

# Generalized Certificate Revocation

Carl A. Gunter and Trevor Jim

July 19, 1999

## Abstract

We introduce a language for creating and manipulating *certificates*, that is, digitally signed data based on public key cryptography, and a system for *revoking* certificates. Our approach provides a uniform mechanism for secure distribution of public key bindings, authorizations, and revocation information. An external language for the description of these and other forms of data is compiled into an intermediate language with a well-defined denotational and operational semantics. The internal language is used to carry out consistency checks for security and optimizations for efficiency. Our primary contribution is a technique for treating revocation data *dually* to other sorts of information using a polarity discipline in the intermediate language.

## 1 Introduction

Public Key Infrastructures (PKI's) have received considerable attention in the last decade in the hope that they can form a foundation for secure electronic commerce. We have developed a programming language, QCM, that can be used as the basis of a general and semantically sound PKI [9]. In this paper we extend QCM to support *certificate revocation*.

A certificate is a digitally signed document. PKI's use certificates to securely distribute data including key bindings ('Alice's key is  $K$ ') and authorizations ('Alice has permission to use the printer'). Because the data in a certificate can become out of date, most PKI's support a mechanism for revoking certificates called the Certificate Revocation List, or CRL. However, CRL's are controversial: some researchers have pointed out that their semantics is not well understood [8, 19], and others have proposed eliminating them entirely [17].

We believe that part of the confusion regarding CRL's and PKI's stems from treating revocation data specially. That is, rather than treating revocation data as just another kind of data to be distributed, PKI's historically have had ad hoc mechanisms for managing CRL's, certificates containing key bindings, and certificates containing authorizations. Consequently, the semantics and distribution of the different kinds of data must be understood separately, rather than as instances of a general concept.

In this paper, we propose an alternative, general framework that can be used to uniformly and securely distribute data of all sorts, including public key bindings, authorizations, and revocation information. The work builds on our Query Certificate Manager (QCM) system [9], a programming language that already treats public key bindings and

authorizations uniformly. The key observation we make is that data used for revocation should be treated *dually* to other sorts of information: for example, to determine that Alice's key is  $K$ , it is necessary to determine that (Alice,  $K$ ) is one of the bindings issued by a certificate authority, and *is not* a binding revoked by the authority. By viewing the two kinds of data dually, and developing a *polarity discipline* for QCM programs that distinguishes them, we are able to use the same mechanisms to distribute both kinds of data securely.

The main accomplishment of the paper is to define a general and semantically sound PKI that supports revocation and that can be feasibly implemented. This is demonstrated using a denotational model to define correctness, an operational semantics that is a simplified version of our distributed implementation, and a polarity discipline that ensures that the operational semantics is sound with respect to the model. From a programming language perspective, some of the novelties of the system are that QCM is designed to work with partial information—QCM does not calculate the exact answer for a program, only a best approximation—and that in order to handle revocation, the QCM implementation manipulates representations of infinite sets.

From a security perspective, our contributions are to provide a semantics of revocation, and to show that revocation information can be distributed by the same mechanisms used to deliver key binding and authorization data. A MITRE study to recommend a PKI for the U.S. Government noted:

Certificate revocation list distribution is by far the biggest cost driver associated with the operation of the PKI. Requiring that every request to the directory service for a certificate be accompanied by a similar request for the CRL on which that certificate may appear places an extremely heavy burden on the directory communications system. . . . Other ways of dealing with the CRL's must be considered. ([1] at 7–4)

The paper is organized as follows. In the next section we explain why certificate revocation is both logically and pragmatically problematic, in the context of related work. In Section 3, we introduce an extension of QCM with non-membership tests. Non-membership tests are a restricted form of negation that enable CRL's to be programmed directly in QCM. In Section 4, we present our polarity discipline and show that polarities can be used to make QCM with revocation monotonic. Section 5, we give an operational semantics for QCM with revocation and show that it is sound, using monotonicity. Section 6 concludes. Proofs are sketched in an Appendix.

