

Certificate Dispersal in Ad-Hoc Networks

Mohamed G. Gouda and Eunjin Jung

Department of Computer Sciences

The University of Texas at Austin

E-mail: {gouda, ejung}@cs.utexas.edu

Abstract

We investigate how to disperse the certificates, issued in an ad-hoc network, among the network nodes such that the following condition holds. If any node u approaches any other node v in the network, then u can use the certificates stored either in u or in v to obtain the public key of v (so that u can securely send messages to v). We define the cost of certificate dispersal as the average number of certificates stored in one node in the network. We give upper and lower bounds on the dispersability cost of certificates, and show that both bounds are tight. We also present two certificate dispersal algorithms, and show that one of those algorithms is more efficient than the other in several important cases. Finally, we identify a rich class of “certificate graphs” for which the dispersability cost is within a constant factor from the lower bound.

1. Introduction

We consider a network where each node u has a private key $rk.u$ and a public key $bk.u$. In this network, in order for a node u to securely send a message m to another node v , node u needs to encrypt the message m using the public key $bk.v$, before sending the encrypted message, denoted $bk.v < m >$, to node v . This necessitates that node u know the public key $bk.v$ of node v .

If a node u knows the public key $bk.v$ of another node v in this network, then node u can issue a certificate, called a certificate from u to v , that identifies the public key $bk.v$ of node v . This certificate can be used by any node in the network that knows the public key of node u to further acquire the public key of node v .

A certificate from node u to node v is of the following form:

$$rk.u < u, v, bk.v >$$

This certificate is signed using the private key $rk.u$ of node u , and it includes three items: the identity of the certificate

issuer u , the identity of the certificate subject v , and the public key of the certificate subject $bk.v$. Any node that knows the public key $bk.u$ of node u can use $bk.u$ to decrypt the certificate from u to v and obtain the public key $bk.v$ of node v .

The certificates issued by different nodes in a network can be represented by a directed graph, called the *certificate graph* of the network. Each node in the certificate graph represents a node in the network. Each directed edge from node u to node v in the certificate graph represents a certificate issued by node u for node v in the network.




Figure 1. A certificate graph example

Figure 1 shows a certificate graph for a network with five nodes: a , b , c , d , and e . According to this graph,

- node a issued two certificates (a, b) , and (a, d)
- node b issued one certificate (b, c)
- node c issued no certificate
- node d issued one certificate (d, e)
- node e issued one certificate (e, b) .

Node a can use the two certificates (a, b) and (b, c) to obtain the two public keys $bk.b$ and $bk.c$, and so can securely send messages to nodes b and c . Also, node a can use the two certificates (a, d) and (d, e) to obtain the public keys $bk.d$ and $bk.e$, and can securely send messages to nodes d and e . Node d can use the three certificates (d, e) , (e, b) , and (b, c) to obtain the public keys $bk.e$, $bk.b$, and $bk.c$, and can securely send messages to nodes e , b , and c .

We assume that the network is ad-hoc and its nodes are mobile so move around. In this case, the issued certificates

need to be dispersed among the nodes in the network such that if a node u approaches another node v and wishes to securely send messages to it, then u can obtain the public key of v from the set of certificates stored either in u or in v (provided there is a directed path from u to v in the certificate graph).

As an example, assume that each node in the certificate graph in Figure 1 stores the certificates in the maximal, shortest-path, incoming tree rooted at the node:

node a stores no certificates	
node b stores the certificates	$(a, b), (d, e), (e, b)$
node c stores the certificates	$(a, b), (d, e), (e, b),$ (b, c)
node d stores the certificates	(a, d)
node e stores the certificates	$(a, d), (d, e)$

Thus, if node a approaches node e , then a can use the two certificates stored in e to obtain the public key of e and securely send messages to e . (Note that node e can never obtain the public key of node a because there is no directed path from node e to node a in the certificate graph in Figure 1).

As an application of this situation, consider the tanks in an armored division. Each tank has a computer and can be viewed as a node in an ad-hoc network. Before the tanks are deployed into the field, a certificate graph needs to be designed to secure the future communications between the tanks in the field. Then the certificates from this certificate graph need to be dispersed amongst the tanks before they are deployed. Later, the tanks are deployed into the field and each of them has a number of certificates in its local storage. Now, if two tanks approach each other in the field, then the two tanks have enough certificates in their local storage so that each of them can compute the public key of the other and two tanks can securely exchange messages.

