Why Hierarchical Key Distribution is Appropriate for Multicast Networks

Chandana Gamage, Jussipekka Leiwo*, and Yuliang Zheng

Peninsula School of Computing and Information Technology
Monash University, McMahons Road, Frankston, VIC 3199, Australia
{chandag,skylark,yuliang}@pscit.monash.edu.au

Abstract. The design rationale for many key distribution schemes for multicast networks are based on heuristic arguments on efficiency, flexibility and scalability. In most instances the choice of key server placement in a multicast network architecture is based on intuitive cryptographic considerations. We use an analytical model of multicast group formation and network growth to look at the selection of a key distribution scheme from a network operation perspective. Thereafter, this model is used to validate the choice of hierarchical (hybrid) key distribution model as the most appropriate.
Keywords. Network security, Multicast networks, Key distribution architectures

1 Introduction

The phenomenal growth of wide area networks, in the form of ubiquitous *Internet*, have given rise to many new applications that are different from the typical one-to-one (unicast) communication model of standard network applications. Many of the new applications in information distribution and collaborative activities such as web-casting, shared white-boards, on-line auctions, etc., have a one-to-many (multicast [6, 7]) model of communications. There are two main reasons that motivate the use of multicast for highly distributed network applications:

1. The number of messages a sender needs to transmit is reduced. This is due to the fact, that a single multicast address represents a large number of individual receivers. This results in a lower processing load for the sender and also simplifies the application design.
2. The number of messages in-transit over the network is reduced. As the correct message delivery is handled by multicast-capable routers, which normally make redundant copies of a message only when transmitting on divergent network links, data meant for a group of receivers is transmitted as a single message for most part of the network. This in turn improves the overall network bandwidth utilization.

* Since Sept. 1999, author has been with Vrije Universiteit, Department of mathematics and computer science, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands, leiwo@cs.vu.nl

Therefore, multicast data transmission provides significant benefits to both the applications and the network infrastructure and consequently is an important network technology for emerging applications. The basic difference between broadcast networks and multicast networks is that in multicast, delivery is to a specifically targeted group. This group may be created based on many metrics such as affiliation to a certain institution, long-duration membership subscriptions, short-duration tickets, etc. Many of the group management functions such as join, leave or re-join that control membership of a multicast group require cryptographic techniques to ensure that integrity of the control process is not compromised by malicious users or intruders. Furthermore, the multicast application itself may require secure data transmission to and from members. As the communication model of multicasting is different from unicast communication, the attacks and threat models are also different for multicast networks and in fact more severe [2].

To provide secure group management services, standard security functions such as identification, authentication and message transmission with confidentiality and integrity are required. The basic support service for secure group management in multicast networks is session key distribution which incorporates the primary functions of member identification, authentication and session key transport. The key distribution schemes described in literature can be classified under three basic models of centralized, distributed or hierarchical as shown in figure 1. In the fully distributed scheme, although shown as a tree, a fixed root may not be physically present.

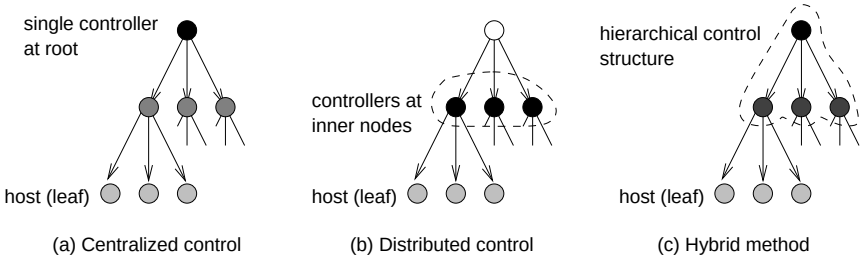

Fig. 1. Standard multicast group control methods. The fully distributed method shown in (b) requires horizontally structured coordination among participating controller nodes

Motivation The cryptographic research literature is replete with sophisticated key distribution architectures for multicasting based on wide ranging assumptions while the networking community have adopted only a handful of techniques in proposed or experimental secure multicasting schemes. The work presented

in this paper was motivated by the inadequate consideration given to network-centric issues when developing solutions that are grounded in cryptography.