## 2 Related Work

In this paper we will refer to any data with a digital signature as a *certificate*. Certificates are tamper-evident (modifying the data makes the signature invalid) and unforgeable (only the holder of the secret, signing key can produce the signature). These properties make certificates useful in conducting secure electronic transactions.

Many different kinds of data can be signed. For example, a certificate may represent a binding (e.g., '*p* is the public key of Alice'), or it may indicate a permission (e.g., 'Alice has permission to use the swimming pool'). In any case, the signer of the certificate, who is known as the *issuing party*, may wish to indicate a term of validity for the certificate. For instance, the certificate giving Alice permission to use the pool could be marked valid until the end of the academic year. The *relying party* who examines a certificate must take this validity period into account: when deciding whether to admit Alice to the swimming pool, the relying party should consider the certificate invalid if the academic year is over, and deny access.

*Revocation* is used to invalidate a certificate prematurely, before the end of its validity period. Revocation might be needed due to key compromise (e.g., the signing key is stolen), change of affiliation (e.g., Alice drops out of school), or many other reasons.

The best known PKI is the ISO Directory, based on the X.500 series of standards [11]. One of these standards, X.509 [12], proposes a hierarchy of *Certificate Authorities (CA's)* which produce *public key certificates* binding principals (like people, devices, or entities) to public keys. In recognition of the practicalities of large systems, X.509 provides a mechanism for dealing with certificates that become too old and must therefore be considered expired, as well as certificates that, for some reason, need to be revoked ahead of their expiration date. The latter are announced on *Certificate Revocation Lists (CRL's)* which are signed by certificate authorities (CRL's are certificates too).

A typical scenario under the X.509 standard would proceed as follows.

- The issuing party provides a certificate with a period during which it is to be considered valid. A serial number included in the certificate uniquely identifies it.
- The issuing party provides a CRL, which is a certificate that contains serial numbers of revoked certificates. The CRL also has a time issued and a time for the next CRL to be issued, so it is possible to determine whether the CRL is current.
- To validate a certificate, the relying party checks the correctness of its signature, checks that the certificate has not expired, and checks to see if the serial number of the certificate is listed on the current CRL of the issuing party. The signature of the CRL itself must also be checked.

The most recent X.509 standard provides for some variations on this protocol, such as *indirect CRL's* that are signed not by the issuing party, but by a third party designated by the issuing party.

*Attribute certificates* are yet a third kind of certificate, used not to bind keys to identities, but rather to use keys for a range of objectives requiring authentication and/or secrecy. For instance, the IETF Working Group, PKIX [20, 10], focuses on developing a system based on CA's that can

support security for network directory services (like DNS), electronic mail (*viz.* extensions of PEM [14]), and web pages. Indeed, the original aim of the certificate authorities was to support *authorization* for the Directory [11, 12]. A directory system like the Lightweight Directory Access Protocol (LDAP) [21] can use the certificate authorities to support control over which principals can access or alter the Directory. Attribute certificates can be used to bind a subject to a set of permissions, thus supporting authorization. How a system should manage attribute certificates is not well understood. For instance, Warwick Ford, the co-chairman of PKIX, notes in [7]:

Attribute certificates represent an important area of electronic commerce technology which is yet to be fully explored or developed. (page 253)

Certified distribution of authorization information is an active development area and... the 'best' approach is far from clear. (page 256)

Support for this critical function has been the subject of recent research on *policy verification engines* that support the use of certificates representing authorization and identity information to state and verify security policies [2, 5, 6, 9]. Although the X.509 system supports revocation for attribute certificates just as it does for public key certificates, none of these new verification engines directly provide such support.

Given this confusion surrounding certificate management, it is not surprising that some have proposed to simplify PKI's by eliminating CRL's completely [17].

The problems with CRL's fall into essentially three categories: (1) whether another mechanism for achieving similar goals would work better; (2) the semantics of CRL's; and (3) the means for distributing CRL's.