In this paper, we discuss three contributions to the problem of certificate dispersal in ad-hoc networks. First, we present tight upper and lower bounds on the average number of certificates stored in one node. Second, we present a somewhat efficient certificate dispersal algorithm that ensures that the average number of certificates stored in a node is small (if not the smallest). Third, we identify rich classes of certificate graphs for which the average number of certificates to be stored in a node approaches the lower bound.

2. Related Work

Several papers have investigated the use of certificates to provide security in traditional and in ad-hoc networks. We summarize the results of these papers in the following paragraphs.

Architectures for issuing, storing, discovery, and validating certificates in traditional networks are presented in [14],

[2], [12], [5], [1], [4], [11], and [6]. There are several limitations in ad-hoc networks that do not allow us to use these results without significant overhead to the networks. In traditional networks, one can assume that the communication between nodes is reliable and reasonably fast. In ad-hoc networks, finding routes between nodes itself is challenging, let alone maintaining the relevant information of the found routes. Also, nodes in ad-hoc networks often have very limited resources. For example, computational capability, storage, and power supply are much less than what most nodes have in traditional networks. To accommodate these limitations, different architectures for issuing, storing, discovery, and validating certificates in ad-hoc networks have been developed.

In [15], Zhou and Haas have presented an architecture for issuing certificates in an ad-hoc network. According to this architecture, the network has k servers. Each server has a different share of some private key rk . To generate a certificate, each server uses its own share of rk to encrypt the certificate. If no more than t servers have suffered from Byzantine failures, where $k \geq 3t + 1$, then the resulting certificate is correctly signed using the private key rk , thanks to threshold cryptography. The resulting certificate can be decrypted using the corresponding public key which is known to every node in the ad-hoc network.

In [10], Kong, Perfos, Luo, Lu and Zhang presented more distributed architecture for issuing certificates. Instead of employing k servers in the ad-hoc network, each node in the network is provided with a different share of the private key rk . For a node u to issue a certificate, the node u forwards the certificate to its neighbors and each of them encrypt the certificate using its share of rk . If node u has at least $t + 1$ correct neighbors (i.e. they have not suffered from any failures), the resulting certificate is correctly signed using the private key rk .

In our paper, we propose an architecture where every node has both a public key and a private key so it can issue certificates for any other node in the network. This architecture is very efficient in issuing and validating certificates but cannot tolerate Byzantine failures. In particular, if one node suffers from Byzantine failure, then this node can successfully impersonate any other node that is reachable from this node in the certificate graph of the network. This vulnerability to Byzantine failures is not unique to our certificate work in ad-hoc networks. In fact, many proposed certificate architectures, e.g. [14], [2], [12], [11], and [6] yield similar vulnerabilities in traditional networks. Recently, we have developed a technique for augmenting the certificates with additional information to ensure that the network can tolerate some degree of Byzantine failures[9].

Perhaps the closest work to ours is [8] where the authors, Hubaux, Buttyán, and Capkun, investigated how to disperse certificates in a certificate graph among the network nodes under two conditions. First, each node stores

the same number of certificates. Second, with high probability, if two nodes meet then they have enough certificates for each of them to compute the public key of the other. By contrast, our work is based on two different conditions. First, different nodes may store different number of certificates, but the average number of certificates stored in one node is minimized. Second, it is guaranteed (i.e. with probability 1) that if two nodes meet then they have enough certificates for each of them to compute the public key of the other.

Later, the same authors have showed in [3] that a lower bound on the number of certificates to be stored in a node is $\sqrt{n} - 1$ where n is the number of nodes in the system. By contrast, we show below that the tight lower bound on the average number of certificates to be stored in a node is e/n , where e is the number of edges in the system.

3. Certificate Dispersal

A *certificate graph* G is a directed graph in which each directed edge, called a *certificate*, is a pair (u, v) , where u and v are distinct nodes in G . For each certificate (u, v) in G , u is called the *issuer* of the certificate and v is called the *subject* of the certificate.