Organization In section 2, we overview several secure multicast schemes to see if their key distribution scheme selection is based on network considerations or cryptographic issues (or a combination of both). In section 3, we develop analytical arguments from a network perspective to validate the choice of hierarchical key distribution as the preferred framework. We make concluding remarks in section 4.

2 Related Work

In general, control and routing tree structure selection (shared trees, shortest-path trees, etc.) and protocol algorithm design for multicasting is based on expected sparseness/denseness of multicast group, efficiency in terms of number of messages, low message propagation delay, ease of recovery from message loss and low overhead in group management. For *secure* multicasting, in which the main design aspect is the key distribution scheme, designers may opt to consider underlying multicast network characteristics or mainly use cryptographic metrics such as number of rounds required for key distribution, size of security control messages and key update/change techniques. Next we briefly review previous work from literature that have taken different approaches to implementing secure multicasting.

A design for a secure key distribution architecture is presented in [14] that is overlaid on the core-based tree (CBT) multicast routing protocol [3]. The justification for the hybrid control structure of [14] in which key distribution centers (KDC) are co-located with routers is based on the favorable characteristics of the multicast protocol rather than on multicast network structure itself. Among the main reasons given for the use of CBT framework for key distribution are the pre-existing scalability properties of the routing protocol, close relationship between grouping structure and router placement and the ability to combine processing workload for router setup and key distribution. Early work on key distribution schemes based closely on underlying multicast protocol structures appeared in [1, 2, 11].

Similarly, the *Iolus* secure multicasting framework [15] is based on a distributed tree of group security intermediaries (GSI) for subtrees and an overall group security controller (GSC) for coordination of GSIs. The collection of these group security agents constitutes a hybrid key distribution architecture. However, the framework is designed to operate over many different multicast protocols including CBT and protocol independent multicasting (PIM) [7]. The distributed registration and key distribution (DiRK) technique presented in [16] is another multicast protocol independent decentralized and distributed model that simply assumes a hybrid model is better suited for large scale multicast groupings. Similar proposals appear in [9]

In contrast, the SecureRing suite of group communication protocols [12] use multicasting and a fully distributed control structure to provide membership

management and message distribution under Byzantine errors but does not depend on any particular characteristics of the underlying multicast routing protocol for efficient or reliable operation. Their scheme uses cryptographic message digests and Byzantine fault detectors among other techniques to achieve efficiency and reliability. Similar cryptographic protocol based work also appear in [4, 8, 10, 13].

In summary, we can see that most key distribution schemes for secure multicasting use the hybrid model of key server placement. While this approach is intuitively reasonable, there are no analytical basis to support the model selection. In the next section, we analyze the growth and formation of multicast groups in wide area networks to provide evidence for the correctness of choosing a hybrid model.

3 Analysis of Key Distribution Agent Placement Models

We start our analysis using a regular tree structure which is more tractable than a general network topology. Consider a multicast distribution tree as shown in figure 2 with arity k and depth D where all the leaf nodes represent hosts that could be potential members of a multicast group. The inner nodes represent routers and the nodes at depth $D-1$ denotes sites (or local clusters). Therefore we have a regular network structure with total number of hosts $M = k^D$ and total number of sites $m = k^{D-1}$.

