

Efficient Fault-Tolerant Certificate Revocation*

Rebecca N. Wright
AT&T Labs – Research
180 Park Avenue
Florham Park, NJ 07932 USA
rwright@research.att.com

Patrick D. Lincoln, Jonathan K. Millen, Andrea I. Lincoln
SRI International
333 Ravenswood Ave
Menlo Park, CA 94025 USA
{millen,lincoln}@cs1.sri.com

June 5, 2000

Abstract

We consider scalable certificate revocation in a public-key infrastructure. We introduce depender graphs, a new class of graphs which support efficient and fault-tolerant revocation. Nodes of a depender graph are participants that agree to forward revocation information to other participants. Our depender graphs are k -redundant, so that revocations are provably guaranteed to be received by all non-failed participants even if up to $k-1$ participants have failed. We present a protocol for constructing k -redundant depender graphs that has two desirable properties. First, it is load-balanced, in that no participant need have too many dependers. Second, it is localized, in that it avoids the need for any participant to maintain the global state of the depender graph. We also give a localized protocol for restructuring the graph in the event of permanent failures.

1 Introduction

Public keys and their certificates eventually become invalid. Most certificates have an expiration date, but for various reasons a certificate may become invalid prior to the expiration date. For example, the secret key may have been lost or compromised. The owner's identifying information, which might include an e-mail address or employer, may have changed. The certificate might have been used to enable organizational privileges which have been withdrawn by the employer. Under these circumstances, there should be some way to revoke the certificate.

1.1 Existing Approaches

Current proposed standards for revocation, as found in the X.509 directory framework [X.509], and the Internet draft standard Public Key Infrastructure [AZ98], involve *certificate revocation*

*Some of the ideas in this paper appeared in exploratory form in: J. Millen and R. Wright, "Certificate Revocation the Responsible Way," *Proc. Computer Security, Dependability, and Assurance: From Needs to Solutions (CSDA'98)*, IEEE Computer Society, (1999), pp. 196–203.

lists (CRLs) maintained on key servers, which act as a repositories for certificates. To revoke a certificate, the owner sends the key server a *revocation notice*, which is a signed message identifying the certificate to be revoked.

Upon receipt of a revocation notice, the key server updates its CRL and no longer gives out the revoked certificate. End users who want to check the validity of a certificate must query the key server, and in response receive all or part of the latest full, signed, CRL, which is updated periodically. Good discussions of revocation technologies can be found in [FL98] and [Mey98]. There are various strategies for reducing communication and storage costs while maintaining timeliness of revocation, such as Kocher's certificate revocation tree [Koc98] and related advances [KAN99, NN98], and methods for reducing server load, as in [Coo99, MJ00].

One can try to reduce the need for revocation by limiting certificates to brief expiration periods, but this increases server load because new certificates must be sent more frequently. Rivest [Riv98] suggested a two-level staged expiration, but this more complex system still requires a "suicide bureau" to maintain revocations due to key compromise. McDaniel and Rubin [MR00] suggest that revocation will remain a necessary part of any PKI.

From a social point of view we want to acknowledge the fact that many certificates are issued by individuals, perhaps using PGP, and they are distributed without the use of a key server [Zim95]. Certificates and revocations might be posted on Web pages to publicize them, but these pages typically do not support key server responsibilities such as CRL maintenance.

1.2 A New Distributed Approach

In this paper, we propose a new method for handling distribution of revocations or certificate updates. In our method, each certificate holder has a list of *dependers*. Revocations and updates for that certificate are sent only to the dependers.

The advantage of having a set of dependers for each certificate is that it narrows the burden of notification to the minimal set of interested parties. However, a solution in which a single root entity sends revocation notices for a particular certificate to all the dependers for that certificate has several disadvantages. If the root entity is a key server with many certificates and many customers, it may be too costly to provide and distribute customized CRL's for each of its customers. On the other hand, if the root entity is an individual, it need only be responsible for sending notices regarding its own certificate, but even so may not have the resources to distribute them to a large list. For example, everyone with a copy of the PGP software has the certificate of its creator Phil Zimmerman, and he would not and could not put everyone on his depender list. Finally, it is not fault-tolerant. For example, if the network link connecting a depender to the root entity is crashed or slow, then the depender will not be able to receive the revocation notice in a timely fashion.