Note that according to this definition a certificate graph is a directed graph that does not have self-loops and does not have multiple edges from any node to any other node.

A directed path $(v_0, v_1), (v_1, v_2), \dots, (v_{k-1}, v_k)$ in a certificate graph G , where the nodes v_0, v_1, \dots, v_k are distinct, is called a *certificate chain* from v_0 to v_k . The *length* of a chain is the number of certificates in the chain. A chain from u to v is *shortest* iff its length is not larger than the length of any other chain from u to v in the same certificate graph.

Let c denote the chain $(v_0, v_1), \dots, (v_{k-1}, v_k)$. Then, each of the chains $(v_0, v_1), \dots, (v_{j-1}, v_j)$, where $1 \leq j \leq k$, is called a *prefix* of c that ends at node v_j . Also, each of the chains $(v_j, v_{j+1}), \dots, (v_{k-1}, v_k)$, where $0 \leq j \leq k-1$, is called a *suffix* of c that starts at node v_j .

A *certificate dispersal algorithm* F is an algorithm that takes as input any certificate graph G and computes a subset of the certificates, denoted $F(G, v)$, for every node v in G such that the following two conditions hold:

i. *Connectivity* :

For every distinct pair of nodes u and v in G , if there is a chain from u to v in G , then there is a chain from u to v in the set $F(G, u) \cup F(G, v)$.

ii. *Completeness* :

For every certificate in G , there is a node v in G such that this certificate is in $F(G, v)$.

Let G be a certificate graph, and F be a certificate dispersal algorithm. The *cost* of using F to disperse the certificates in G among the nodes of G , denoted $c.(F, G)$, is com-

puted as follows:

$$c.(F, G) = \frac{1}{n} \left(\sum_{v \in G} |F.(G, v)| \right)$$

where n is the number of nodes in G , and $|F.(G, v)|$ denotes the number of certificates in the set $F.(G, v)$ assigned by F to node v . Note that $c.(F, G)$ is the average number of certificates assigned by F to a node in G .

The *dispersability cost* of a certificate graph G , denoted $d.G$, is computed as follows:

$$d.G = \min_F c.(F, G)$$

A certificate dispersal algorithm F_e is *efficient* iff for every certificate graph G ,

$$c.(F_e, G) = d.G$$

It follows from this definition that if F_e is an efficient certificate dispersal algorithm and F is a certificate dispersal algorithm then for every certificate graph G ,

$$c.(F_e, G) \leq c.(F, G)$$

Lemma 1 : (*Upper Bound on Dispersability Cost*)

For every certificate graph G with n nodes,

$$d.G \leq n - 1$$

Proof :

In the next section, we present a certificate dispersal algorithm F_{full} that assigns to every node v in a certificate graph G , the certificates in a maximal outgoing tree rooted at v . Because each maximal outgoing tree in a certificate graph G has at most $n - 1$ certificates, where n is the number of nodes in G , we have

$$|F_{full}(G, v)| \leq n - 1$$

$$c.(F_{full}, G) = \frac{1}{n} \left(\sum_{v \in G} |F.(G, v)| \right) \leq (n - 1)$$

Let F_e be any efficient dispersal algorithm; hence

$$d.G = c.(F_e, G) \leq c.(F_{full}, G) \leq n - 1$$

□

Lemma 2 : (*Lower Bound on Dispersability Cost*)

For every certificate graph G with n nodes and e certificates,

$$d.G \geq \frac{e}{n}$$

Proof :

Let F_e be any efficient certificate dispersal algorithm. From the completeness condition of a dispersal algorithm, every

edge in a certificate graph G is assigned by F_e to some node in G . Thus,

$$\sum_{v \in G} |F_e.(G, v)| \geq e$$

where e is the number of certificates in G . We conclude

$$d.G = c.(F_e, G) = \frac{1}{n} \left(\sum_{v \in G} |F_e.(G, v)| \right) \geq \frac{e}{n}$$

□

Lemma 3 : (*Achieving the two Bounds on Dispersability Cost*)

There is a certificate graph G with n nodes and e certificates such that the following two conditions hold.