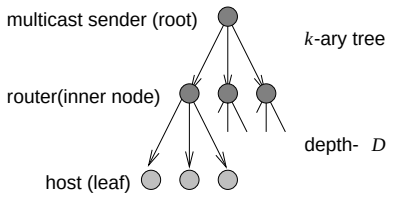

Fig. 2. The basic k-ary tree used to model the multicast distribution tree

3.1 Clustering of Hosts in the Multicast Distribution Tree

First we look at the effect on key distribution schemes due to the clustering of hosts. When we select a number of hosts to create a multicast group (say, of total size n), they could be arbitrarily distributed among several sites. While a single member multicast group will have a node from only a single site, a two member multicast group can select nodes from one or two distinct clusters. Following this argument, we can determine the best possible and worst possible clustering

of hosts in sites when creating a multicast group. The plots of the two curves (equations 1 and 2) are shown in the graph of figure 3.

$$\text{Best case curve: } m = \left\lceil \frac{n}{k} \right\rceil \tag{1}$$

$$\text{Worst case curve: } m = \min\{n, k^{D-1}\} \tag{2}$$

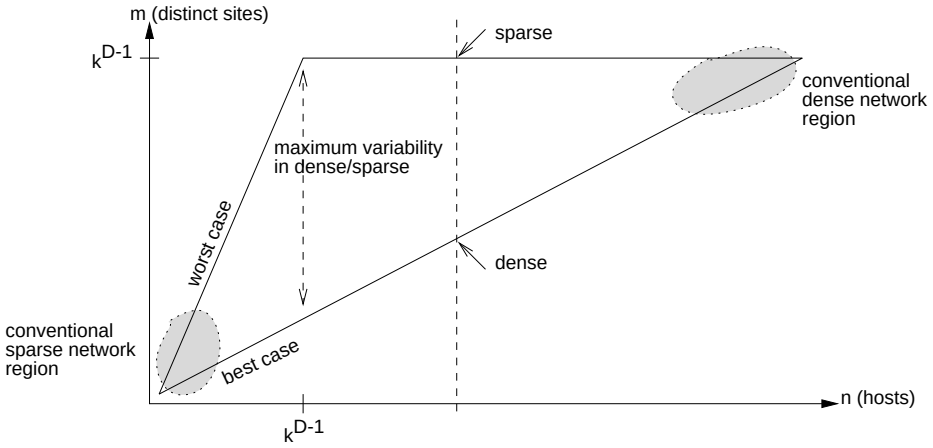

Fig. 3. The graph of number of distinct sites vs. number of hosts in the multicast group shows the allowable variability in denseness/sparseness for a multicast group of given size

We make following observations on the distribution of hosts in sites when setting up a multicast group as derived from the uniform tree structure:

1. The *conventional sparse region* was defined based on the observation that it has relatively very small number of hosts in the group and therefore even in *worst case* can only get distributed into few sites. The recommend key distribution architecture for this scenario is the centralized model. In using any other model, the multicast network will be needlessly using key distribution (sub)agents in inner nodes where most will be unused. In this region, the main issue is efficient use of security agents and not scalability.
2. The *conventional dense region* was defined based on the observation that it has relatively very large number of hosts in the group and therefore even in the *best case* can easily get distributed to nearly all the sites. In this instance, the recommended architecture is the distributed model. Any other model will create a bottleneck situation at the root affecting performance and also make it difficult for the key distribution architecture to scale with

the growth of the multicast network. In this region, the main issues concern both efficiency and scalability.

3. From a practical sense, the most interesting region is the middle area where the variability range is significant. Essentially, this means we might have either a densely populated or sparsely populated multicast network depending on the host distribution among site. Given the large range of sites (m) to which a multicast group of given size (n) can form into, it is quite impractical to discuss an average case scenario. The standard approach would be to use the hierarchical model as the key distribution agent architecture.

3.2 Total Size of the Multicast Distribution Tree

Next we look at the effect of clustering of hosts on the total size of the multicast distribution tree. For the purpose of analyzing the cost of message distribution, we assume a fixed transmission cost for any link in the multicast tree. For a multicast distribution tree represented as a uniform tree structure, the lowest total cost is obtained when hosts are densely located in the smallest possible number of sites as shown in figure 4 (a). The total size of the distribution tree L for a multicast group with n members is obtained by progressively counting the total number of links in all the full sub trees below a given level from top to bottom as shown in equation 3. The quantity ϕ_l denotes the total number of nodes counted prior to level l and the value ρ_l accounts for the link traversed when moving to the next level below to process a partially filled sub tree.

