Naor, and J. Lotspiech. Revocation and Tracing Schemes for Stateless Receivers. In Advances in Cryptology - CRYPTO 2001. Springer-Verlag Inc. LNCS 2139, 2001, 41-62.
[24] P. Rohatgi. A compact and fast hybrid signature scheme for multicast packet authentication. In Proc. of the 6th ACM CCS'99, pp. 93-100, 1999.
[25] J. Park, E. Chong, and H. Siegel, Efficient Multicast Packet Authentication Using Signature Amortization. In Proc. of IEEE Symposium on Security and Privacy, 2002.
[26] A. Perrig, R. Canetti, D. Song, and J. Tygar. Efficient and secuire source authentication for multicast. In Network and Distributed System Security Symposium, NDSS'01, Feb. 2001.
[27] A. Perrig, D. Song, D. Tygar. ELK, a new protocol for efficient large-group key distribution. In Proc. of IEEE Symp. on Security and Privacy 2001, Oakland CA, May 2001.
[28] S. Setia, S. Koussih, S. Jajodia, and E. Harder. Kronos: A Scalable Group Re-Keying Approach for Secure Multicast. In Proc. of the IEEE Symposium on Security and Privacy, Oakland CA, May 2000.
[29] J. Staddon, S. Miner, M. Franklin, D. Balfanz, M. Malkin and D. Dean. Self-Healing Key Distribution with Revocation. In Proc. of the IEEE Symposium on Security and Privacy, oakland, CA, May 2002.
[30] S. Setia, S. Zhu and S. Jajodia. A Comparative Performance Analysis of Reliable Group Rekey Transport Protocols for Secure Multicast. In Performance Evaluation 49(1/4): 21-41 (2002), special issue Proceedings of Performance 2002, Rome, Italy, Sept 2002.
[31] D. Song, D. Zuckerman, and J. Tygar. Expander Graphs for Digital Stream Authentication and Robust Overlay Networks. In Proc. of IEEE Symp. on Security & Privacy, 2002.
[32] C. Wong, M. Gouda, S. Lam. Secure Group Communication Using Key Graphs. In Proc. of SIGCOMM 1998, Vancouver, British Columbia, 68-79.
[33] D. Wallner, E. Harder and R. Agee. Key Management for Multicast: Issues and Architecture. Internet Draft, draft-wallner-key-arch-01.txt, September 1998.
[34] C. Wong and S. Lam. Digital signatures for flows and multicasts. IEEE/ACM Transactions on Networking, vol. 7, pp. 502-513, 1999.
[35] R. Wright, P. Lincoln, and J. Millen. Efficient Fault-Tolerant Certificate Revocation. In Proc. of ACM CCS 2000.
[36] H. Yang, H. Luo, Y. Yang, S. Lu, and L. Zhang HOURS: Achieving DoS Resilence in an Open Service Hierarchy. In Proc. of Dependable Systems and Networks (DSN), 2004.
[37] Y. Yang, X. Li, X. Zhang and S. Lam. Reliable group rekeying: Design and Performance Analysis. In Proc. of ACM SIGCOMM 2001.
[38] B. Zhang, S. Jamin, and L. Zhang. Host Multicast: a Framework for Delivering Multicast to End Users. In IEEE Infocom 2002.