- i. $d.G = n - 1$
- ii. $d.G = \frac{e}{n}$

Proof :

Let G be a fully connected certificate graph; i.e., for any two distinct nodes u and v in G , there is a certificate from u to v and a certificate from v to u . Thus, the number of certificates e in G is $n(n - 1)$, and the following relations hold in G .

$$n - 1 = \frac{n(n - 1)}{n} = \frac{e}{n}$$

Hence the upper and lower bounds on dispersability cost meet at G , and $d.G$ is equal to each of the two bounds. Therefore, the two conditions of Lemma 3 hold for G .

□

It is also possible to construct certificate graphs that meet the upper bound on dispersability cost but not its lower bound, and to construct certificate graphs that meet the lower bound on dispersability cost but not its upper bound.

As an example, consider a *ring* certificate graph G_0 in Figure 2. This graph has n nodes and n certificates arranged in a directed ring. It is straightforward to show that any dispersal algorithm will assign to every node in the graph at least $n - 1$ certificates. Thus, $d.G_0 \geq n - 1$, and G_0 meets the upper bound on dispersability cost but not its lower bound.




Figure 2. A ring certificate graph

As a second example, consider an *hourglass* certificate graph G_1 in Figure 3. This graph has n nodes and $n - 1$ certificates, where n is odd, arranged in an hourglass shape with




Figure 3. An hourglass certificate graph

one center node, $(n - 1)/2$ input nodes, and $(n - 1)/2$ output nodes. There is a certificate dispersal algorithm F that assigns certificates to every node v in G_1 as follows.

- i. If v is the center node, then $F.(G_1, v) = \{\}$.
- ii. If v is an input node, then

$$F.(G_1, v) = \{(v, \text{center node})\}.$$

- iii. if v is an output node, then

$$F.(G_1, v) = \{(\text{center node}, v)\}.$$

Thus,

$$\begin{aligned} c.(F, G_1) &= \frac{n - 1}{n} \\ &= \text{the lower bound on dispersability cost} \\ &= d.G_1 \end{aligned}$$

Hence, G_1 meets the lower bound on dispersability cost but not its upper bound.

The above discussion suggests the following two problems which we explore in the rest of this paper.

Problem 1 :

Develop an efficient certificate dispersal algorithm F_e .

□

Problem 2 :

Identify rich classes of certificate graphs whose dispersability costs meet the lower bound or are within a constant factor of this lower bound.

□

Problem 1 remains open: Instead of solving Problem 1, we present two certificate dispersal algorithms F_{full} and F_{half} (in Sections 4 and 5 respectively), and show that in several important cases F_{full} is not as efficient as F_{half} . We then present a solution of Problem 2 in Section 6.

4. Full Tree Algorithm for Certificate Dispersal

Before we introduce our first certificate dispersal algorithm, we need to introduce the following definition of chain sets.

Let G be a certificate graph and v be a node in G . A *chain set* for v , denoted $S.v$, is a set of chains in G that satisfies the following three conditions.

- i. If G has no chains that starts at v , then $S.v$ is empty.
- ii. If G has a chain from v to w , then $S.v$ has exactly one shortest chain from v to w .
- iii. If $S.v$ has a chain, then $S.v$ also has every nonempty prefix of this chain.




Figure 4. The diamond certificate graph

As an example, consider the *diamond* certificate graph in Figure 4. In this graph, there are no certificate chains that start at node e or f , and the chain sets for node e and f are both empty:

$$S.e = \{\}, S.f = \{\}$$

The chain set for node d has two chains:

$$S.d = \{ \langle (d, e) \rangle, \langle (d, f) \rangle \}$$

Also the chain set for each of the two nodes b and c has three chains:

$$S.b = \{ \langle (b, d) \rangle, \langle (b, d); (d, e) \rangle, \langle (b, d); (d, f) \rangle \}$$

$$S.c = \{ \langle (c, d) \rangle, \langle (c, d); (d, e) \rangle, \langle (c, d); (d, f) \rangle \}$$

The chain set for node a has five chains:

$$S.a = \{ \langle (a, b) \rangle, \langle (a, c) \rangle, \langle (a, b); (b, d) \rangle, \langle (a, c); (c, d) \rangle, \langle (a, b); (b, d); (d, e) \rangle, \langle (a, c); (c, d); (d, f) \rangle \}$$