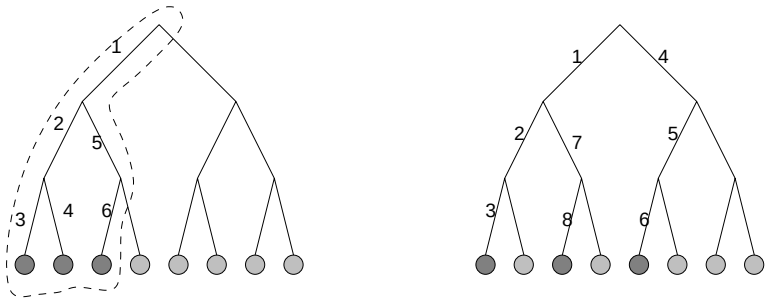

(a) Shortest path length grouping for multicast hosts

(b) Longest path length grouping for multicast hosts

Fig. 4. The best and worst case grouping with respect to number of network links over which messages should pass are given by (a) depth-first search tree and (b) breadth-first search tree

Best case curve:

$$L(n) = \sum_{l=1}^{D} \left\{ \left\lfloor \frac{n-\phi_{l-1}}{k^{D-l}} \right\rfloor \times \left(\left(\sum_{j=1}^{D-l} k^j \right) + 1 \right) + \rho_l \right\}$$

$$\text{where } \phi_l = \begin{cases} 0 & l = 0 \\ \phi_{l-1} + \left\lfloor \frac{n-\phi_{l-1}}{k^{D-(l-1)}} \right\rfloor \times k^{D-(l-1)} & l > 0 \end{cases} \tag{3}$$

$$\text{and } \rho_l = \begin{cases} 0 & n = \phi_{l-1} \\ 1 & otherwise \end{cases}$$

The highest total cost for a multicast distribution tree occurs when the hosts are sparsely distributed among as many sites as possible as shown in figure 4 (b). The distribution is limited by the saturation value φ shown in equation 4 which the maximum number of clusters possible. Several sample plots of the two curves are shown in figure 5.

Worst case curve:

$$L(n) = \begin{cases} nD & n \leq k \\ kD + \sum_{i=1}^{\varphi} \left((k^{i+1} - k^i)(D-i) \right) + (n - k^{\varphi})(D - \varphi) & k^i < n \leq k^{i+1} \end{cases}$$

$$\text{where } \varphi = \left\lfloor \frac{\ln n}{\ln k} \right\rfloor \tag{4}$$

Previously we have discussed non-random clustering of hosts to form a multicast distribution tree in order to study the worst case and best case costs of the delivery tree. Next we look at the random formation of a multicast tree to analyze the total delivery cost for average case. When a host is selected at the leaf level of the tree to form a multicast group, at level l, a route through one of k^l links need to be selected. Therefore, the probability that a given link at level l is in the multicast delivery tree is $\frac{1}{k^l}$. Furthermore, the probability of a link being used in the delivery tree after n hosts have been selected at leaf level is $1 - \left(1 - \frac{1}{k^l}\right)^n$. If hosts are being selected at random at leaf level to form the multicast group, the average number of links at level l that will be included in the delivery tree is $k^l \left(1 - \left(1 - \frac{1}{k^l}\right)^n\right)$. Finally, assuming the link selection process to be a set of independent events, the total size of the multicast tree for a group with n members can be expressed as equation 5 (this result appears in [17] also).

$$L(n) = \sum_{l=1}^{D} k^l \left(1 - \left(1 - \frac{1}{k^l} \right)^n \right) \tag{5}$$