### 2.1 Can We Do Without CRL's?

CRL's are less likely to be needed if certificates have short validity periods. For example, if passes to the swimming pool were issued once each semester rather than once each academic year, then on average only half as many entries would need to appear on the CRL. If Alice's swimming privileges were revoked during the first semester, then the serial number for her certificate would only need to appear in the CRL until the end of the first semester, and not until the end of the year. After the first semester, the certificate would be rejected because of expiration anyway, so it does not need to be in the CRL.

Carrying this to an extreme, if Alice was required to obtain a fresh certificate with a very short expiration period each time she went to the swimming pool, then the issuer could revoke her privileges very quickly simply by refusing to issue new certificates. This is the approach taken by OCSP [16], a proposed alternative to X.509 CRL's. The drawback to this approach is that it requires many more certificates than the X.509 approach, potentially placing an unacceptable burden on certificate servers. The efficiency of the CRL approach is based on an assumption that the number of serial numbers in the CRL will be small compared to the number of valid unrevoked certificates, and that a reasonable interval of validity for CRL's will be acceptable to the stakeholders in the protocol. To see this issue, suppose that the CRL for the swimming pool is issued every 24 hours. Then Alice may still get in a long swim between the time her certificate was revoked and the issuance of the CRL with the

serial number of her certificate. This is called the *window of vulnerability*, and its acceptable size will undoubtedly depend on the resources at stake. Technical efficiency issues concerning the size of the window have been explored as well [15]. These are closely related to the problem of CRL retrieval, which we will discuss shortly.

QCM lets the user decide whether CRL's should be used, and what the size of the window of vulnerability should be.

## 2.2 What Does Revocation Mean?

Consider a certificate,

' $p$  is the public key of Alice' (signed  $q$ ).

Suppose this certificate has serial number  $n$ , and  $n$  appears on a CRL. What does this mean? Fox and LaMacchia point out at least three possible interpretations [8]:

1.  $p$  is not the public key of Alice;
2.  $q$  can no longer vouch for whether  $p$  is the public key of Alice; or
3. the signing key of  $q$  has been compromised.

The relying party should act quite differently depending on what interpretation it takes. For example, suppose we have two additional certificates:

A. ' $p$  is the public key of Alice' (signed  $r$ ).

B. ' $r$  is the public key of Bob' (signed  $q$ ).

If we take interpretation (1), then we should discard A, but we may still continue to trust B. On the other hand, if 2 holds then we may choose to trust both A and B. Finally, if 3 holds, then we may trust A but not B.

The ISO X.509 protocol [12] recognizes this subtlety and provides for a list of reasons that can be attached to CRL entries. Reasons include *superseded*, *cessation of operation*, and *CA compromise*. However, nothing is said about what one is expected to do when any given reason is encountered. The Internet X.509 protocol [10] adds an additional reason with a more operational flavor called *certificate hold* which might have a significance like telling the relying party to take Alice's certificate away from her when she next attempts to use the swimming pool. Fox and LaMacchia in [8] note that: "Until we can come up with a better mechanism for moving revocation information around, and binding that revocation information to the proper statement undone by the revocation, use of revocation information can add significant ambiguity to the chain-building process." If it adds ambiguity even to chains of public key certificates, then it may be even more ambiguous when applied to attribute certificates generally.

To avoid these confusions, we believe a PKI with revocation needs a formal semantics. The only other work we know of along these lines is Stubblebine [19], who extends BAN logic with time intervals and a form of revocation. Stubblebine does not give a semantic model or implementation, however. A semantic model would be helpful in avoiding security breaches. For example, the following scenario was possible in SDSI 1.1 [18]:

$$\begin{aligned} school &= teachers \cup admin \cup students, \\ employees &= school - students. \end{aligned}$$

This defines two groups, *school* and *employees*, in terms of some other groups, including *students*. Alice could be issued a student id, that is, a certificate 'Alice  $\in$  *students*', good for the entire school year. If Alice drops out during the year, the administration could revoke her privileges by issuing a new certificate, 'Alice  $\notin$  *students*.' Now, if Alice gets her hands on both certificates, she can prove that she is an employee: since Alice  $\in$  *students*, by the first definition we have Alice  $\in$  *school*. And then since Alice  $\notin$  *students*, we have Alice  $\in$  *employees*, even though Alice was never a teacher or administrator. The problem, of course, is that the two certificates are contradictory: they do not have a model. We have used our semantics as a guide to design restrictions into QCM that guarantee that such contradictions can never occur, while still permitting revocation to be expressed.