The following two comments are in order. First, each chain set $S.v$ for a node v defines a maximal, shortest-path, outgoing tree rooted at node v in the certificate graph. Second, it is possible to have two or more distinct chain sets for a node. For example, a second chain set for node a in the certificate graph in Figure 4 is as follows:

$$\{ \langle (a, b) \rangle, \langle (a, c) \rangle, \langle (a, b); (b, d) \rangle, \langle (a, b); (b, d); (d, e) \rangle, \langle (a, b); (b, d); (d, f) \rangle \}$$

Using the above definition of a chain set, we are now ready to present our first certificate dispersal algorithm, called the *full tree algorithm* and denoted F_{full} . This algorithm assigns to every node v all the certificates in a chain set $S.v$ for v . In other words,

$F_{full} \cdot (G, v)$ = the set of all certificates that exist in a chain set $S.v$ for v .

Lemma 4 :

F_{full} is a certificate dispersal algorithm.

Proof :

We show that F_{full} satisfies the two conditions of a certificate dispersal algorithm, connectivity and completeness. First, if there is a chain from u to v in G , then at least one of the shortest chains from u to v is in $S.u$ by condition *ii* in the definition of chain set. Second, any certificate (u, v) in G is in $S.u$ since it is the shortest chain from u to v . By the definition of F_{full} , the certificate (u, v) is in $F_{full} \cdot (G, u)$. Therefore, F_{full} satisfies two properties of connectivity and completeness. \square

Next, we show that the dispersal algorithm F_{full} is far from being efficient. First, we show in Lemma 5 that the cost of applying F_{full} to any strongly connected certificate graph meets the upper bound on dispersability cost. Second, we show in Lemma 6 that the cost of applying F_{full} to any hourglass certificate graph is within a factor of four from the upper bound on dispersability cost (even though the dispersability cost of an hourglass graph meets the lower bound).

Lemma 5 :

For any strongly connected certificate graph G with n nodes,

$$c.(F_{full}, G) = n - 1$$

Proof :

The certificate dispersal algorithm F_{full} assigns, to every node v in a certificate graph G , the certificates in a maximal outgoing tree rooted at v . If G is strongly connected, then any maximal outgoing tree is in fact a spanning tree with $(n - 1)$ certificates, where n is the number of nodes in G . Therefore, for any node v in G ,

$$|F_{full} \cdot (G, v)| = n - 1$$

$$c.(F_{full}, G) = \frac{1}{n} \left(\sum_{v \in G} |F_{full} \cdot (G, v)| \right) = n - 1$$

\square

Lemma 6 :

For any hourglass certificate graph G with n nodes (see Figure 3),

$$c.(F_{full}, G) = \frac{n^2 + 2n - 3}{4n} \sim \frac{n}{4}$$

Proof :

Recall that any hourglass certificate graph G has one center node, $\frac{n-1}{2}$ input nodes, and $\frac{n-1}{2}$ output nodes.

$$|F_{full} \cdot (G, \text{center})| = \frac{n-1}{2}$$

For every input node v ,

$$|F_{full} \cdot (G, v)| = \frac{n-1}{2} + 1 = \frac{n+1}{2}$$

For every output node v ,

$$|F_{full}(G, v)| = 0$$

Thus,

$$\begin{aligned} c.(F_{full}, G) &= \frac{1}{n} \left(\frac{n-1}{2} + \left(\frac{n-1}{2} \right) \left(\frac{n+1}{2} \right) \right) \\ &= \frac{n^2 + 2n - 3}{4n} \sim \frac{n}{4} \end{aligned}$$

□

5. Half Tree Algorithm for Certificate Dispersal

Before we introduce our second certificate dispersal algorithm, we need to introduce the following definition of consistent chain sets.

Let $S.u$ and $S.v$ be two chain sets for nodes u and v , respectively, in a certificate graph G . $S.u$ and $S.v$ are *consistent* iff for every two nodes x and y in G , if $S.u$ has a subchain that starts at x and ends at y and $S.v$ also has a subchain that starts at x and ends at y then these two subchains are identical.

A collection of chain sets $\{S.v | v \text{ is a node in } G\}$ is *consistent* iff any two chain sets in the collection are consistent.