The set of graphs in figure 6 plots the curves for best, average and worst case scenarios for the same k and D. As can be seen from the graphs, the average cost of the multicast delivery tree is closer to the worst case cost for small (and therefore sparse) groups and tends toward best case cost for large (and therefore dense) groups. This result is intuitively correct and validates the expressions

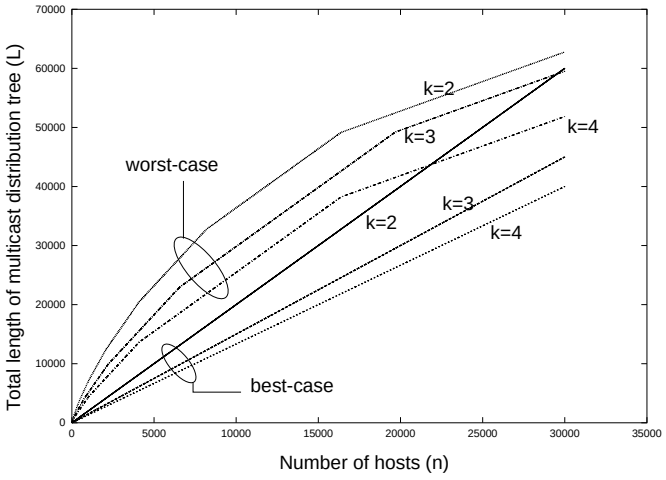

Fig. 5. The graph plots the total size of the multicast distribution tree (in terms of inter-node links) vs. number of hosts assembled in to both worst case and best case multicast groupings. The regular trees have k values 2 ($D = 15$), 3 ($D = 10$) and 4 ($D = 8$)

developed previously to analyze the structure of the multicast tree with respect to clustering of hosts. However, as can be seen from the graphs, the accuracy of the average case curve is lost as the number of hosts increase where the curve dips below the best case result.

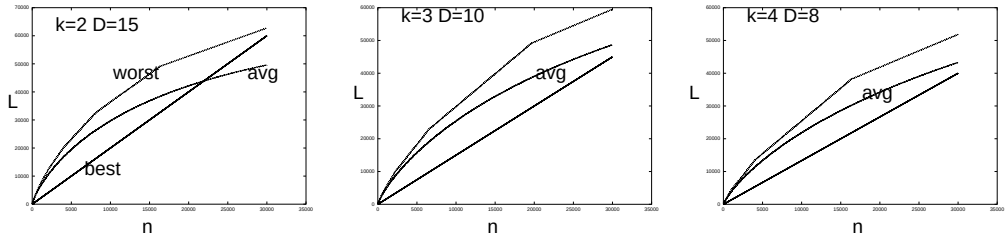

Fig. 6. Total size of the multicast distribution tree vs. number of hosts

The outcome of the foregoing analysis is that for most values of the multicast group size (n), the total size of the multicast distribution tree (L) can vary widely. This behavior again leaves the hierarchical key distribution architecture as the preferred option.

3.3 Applicability of Results to General Multicast Trees

Our analysis so far was based on uniform multicast distribution trees. However, practical multicast distribution structures normally take the shape of irregular trees. An important question at this point is, how relevant the results of an analysis based on uniform trees to real multicast networks? To answer this question, we look at the results obtained by Chuang and Sirbu [5] on the relationship between multicast distribution tree size and size of the membership for general multicast networks. According to the Chuang-Sirbu scaling law, the normalized multicast tree cost is directly proportional to the 0.8 power of the group size (shown in equation 6) for randomly selected group members. The normalized tree cost is obtained as the ratio between total multicast distribution tree length (L_m) and average unicast delivery path length (L_u).

$$\left[\frac{L_m}{L_u}\right]_{general} \propto n^{0.8} \tag{6}$$

We can compute the normalized tree cost for the uniform distribution tree with random member selection using equation 5. The average unicast tree length in this case is the tree depth D. Therefore, for the uniform multicast tree, the normalized tree cost can be given as equation 7.