In our system, rather than having a centralized revocation server who sends revocations to all end users, participants who wish to be dependers for a particular certificate register as dependers with other participants. The participants can then be considered to form a *depender graph*. A participant agrees to forward any revocations or other updates she receives to her dependers. The source of a revocation notice sends it to dependers registered directly with it; those dependers then forward the revocations to their dependers, and so on.

The simplest kind of depender graph is a tree. For example, we could make a rule saying that anyone who relays a certificate should put the recipient on a depender list. That is, if A sends a certificate to B , then A puts B on A 's depender list for that certificate, regardless of who owns the certificate or where it came from. However, this simple scheme has the difficulty that it depends on the correct and prompt operation of participants, and that a participant who distributes a

certificate to many users will also be bound to distribute revocations to them. Furthermore, it is even more vulnerable to failures than the centralized root entity scheme since there is generally only one path by which a revocation notice can be forwarded.

In order to provide tolerance of up to $k - 1$ crashed, slow, or misbehaving participants (or the network links connecting them), we require participants to register as dependers with at least k other participants. This straightforward idea has several desirable properties:

- It is workable for individuals.
- It is “server-light,” so that massive institutional facilities are not required.
- It is decentralized.
- It is survivable in the event of typical computer and network failures.
- It supports prompt revocation, even if some of components exhibit extraordinary delays.
- It requires only a realistic workload for those using the system.
- The workload is allocated in proportion to the self-interest of users.
- It makes it practical to distribute revocation information immediately, rather than delaying for a periodic CRL publication schedule.

Although we focus on using depender graphs to distribute revocations, they can also be used to distribute frequent short-lived certificates or other kinds of certificate updates.

In order to join a depender graph, a participant needs to find k other participants to depend on. We present joining protocols that are *load-balanced*, in that no participant need have too many dependers, and *localized*, in that no global state is maintained and participants need only maintain information about a few other participants. We also give a localized protocol for restructuring the graph in the event of permanent failures.

We define depender graphs and prove their fault tolerance properties in Section 2. We present depender graph construction protocols in Section 3. In Section 4, we present algorithms to reconfigure the graph around permanent failures. We present further discussion in Section 5 and conclude in Section 6.

2 Depender Graphs

For a given certificate, we view certificate-holding participants in a network as nodes in a directed graph, called a *depender graph*, where there is an edge from A to B if B is on A 's depender list for that certificate. In that case we say that B depends on A , and that A is a parent of B . We will always construct depender graphs to be *acyclic* and *rooted*, and we say B is *below* A in a depender graph if there is a path from A to B . The root of the depender graph—usually the certificate owner or some kind of certificate server—is the source of revocation or update information about the certificate. When the root initiates a certificate revocation or update notice, it sends the notice to its dependers, called *root-dependers*. In turn, each node receiving the notice forwards it to its dependers.

In general, different certificates have different depender graphs, though these graphs may share some common subgraphs. In practice, multiple depender graphs might have significant overlap, and

some operations on them could be combined for efficiency. We do not discuss such optimizations further in this paper.

In order to avoid spurious revocations, revocation notices are typically authenticated in some way. In our setting, we propose that revocation notices of an individual’s public key should be signed by the corresponding private key; forwarded revocation notices maintain the initial signature. A correct signature then implies that the revocation notice either came from the owner of the key and should therefore be trusted, or the revocation notice came from someone who knows the private key (or who knows how to forge signatures from that key), in which case the key is by definition compromised and should be revoked. One advantage of this method is that since the key used to verify the revocation notice is the same as the key that is being revoked, a user will always be able to check the signature on revocation notices for certificates she has. The only case where this is not possible is when a public key is being revoked because the private key has been lost. In this case, we require the user to first obtain a new set of keys and use these to authenticate the revocation message; this has the disadvantage of requiring the new key to be disseminated before the old key can be revoked. However, it has the advantage that since revocations are digitally signed, and discarded if the authentication of the signature fails for any reason, malicious or arbitrary failures have the same effect as crash or omission failures, in which the wrong messages are sent.

We would like depender graphs to be fault-tolerant. Obviously, we cannot expect information from the root to be sent if the root has failed. However, the temporary or permanent failure of fewer than k non-root nodes should not prevent a revocation notice sent by the root from reaching any non-failed certificate holder in a timely fashion. To obtain fault tolerance, we consider the following k -redundancy property: a rooted directed acyclic graph is k -redundant if even after the removal of any set of $k - 1$ non-root nodes, there is a path from the root to every remaining node. (In Section 5.2, we address methods for making the root itself fault-tolerant if desired.)