We are now ready to present our second certificate dispersal algorithm, called the *half tree algorithm* and denoted F_{half} . This algorithm takes as input a consistent collection of chain sets $\{S.v | v \text{ is a node in a certificate graph } G\}$ and computes a set of certificates $F_{half}(G, v)$ for every node v in G . Algorithm F_{half} is defined as follows.

```

1: for every nonempty  $S.v$  in the consistent collection
   of chain sets do
2:   let  $c$  denote the longest chain  $\langle (v_0, v_1); \dots; (v_{k-1}, v_k) \rangle$ 
     in  $S.v$ : note that  $v_0 = v$ ;
3:   let  $x := \lfloor \frac{k}{2} \rfloor$ ;
4:   find the largest  $y$ ,  $0 \leq y \leq k$ , such that all
     certificates in the prefix  $\langle (v_0, v_1); \dots; (v_{y-1}, v_y) \rangle$ 
     are already in  $F_{half}(G, v)$ ;
5:   if  $x \leq y$ 
6:   then
     store the certificates in every prefix of
     the subchain  $\langle (v_y, v_{y+1}); \dots; (v_{k-1}, v_k) \rangle$ 
     in  $F_{half}(G, w)$  where  $w$  is the node at
     which the prefix ends;
7:   else
7a:    store the certificates in the prefix
      $\langle (v_y, v_{y+1}); \dots; (v_{x-1}, v_x) \rangle$  in  $F_{half}(G, v)$ ;

```

```

7b:    store the certificates in every prefix of
     the subchain  $\langle (v_x, v_{x+1}); \dots; (v_{k-1}, v_k) \rangle$ 
     in  $F_{half}(G, w)$  where  $w$  is the node at
     which the prefix ends;
     endif;
8:    remove chain  $c$  from  $S.v$ ;
9:  enddo;

```

Lemma 7 :

F_{half} is a certificate dispersal algorithm.

Proof :

First, if there is a chain between nodes u and v , then at least one of the shortest chains from u to v is stored in $S.u$. All the certificates in the chain from u to v will be stored in u and v by the definition of F_{half} . Second, any certificate (u, v) in G will be stored in $S.u$ since it will be the shortest chain from u to v . By the definition of F_{half} , the certificate (u, v) is stored either in u or in v . Therefore, F_{half} satisfies two properties of certificate dispersal algorithm. □

Next, we show in Lemma 8 that in the important case of strongly connected certificate graphs, F_{half} is not less efficient than F_{full} , and in some instances, F_{half} is in fact more efficient than F_{full} . Then in Lemma 9, we show that in the important case of tree certificate graphs, F_{half} is not less efficient than F_{full} , and in some instances, F_{half} is in fact more efficient than F_{full} . In Lemma 10, we show that in the case of the hourglass certificate graphs F_{half} achieves the lower bound on dispersability cost which is much less than what F_{full} achieves.

Lemma 8 :

For any strongly connected certificate graph G ,

$$c.(F_{half}, G) \leq c.(F_{full}, G)$$

For some strongly connected certificate graph G ,

$$c.(F_{half}, G) < c.(F_{full}, G)$$

Proof :

Let G be any strongly connected certificate graph, and v be any node in G . The certificates in the set $F_{half}(G, v)$ define a graph G' , which is a subgraph of the original graph G . In G' , there can be at most one path from any node to node v , and at most one path from node v to any other node. Graph G' satisfies exactly one of the following two conditions.

- i. G' has no cycle.
- ii. G' has a cycle, but it has at most $n - 1$ nodes.

In the first case, the number of certificates in G' is at most $n - 1$, since there is no cycle in G' . In the second case, the number of certificates in G' is also at most $n - 1$, which is

the number of certificates if all the $n - 1$ nodes participate in the cycle. Therefore, $|F_{half} \cdot (G, v)| \leq n - 1$.

$$c.(F_{half}, G) = \frac{1}{n} \sum_{v \in G} |F_{half} \cdot (G, v)| \leq \frac{n(n-1)}{n} = n - 1$$

Because G is strongly connected, $c.(F_{full}, G) = n - 1$ by Lemma 5. Therefore,

$$c.(F_{half}, G) \leq c.(F_{full}, G)$$

This completes our proof of the first part of the lemma.