$$\left[\frac{L_m}{L_u}\right]_{uniform} = \frac{L(n)}{D} = \frac{1}{D}\sum_{l=1}^{D} k^l \left(1 - \left(1 - \frac{1}{k^l}\right)^n\right) \tag{7}$$

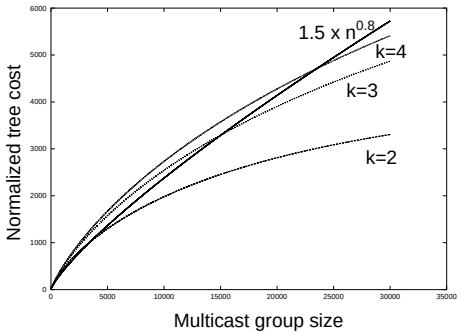

Fig. 7. The graph of normalized distribution tree cost vs. multicast group size with constant of proportionality for Chuang-Sirbu curve set at 1.5

The graph in figure 7 shows that the shape of normalized distribution tree curves for different k values of uniform trees follows that of the general curve due to Chuang-Sirbu scaling law for the range of n in which the average curve

lies between best case and worst case curves of figure 5. The selection of the proportionality constant is admittedly arbitrary, but its function is simply to scale the curves with no distortion of the shape. As shown in the log-scale graph of figure 8, the value was selected for a close fit with plots for uniform trees. The implication of this matching of curves representing theoretical multicast networks to a curve of general multicast networks is, we can expect that for most group membership sizes (n), the average distribution cost (L) of *real* multicast networks also to be in the approximate middle of best case (dense) and worst case (sparse) values.

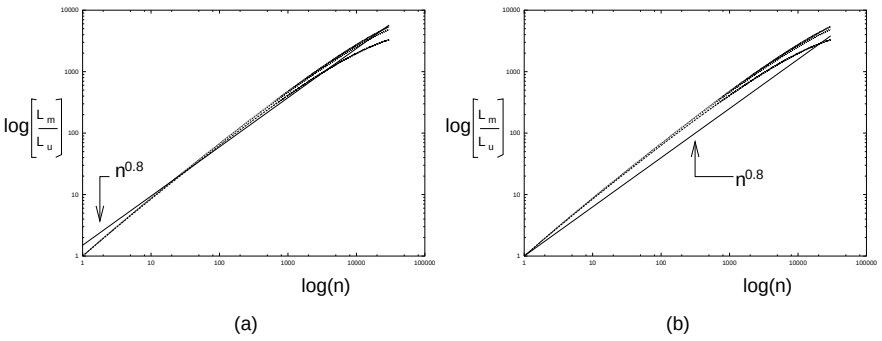

Fig. 8. The graph of normalized distribution tree cost vs. multicast group size with constant of proportionality for Chuang-Sirbu curve set at (a) 1.5 and (b) 1.0

In summary, the significance of this average total distribution cost curve of real multicast networks not being closer to sparse or dense formation of groups is that it is not meaningful to use a centralized or fully distributed control structure for key distribution. This in turn provides an analytical basis for using the hybrid control structure for key distribution.

4 Conclusion

A key distribution framework provides the backbone for any secure multicast architecture. Although the most widely used model for key distribution is the hybrid scheme, the reasons for its selection are usually heuristic arguments of flexibility and scalability. In this work we have used a different approach to validate the use of hybrid model by providing analytical arguments to exclude the use of both centralized and fully distributed control models. Although this work is based on key distribution in multicast networks, the results are applicable in other contexts such as loss recovery where a hierarchical control structure may be used.

References

[1] A. Ballardie. Scalable multicast key distribution. RFC 1949, Network Working Group, May 1996.

[2] A. Ballardie and J. Crowcroft. Multicast-specific security threats and countermeasures. In *Proceedings of the Internet Society Symposium on Network and Distributed System Security (NDSS'95)*, pages 2–16, San Diego, CA, February 1995. IEEE Computer Society Press.