We show below that the global property of k -redundancy can be achieved by ensuring a local property—that every node except for the root and its dependers has k parents in the graph; this is called the k -parent property. We refer to a rooted, directed, acyclic graph with the k -parent property as a k -rdag.

In order to prove the fault tolerance properties of k -rdags, we need some basic graph theoretic definitions, slightly modified to take into account the rooted nature of our depender graphs. A set of nodes is *root-avoiding* if it does not contain the root. A *cut set* is a root-avoiding set of nodes whose removal disconnects some remaining node from the root. Two or more paths from A to B are *pairwise internally node-disjoint* if no two of the paths have any nodes in common except A and B . In any rooted, finite, acyclic graph, it is possible to define a *rank* function on nodes such that every edge goes to a node of greater rank than the one it is from (so edges are rank-increasing). For example, the rank of a node can be the length of the longest path from the root to that node.

Theorem 1 *Let G be a k -rdag. Then G is k -redundant.*

Proof: Let G be a k -rdag and let C be a cut set of G . Note that if every cut set contains at least k nodes, then any set of $k - 1$ or fewer non-root nodes is not a cut set, so all remaining nodes are connected to the root, and the graph is k -redundant. Hence, it suffices to show that C has at least k nodes. Let x be a node that is disconnected from the root in $G - C$, and define the *neighborhood* of x to be the set of nodes y on paths from the root to x in G such that no path from y to x has a node in C . These are the nodes between C and x .

Note that since C disconnects x from the root, x is not the root or a neighbor of the root, and therefore x has k parents by assumption. If the neighborhood of x is empty, then every parent of x must be in C , and hence C has at least k nodes, and we are done. Otherwise, find a node y in

the neighborhood of x of minimum rank. By the definition of a neighborhood, y also is not the root or a neighbor of the root. Hence, by the k -parent property, y has k parents. Those parents must all be in C , for one that is not would be in the neighborhood of x and have rank less than y , a contradiction. Thus, C has at least k nodes, completing the proof. ■

The following more explicit result will be helpful when we consider the efficiency of revocation distribution.

Theorem 2 *Every k -rdag has k pairwise interior node-disjoint paths from the root to any node.*

Proof: Let G be a k -rdag. If the root and a root-neighbor are both active, then there is always a path between them (consisting of the single edge that connects them). Suppose x is not the root of G , and is not a root-neighbor in G . Then by the argument in the proof of Theorem 1, it follows that any cut set that disconnects x from the root is of size at least k . By Menger's Theorem (cf. [Har69]), it further follows that there are k pairwise interior node-disjoint paths from the root to x . ■

3 Depender Graph Construction

Depender graphs grow as new nodes join the graph. We envision that a new node will join the graph for a particular certificate when it receives the certificate from one of the nodes already in the graph. In order to maintain the k -parent property, the joining node must either depend on the root or find k nodes to depend on that are already in the graph.

3.1 Necessary and Sufficient Conditions

We first address the conditions necessary to ensure that there are always enough available parents without overloading participants with too many dependers. Hence, a restriction on the choice of parents is that there is limit on the number of *depender slots*, the maximum number of dependers a node is willing to support. It is clear that if nodes are not willing to have enough depender slots, then it will not always be possible to add new nodes to the graph, since once the root's depender slots are full, each new node requires k available parent slots in order to join the graph. We can show that it is enough for each new node to have k depender slots. Define a *kernel* as k nodes that have at least 1, 2, ..., k slots available, respectively.

Theorem 3 *A k -rdag can be constructed from any number of nodes with k depender slots.*

Proof: Begin with the root and make the next k nodes root-dependers. Subsequent nodes need to find k parents. We claim that when a kernel exists, another node with k depender slots can always be added to the graph, and there will still be a kernel; that is, the existence of a kernel is an invariant.

Note first that just after the k root-dependers are added, the k root-dependers each have all their k slots available, more than satisfying the requirement for a kernel. (In fact, the root-dependers form a kernel even if the i th root-depender has only i slots.)

For the proof of invariance, assume that a kernel exists. We can add a new node and give it k parents by taking one parent from each of the kernel nodes. This preserves the existence of a kernel, since the original kernel nodes now have at least 0, 1, ..., $k - 1$ slots available and the new node can be added to the kernel with its k available slots. ■




Figure 1: The $k = 3$ Triangular Scheme

The kernel-based algorithm for adding nodes to a depender graph used in the proof above is called a *triangular* scheme. The result of adding eight nodes to a root using such a scheme is illustrated in Figure 3.1 for $k = 3$. To emphasize the regular construction of the graph, the root dependers are shown with additional root-depender parents, though those edges are not necessary.