Figure 5. The two-ring certificate graph

To prove the second part of the lemma, consider the two-ring certificate graph G'' in Figure 5. This graph is strongly connected and has three nodes. Then by Lemma 5,

$$c.(F_{full}, G'') = n - 1 = 2$$

By applying F_{half} to G'' , we get

$$\begin{aligned} F_{half} \cdot (G'', u) &= \{(u, v), (v, u)\} \\ F_{half} \cdot (G'', v) &= \{\} \\ F_{half} \cdot (G'', w) &= \{(v, w), (w, v)\} \end{aligned}$$

Therefore,

$$c.(F_{half}, G'') = \frac{1}{3}(2 + 0 + 2) = \frac{4}{3} < c.(F_{full}, G'')$$

□

Lemma 9 :

For every tree certificate graph T ,

$$c.(F_{half}, T) \leq c.(F_{full}, T)$$

For any complete tree certificate graph G ,

$$c.(F_{half}, G) < c.(F_{full}, G)$$

(Due to space constraints, the proof of Lemma 9 is in [7].)

□

Lemma 10 :

For any hourglass certificate graph G with n nodes and e certificates (see Figure 3) where n is odd,

$$c.(F_{half}, G) = \frac{e}{n} < c.(F_{full}, G)$$

Proof :

Recall that an hourglass certificate graph G with n nodes has one center node, $\frac{n-1}{2}$ input nodes, and $\frac{n-1}{2}$ output nodes. Applying F_{half} to this certificate graph, we get

$$\begin{aligned} \text{for every input node } u, \quad & F_{half} \cdot (G, u) = \{(u, c)\} \\ \text{for the center node } c, \quad & F_{half} \cdot (G, c) = \{\} \\ \text{for every output node } w, \quad & F_{half} \cdot (G, w) = \{(c, w)\} \end{aligned}$$

Therefore,

$$c.(F_{half}, G) = \frac{n-1}{n} = \frac{e}{n} < \frac{n}{4} \sim c.(F_{full}, G)$$

□

6. Certificate Graphs with Small Dispersability Costs

In this section, we consider Problem 2, stated in Section 2, and identify a class of strongly connected certificate graphs that have small dispersability costs. This class is based on the star certificate graph G in Figure 6.




Figure 6. A star certificate graph

The star graph consists of one center node and m satellite nodes. In this graph, the center node is connected with each of the satellite nodes by a directed ring. The best way to disperse the certificates in this graph is to assign the two certificates in each ring to the satellite node in that ring, while assigning no certificates to the center node. Thus,

$$d.G = \frac{1}{m+1}(2m)$$

Because the number of nodes in this graph is $(m+1)$ and the number of certificates in this graph is $2m$, the dispersability cost of this graph achieves its lower bound.

Unfortunately, the star certificate graph has a security problem. If the private key of the center node is compromised by an adversary, then this adversary can impersonate any node as it communicates with any other node in the system. Thus, although the dispersability cost of this graph achieves the lower bound, the security problem of this graph compels us to seek other certificate graphs. Next, we discuss how to generalize this star graph into a class of graphs whose dispersability costs are small and whose security problems are not so severe.

In our generalization of the star certificate graph, each of the m rings in the graph has k satellite nodes (beside the

center node that exists in every ring). We refer to this generalized certificate graph as (m, k) -star. Figure 7 shows an $(m, 2)$ -star.




Figure 7. An $(m, 2)$ -star certificate graph

An efficient way to disperse the certificates in the (m, k) -star G is to assign the $(k + 1)$ certificates in a ring to every satellite node in that ring, while assigning no certificates to the center node. Thus,

$$d.G = \frac{1}{mk + 1}(mk(k + 1))$$

The number of nodes in the (m, k) -star is $mk + 1$, and the number of certificates in this graph is $m(k + 1)$. Thus, the lower bound on the dispersability cost for the (m, k) -star is $\frac{m(k+1)}{mk+1}$. Therefore, the dispersability cost of an (m, k) -star is within a factor of k from its lower bound.

It is straightforward to show that an (m, k) -star, where $k \geq 2$, has better security properties than the original $(m, 1)$ -star. In particular, if the private key of the center node is compromised by an adversary then this adversary will not be able to impersonate any satellite node in a ring while it communicates with any other satellite node in the same ring.