## 2.3 Distributing CRL's

In the first version of X.509, CRL's were provided monolithically via a list of serial numbers for a given CA.<sup>1</sup> However, if the number of revoked certificates became large and the list needed to be refreshed frequently to keep down the window of vulnerability, then significant burdens could be placed on the CRL server and the network. Much of the work on CRL's has been aimed at addressing this problem, and we sketch only enough of the story to try to convince the reader that the distribution and retrieval of CRL's is a non-trivial issue and remains an open problem. We believe this is also true of attribute certificates, but will not attempt to argue it here.

A first observation is that CRL's can be divided according to who might be interested in them. For instance, a single CRL *distribution point* could be partitioned into a family of CRL's representing different reasons for revocation, possibly issued at different rates. This might allow a short list of compromised keys to be issued with a short validity period while a longer list of some lower-priority revocations is updated less frequently. It also facilitates 'on demand' retrieval of the CRL. For instance, only the CRL for people with names in a certain range may need to be retrieved to check a given certificate.

A second observation is that most of the entries in a CRL remain the same as CRL's are updated, so it is possible to issue a *delta CRL* indicating only the changes. Of course, this means that checking a certificate entails having all of the deltas from the last full release. This also opens the possibility of having relying parties choose different strategies for validation. For instance, a low-risk decision might not require checking the delta CRL's.

A third observation is that multiple CA's could use a single CRL distribution point, maintained by a third party. These are the indirect CRL's we mentioned earlier.

As a final observation, there are a number of ways to obtain CRL's. While a relying party may choose to query a distribution point, it is equally possible for a distribution point to 'push' a CRL to a collection of potential relying parties. If properly done, this could enable efficient distribution via unreliable network transport using UDP unicast or multicast [13, 15].

All of these observations on the distribution of CRL's apply equally to the distribution of authorizations (ACL's) and key bindings (public key directories). However, in the

<sup>1</sup>To be more precise, there were two such lists, one for certificates for other CA's and one for the certificates of end users.

PKI's that we are aware of, CRL's are always treated specially, with their own distribution mechanism. In QCM, the mechanisms for distributing revocation data are the same as those for distributing other kinds of data. In this way, when we implement an improved distribution mechanism, it applies to all kinds of data immediately.

### 3 QCM with Revocation

QCM is a programming language for securely specifying and evaluating distributed tables; its syntax is based on a calculus of set comprehensions [3, 4]. QCM programs can express *security policies* much like those of PolicyMaker [2], as well as the *groups* of SDSI [18]. QCM improves on SDSI and PolicyMaker by supporting automatic retrieval of remotely-stored certificates as part of group or policy enquiries; SDSI and PolicyMaker leave certificate retrieval to a separate, as-yet-undefined mechanism. Conversely, while systems like X.509's Directory [11] and LDAP [21] provide certificate distribution, they are not integrated with verification (certificate use) as in QCM.

Users of QCM write programs (security policies) in a high level language that we call the *external* language. Our implementation translates programs in the external language into an *internal* language that facilitates query decomposition and optimization. In this section, we will introduce the internal language and define its denotational semantics. We focus on the features we need for key compromise and revocation, omitting some of the other mechanisms of the language to simplify the presentation. These omitted features are described in other papers [9, 13].