Note that a kernel may not be unique, and there may exist other nodes with additional available slots, because some nodes, such as those designed to be key servers, may support more than the minimum assumed k dependers for each certificate.

The triangular scheme always has $1 + 2 + \dots + k = (k^2 + k)/2$ slots available once all the root-dependers have been added. This may sound excessive, since adding a node only requires finding k slots (in different parents), but we can show that this number $(k^2 + k)/2$ is minimal.

Theorem 4 *In order to add k non-root-depender nodes, a k -rdag must have at least $(k^2 + k)/2$ slots available.*

Proof: Consider adding a new set S of k nodes. The first node in S to be added must depend on k other nodes. So there must be at least one slot open in k other nodes at the beginning of the process of adding the S nodes. Also, by the end of adding all nodes in S , k^2 slots have been used. Each of the k additions need to depend on k nodes, some of which may be in S . The maximum number of slots that may be used in the set S (with members of S depending on earlier members of S) is $(k^2 - k)/2$. Since k^2 total slots are used in adding S , that means there must have been at least $k^2 - (k^2 - k)/2 = (k^2 + k)/2$ slots at the beginning of the process of adding S . ■

Hence, the triangular scheme is optimal in the sense of having the fewest sustainable number of available slots. Note, though, that there are other ways of achieving the same optimal number of available slots if some nodes are willing to support more than k dependers.

3.2 A Localized Protocol for Node Addition

One motivation of forwarding certificates and recording dependers for later revocation is that it is distributed and decentralized character, so that it is not necessary for the root to communicate with all the nodes holding its certificate. Adding nodes with a triangular scheme seems to destroy this advantage by requiring participants to keep track of which nodes are in the current kernel. However, it is not necessary to do so, because it turns out the existence of a kernel can be maintained without knowing where it is.

Specifically, if there is a kernel and the parents of a new node are taken to be *any k nodes with available slots*, a kernel exists after the addition of the node. To see this, note that where kernel nodes are taken, an argument as in the proof of Theorem 3 shows that the new node plus all but one node from the old kernel form a new kernel. Where a non-kernel node is taken, the kernel node that “should” have been taken is still available to fill its role in the new kernel. Hence, the existence of a kernel is preserved. This flexibility in choosing parents makes it possible to consider optimization goals, such as minimizing the average path length in the depender graph.

Theorem 5 shows that if there is a kernel, then one can find k available depender slots in k distinct nodes by tracing down in the graph from any initial “search set” of k nodes.

Theorem 5 *If G is a k -rdag, then there is an available parent set below any set of k nodes.*

Proof: Let S be a set of k nodes. Induct on the maximum length (counting the number of edges) of a path that begins in S and ends outside S . If the maximum is 0 then the S nodes have no dependers outside S , so each node in S can have at most $k - 1$ dependers (all the other nodes in S), each node in S has at least one available slot, and thus S can be the parent set.

For the induction step, suppose the maximum such path length is n . If every node in S has an available slot, the k nodes in S can serve as parents. Otherwise, some node has no available slots, so it has a set S' of k dependers. The set S' has maximum path length smaller than n , and is below the original set S , so by induction and transitivity of below there exists an available parent set below the given k nodes. ■

Theorem 5 suggests a localized protocol for adding new nodes, for which each node in the graph keeps track only of its parents and its dependers. Given a new node, we begin by identifying a single node already in the graph as a “starting node”; typically, the starting node would be a participant from whom a new participant has just learned a certificate. If the starting node does not have k parents, it must be the root or a root-depender. In that case, either the new node can be a root-depender, or if there are already k root-dependers, take those k nodes as a search set and apply Theorem 5. Otherwise, the starting node has k parents that can be taken as a search set.

It might be desired to choose parents in such a way that the path lengths from the root to each new node are minimized. The construction in Theorem 5 does not satisfy that property. To minimize path length, one would instead traverse back up the parent links and take depender slots from the highest available nodes. However, this would either require nodes to maintain more information about where in the graph the available slots are, or would require a new participant to traverse more of the graph in the worst case.