[3] A. Ballardie, P. Francis, and J. Crowcroft. Core based trees (CBT): An architecture for scalable inter-domain routing. *SIGCOMM Computer Communication Review*, 23(4):85–95, October 1993. Conference Proceedings of the Communication Architectures, Protocols and Applications Conference.

[4] R. Canetti, J. Garay, G. Itkis, D. Micciancio, M. Naor, and B. Pinkas. Multicast security: A taxonomy and efficient constructions. In *Proceedings of IEEE INFOCOM'99*, New York, NY, March 1999.

[5] J. C.-I. Chuang and M. A. Sirbu. Pricing multicast communication: A cost-based approach. In *Proceedings of the 8th Annual Internet Society Conference (INET'98)*, Geneva, Switzerland, July 1998. ISOC.

[6] S. E. Deering and D. R. Cheriton. Multicast routing in datagram internetworks and extended LANs. *ACM Transactions on Computer Systems*, 8(2):85–110, May 1990.

[7] S. E. Deering, D. Estrin, D. Farinacci, V. Jacobson, C.-G. Liu, and L. Wei. An architecture for wide-area multicast routing. *SIGCOMM Computer Communication Review*, 24(4):126–135, October 1994. Conference Proceedings of the Communication Architectures, Protocols and Applications Conference.

[8] L. Gong. Enclaves: Enabling secure collaboration over the Internet. In *Proceedings of the 6th USENIX Security Symposium*, pages 149–159, San Jose, CA, July 1996. USENIX.

[9] L. Gong and N. Shacham. Multicast security and its extension to a mobile environment. *ACM-Baltzer Journal of Wireless Networks*, 1(3):281–295, October 1995.

[10] L. Gong and N. Shacham. Trade-offs in routing private multicast traffic. In *Proceedings of IEEE GLOBECOM'95*, Singapore, November 1995.

[11] H. Harney, C. Muckenhirn, and T. Rivers. Group key management protocol (GKMP) architecture. Internet Draft, 1994.

[12] K. P. Kihlstrom, L. E. Moser, and P. M. Melliar-Smith. The SecureRing protocols for securing group communication. In *Proceedings of the 31st Annual Hawaii International Conference on System Sciences (HICSS-31)*, volume 3, pages 317–326, Kona, Hawaii, January 1998. IEEE Computer Society Press.

[13] D. Malkhi, M. Merrit, and O. Rodeh. Secure reliable multicast protocols in a WAN. In *Proceedings of the 17th International Conference on Distributed Computing Systems (ICDCS'97)*, pages 87–94, Baltimore, MD, May 1997. IEEE Computer Society Press.

[14] K. Matsuura, Y. Zheng, and H. Imai. Compact and flexible resolution of CBT multicast key-distribution. In Y. Masunaga, T. Katayama, and M. Tsukamoto, editors, *Proceedings of the Second International Conference on Worldwide Computing and Its Applications (WWCA'98)*, volume 1368 of *Lecture Notes in Computer Science*, pages 190–205, Tsukuba, Japan, March 1998. Springer-Verlag.

[15] S. Mittra. Iolus: A framework for scalable secure multicasting. *SIGCOMM Computer Communication Review*, 27(4):277–288, October 1997. Conference Proceedings of the Communication Architectures, Protocols and Applications Conference.

[16] R. Oppliger and A. Albanese. Participant registration, validation, and key distribution for large-scale conferencing systems. *IEEE Communications Magazine*, 35(6):130–134, June 1997.

[17] G. Phillips, S. Shenker, and H. Tangmunarunkit. Scaling of multicast trees: Comments on the Chuang-Sirbu scaling law. *SIGCOMM Computer Communication Review*, 29(4), October 1999. Conference Proceedings of the Communication Architectures, Protocols and Applications Conference.