7. Conclusion

In this paper, we introduce the problem of certificate dispersal. Tight lower and upper bounds on the cost of the certificate dispersability are given along with example certificate graphs that achieve both bounds. Two certificate dispersal algorithms, F_{full} and F_{half} , that can reach these bounds for certain graphs are devised. The algorithm F_{full} makes each node store a maximal shortest-path outgoing tree, whereas F_{half} makes each node store half of an outgoing tree and half of an incoming tree. We show that F_{half} performs better than F_{full} in strongly connected certificate graphs and in tree certificate graphs. We also present a class of certificate graphs whose dispersability costs is within a constant factor from the lower bound.

The dispersal algorithm discussed in this paper assume that all the directed paths between two nodes are equally good. In some applications this may not be true, as discussed in [13]. Further research is needed to adapt the dispersal algorithms in this paper to these applications.

Acknowledgement This work was supported in part by the Defense Advanced Research Projects Agency (DARPA) Contract F33615-01-C-1901, and in part by three IBM Faculty Partnership Awards for the academic years of 2000–2001, 2001–2002, and 2002–2003, and by the Texas Advanced Research Program, Texas Higher Education Coordinating Board, under Grant TARP 14-970823.

References

- [1] M. Blaze, J. Feigenbaum, J. Ioannidis, and A. Keromytis. The keynote trust-management system version 2. RFC 2704, 1999.
- [2] S. Boeyen, T. Howes, and P. Richard. Internet X.509 public key infrastructure operational protocols - LDAPv2. RFC 2559, 1999.
- [3] S. Capkun, L. Buttyán, and J.-P. Hubaux. Self-organized public-key management for mobile ad hoc networks. *IEEE Transactions on Mobile Computing*, 2(1):52–64, 2003.
- [4] D. Clarke, J.-E. Elien, C. Ellison, M. Fredette, A. Morcos, and R. Rivest. Certificate chain discovery in SPKI/SDSI. *Journal of Computer Security*, 9(4):285–322, 2001.
- [5] C. Ellison, B. Frantz, B. Lampson, R. Rivest, B. Thomas, and T. Ylonen. SPKI certificate theory. RFC 2693, 1999.
- [6] E. Freudenthal, T. Pesin, L. Port, E. Keenan, and V. Karamcheti. dRBAC: distributed role-based access control for dynamic coalition environments. In *Proceedings. 22nd International Conference on Distributed Computing Systems, 2002*, pages 411–420, 2002.
- [7] M. G. Gouda and E. Jung. Certificate dispersal in ad hoc networks. Technical report, 2003.
- [8] J.-P. Hubaux, L. Buttyán, and S. Capkun. The quest for security in mobile ad hoc networks. In *Proceedings of the 2001 ACM International Symposium on Mobile ad hoc networking & computing*, pages 146–155. ACM Press, 2001.
- [9] E. Jung and M. G. Gouda. Strongly chained certificates. *in preparation*, 2003.
- [10] J. Kong, P. Zerfos, H. Luo, S. Lu, and L. Zhang. Providing robust and ubiquitous security support for wireless mobile networks. In *Ninth International Conference on Network Protocols (ICNP'01)*, pages 251–260, 2001.
- [11] N. Li, W. H. Winsborough, and J. C. Mitchell. Distributed credential chain discovery in trust management: extended abstract. In *Proceedings of the 8th ACM conference on Computer and Communications Security*, pages 156–165. ACM Press, 2001.
- [12] M. Myers, R. Ankney, A. Malpani, S. Galperin, and C. Adams. X.509 internet public key infrastructure online certificate status protocol - OCSP. RFC 2560, 1999.
- [13] M. K. Reiter and S. G. Stubblebine. Authentication metric analysis and design. *ACM Transactions on Information and System Security (TISSEC)*, 2(2):138–158, 1999.
- [14] R. L. Rivest and B. Lampson. SDSI – A simple distributed security infrastructure. Presented at CRYPTO'96 Rumpsession, 1996.
- [15] L. Zhou and Z. J. Haas. Securing ad hoc networks. *IEEE Network*, 13(6):24–30, 1999.