4 Reconfiguring After Failures

When a node wants to drop out of a depender graph, or is otherwise discovered (somehow) to have failed permanently, we would like to be able to restructure the depender graph so that the

k -redundancy property is maintained on the new graph. If such reconfigurations are done, the fault tolerance of our system over time can be much more than k , as long as there are not more than $k - 1$ failed nodes between reconfigurations. We sketch a protocol for reconfiguring the graph if only crash failures can occur. If malicious failures can occur, the reconfiguration protocol needs to be made robust in order to tolerate them.

A node’s role as a depender and as a parent for its dependers can be taken over completely by one of the following:

- the last node added, which has no dependers,
- the next node added, if it is feasible to wait,
- one of its dependers (whose role will have to be taken over recursively),
- k slots (found by the protocols of Section 3.2) using one of its parents or dependers as a starting node.

There are several details that have to be addressed. If the protocol to determine which node will replace the failed node is decentralized, then there is a problem with asynchrony. Without a global ordering, if two nodes try to take the place of a failed node, each might start replacing the failed node in its parent’s dependers’ slots. Then both new nodes might get half way through, having only $k/2$ parents. In theory, the dependers of the failed node could run a mutual exclusion or priority algorithm. In practice, some canonical ordering such as one based on IP address could be used.

In order to carry out a replacement, it is necessary that the topological information stored in the failed node (its parent and depender addresses) has not been lost. It could be saved in another “caretaker” node, for example the first depender of that node in some canonical ordering (such as IP address in a practical setting). However, it cannot be saved indefinitely, because then it would become part of the topological information of the caretaker node, which would have to be saved in another node, leading to an unacceptable regress of accumulating redundant information. If we assume that at most ℓ permanent failure occurs between reconfigurations, it is sufficient for each node to store ℓ levels of topological information.

A topic for further research is to repair known failures gradually. A failure is “known” if a node becomes aware that one of its parents or dependers is no longer active. The unsolved question here is how to use local information to find and take advantage of available slots.

5 Discussion

In this section, we briefly discuss a number of issues and possible extensions where further research is called for.

5.1 Link/Transport connectivity

If two paths are node-disjoint, then they are also edge-disjoint. Thus, our depender graphs are tolerant against the failure of $k - 1$ node or edge failures. However, in a real network, links between different nodes are not independent. Often many links go through the same switching node in an underlying communication infrastructure. Thus, the failure of one switching node may result in the failure of many edges in a depender graph.

It would therefore be desirable to assure that links from a node to its k parents are independent (so it takes k failures of lower-layer switching nodes to break them all). If in addition there are

k independent paths from the root to its dependers, then an inductive argument shows that it takes k failures of the underlying network components to cut all paths to a node. A weaker version guarantees k -redundancy for non-root-depender nodes so long as each link from the root to a root-depender is independent of all other links in the graph. We can still show by induction that it takes k failures to cut off a non-root-depender.

Checking independence of transport paths can be done using network monitoring tools such as “traceroute.” However, in practice this information is rather dynamic and may be difficult to keep a handle on.

5.2 Distributing root authority

In some settings, it is desirable for the root authority to be distributed among multiple parties, so that it takes at least t of these parties to send out a valid revocation notice. This can be achieved by distributing the functionality of the root into multiple parties and using threshold signatures [DF91, GJKR96] so that the correct participation of t parties is necessary and sufficient to create a valid revocation. If this new “distributed root” consists of at least $k + t - 1$ parties, then this also provides crash fault tolerance for up to $k - 1$ of the root parties.

In order for the threshold signatures to work, the root-dependers must now have at least $k + t - 1$ the root parties as parents; other nodes still need k parents as before. When a revocation notice is sent, it is signed using the threshold signature scheme. Each root node sends its partial signature to the root-dependers. The root-dependers reconstruct the signed revocation notice, and if it is a valid signature, they proceed as before by forwarding the signed revocation. Since the resulting depender graph has its normal properties with respect to this distributed root, it still enjoys the k -redundancy property with respect to it.

5.3 Global Optimizations

For distribution of revocation notices, the k -redundancy property can be exploited simply by having each node forward a notice to all its dependers. In the general case, this is the best that can be done. However, there are several situations in which global information about the graph could be used to reduce or eliminate unnecessary network traffic while still ensuring revocations are distributed properly.

For example, if the graph in fact more than k disjoint paths to some nodes, it might be possible to remove or ignore some of the edges of the graph. Similarly, if not all nodes need to receive each update, then some edges can be removed. Given a particular destination node, Theorem 2 says that there are k pairwise interior node-disjoint paths from the root to that node, so that using only the edges in these paths would eliminate unnecessary traffic while preserving k -redundancy with respect to that one destination node. When only some subset of nodes needs to receive a revocation notice, the goal would be to find a minimal set of edges that include k disjoint paths to each node in the subset. Finally, in the case that something more is known about which failure configurations can occur than just that any $k - 1$ nodes might simultaneously fail, it might be possible to ensure that each node has always at least one path from the root through no failed nodes without having k disjoint paths to each node.

To go one step further, depender graphs for multiple certificates could take advantage of structure sharing, so that where participants have certificates in common, messages relating to those certificates could be coalesced. This could be done not just for revocation notices, but for the depender graph construction protocol itself. Structure sharing could be further facilitated by protocols to merge multiple depender graphs with common nodes into one graph with the maximal

number of common edges.

6 Conclusions

We have presented depender graphs, which provide a locally manageable, scalable, efficient, and fault-tolerant method of certificate revocation in a public-key infrastructure.

Due to their fault tolerance and localized construction protocols, k -rdags may find useful applications elsewhere. As described in this paper, they are most useful for environments in which only crash or delay failures occur, or if the information to be sent is digitally signed or otherwise verifiable, as in the case of certificate revocations. However, even in the case of malicious failures during the distribution of information, k -rdags can tolerate up to $(k - 1)/2$ failures by using voting.

Other possible applications that might benefit from depender graphs include fault-tolerant multicast backbone (MBone) trees, maintaining location information for a mobile host as it moves from one base station to another, and distributing routing information in the Internet such as reachability information exchanged by the BGP protocol.

Acknowledgments

We thank Patrick McDaniel for helpful discussions.

References

- [AZ98] C. Adams and R. Zuccherato, “Internet X.509 Public Key Infrastructure Data Certification Server Protocols,” Internet Draft, PKIX Working Group, 1998.
- [Coo99] D. Cooper, “A Model of Certificate Revocation,” *Proc. 15th Annual Computer Security Applications Conference*, 1999, 256–264.
- [DF91] Y. Desmedt and Y. Frankel, “Shared generation of authenticators and signatures,” In *Advances in Cryptology—CRYPTO ’91, Lecture Notes in Computer Science* 576, 457–469, Springer-Verlag, 1992.
- [FL98] B. Fox and B. LaMacchia, “Certificate Revocation: Mechanics and Meaning,” *Proc. Financial Cryptography ’98*, LNCS 1465, 1998, 158–164.
- [GJKR96] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin, “Robust Threshold DSS Signatures,” In *Advances in Cryptology—CRYPTO ’96, Lecture Notes in Computer Science* 1070, 354–371, Springer-Verlag, 1996.
- [Har69] F. Harary, *Graph Theory*, Addison-Wesley, Reading, MA, 1969.
- [KAN99] H. Kikuchi, K. Abe, and S. Nakanishi, “Performance Evaluation of Certificate Revocation Using k -Valued Hash Tree,” *Proc. ISW’99*, LNCS 1729, 1999, 103–117.
- [Koc98] P. Kocher, “On Certificate Revocation and Validation,” *Proc. Financial Cryptography ’98*, LNCS 1465, 1998, 172–177.
- [MJ00] P. McDaniel and S. Jamin, “Windowed Certificate Revocation,” *Proc. IEEE Infocom 2000*, IEEE, 2000, 1406–1414.

- [MR00] P. McDaniel and A. Rubin, “A Response to ‘Can We Eliminate Certificate Revocation Lists?’ ”, *Proc. Financial Cryptography 2000*, to appear.
- [Mey98] M. Myers, “Revocation: Options and Challenges,” *Proc. Financial Cryptography '98*, LNCS 1465, 1998, 165-171.
- [NN98] M. Naor and K. Nissim, “Certificate Revocation and Certificate Update,” *Proc. 7th USENIX Security Symposium*, 1998, 217–228.
- [Riv98] R. Rivest, “Can we eliminate certificate revocation lists?” *Proc. Financial Cryptography '98*, LNCS 1465, 1998, 178-183.
- [X.509] “The Directory-Authentication Framework,” CCITT Recommendation X.509.
- [Zim95] P. Zimmermann, *The Official PGP User's Guide*, MIT Press, 1995.