#### 3.1 Introduction to QCM

We will introduce QCM by example. Suppose a research group  $L$  is collaborating with a group  $R$  at a remote site and wishes to maintain an appropriate Access Control List for the resources at  $L$  to be used in the project.  $L$  can define a set, named ACL, of the permitted users with a QCM definition:

$$\text{ACL} = \text{LocalUsers} \cup K_R\$ACL. \quad (\dagger)$$

Here  $K_R$  is the public key for the group  $R$ . The notation  $K_R\$ACL$  is pronounced, " $K_R$ 's ACL," and it is the global name of a set ACL defined by  $R$ .  $K_R$ 's ACL is distinct from the ACL defined by  $(\dagger)$ , which is known globally as  $K_L$ 's ACL. By qualifying names with keys, QCM ensures that the sets defined by different principals will not be confused. Moreover, the QCM implementation will discard any information regarding  $K_R\$ACL$  unless it comes with a signature verified by  $K_R$ ; this means that only  $R$ , the holder of the secret, signing key, can convince QCM that a user is in  $K_R\$ACL$ .

The definition  $(\dagger)$  says that  $K_L$ 's ACL is the union of a set, LocalUsers, and the set  $K_R\$ACL$ . LocalUsers is defined separately by  $L$ , for example, its members can be listed explicitly:

$$\text{LocalUsers} = \{K_L\}.$$

After  $L$  has made these definitions, a QCM evaluator on the local machine can be queried with set expressions involving ACL and LocalUsers. For example, if QCM were asked the query 'ACL?' it would eventually return a set  $\{K_L, \dots\}$ . Exactly how it does this is largely hidden from the user. LocalUsers is easy to obtain, of course, but to obtain  $K_R\$ACL$

QCM might use a variety of different strategies. The most straightforward would be to send a message to a QCM evaluator at  $R$ 's site. An optimization would be to cache the response. An extension of this optimization is to mirror the set  $K_R\$ACL$  locally at  $L$ . This option breaks into two possibilities: one in which  $R$  'pushes' the ACL (or a delta of it) whenever it changes, the second in which  $L$  'pulls' the ACL (if it has changed) whenever it needs to use it. In each case, it is essential to secure the integrity of the communications, so QCM signs messages and verifies signatures when appropriate. QCM automatically and seamlessly supports all of these mechanisms, as well as other commonly-used mechanisms like online versus offline signing [9, 13].

Usually, the entire ACL is not needed; a more typical query would ask whether  $K_{\text{Alice}}$  was a member of the ACL. To keep the language simple, QCM does not have such Boolean queries, but they can be encoded as set queries. For example to see if  $K_{\text{Alice}}$  is a member of the ACL, it is sufficient to ask the query

$$\{\text{"yes"} \mid x \in \text{ACL}, x = K_{\text{Alice}}\}. \quad (\ddagger)$$

This will evaluate to the set  $\{\text{"yes"}\}$  if  $K_{\text{Alice}} \in \text{ACL}$ , and if  $K_{\text{Alice}} \notin \text{ACL}$ , it will evaluate to the empty set,  $\{\}$ .

We have not yet discussed how QCM supports certificates. In QCM a certificate is a signed statement about names, for example:

$$'K_{\text{Alice}} \in \text{ACL}' \text{ (signed } K_R) \quad (\S)$$

This certificate is a statement by  $K_R$  attesting that  $K_{\text{Alice}}$  is in its ACL. When QCM receives  $(\S)$ , it will believe that  $K_{\text{Alice}} \in K_R\$ACL$  (after verifying the signature, of course). So, if QCM is asked the query  $(\ddagger)$  and is given the certificate  $(\S)$  at the same time, it can evaluate the query to  $\{\text{"yes"}\}$  without sending any messages to  $R$ .

If QCM is asked the query 'ACL?' and is given  $(\S)$  at the same time, it behaves in exactly the same way: it believes that  $K_{\text{Alice}} \in K_R\$ACL$ , and it evaluates the query *without sending any messages to  $R$* . The answer returned is  $\{K_L, K_{\text{Alice}}\}$ , regardless of whether there are any additional keys in  $K_R\$ACL$ . This illustrates an important point: the answer returned by QCM is not guaranteed to be an exact answer. However, the answer returned is always a lower bound (subset) of the 'real' answer. This is guaranteed by careful restrictions built into the language.

This is a sensible choice, for two reasons. First, it is conservative: some people who should be approved may be denied access, but no one who should be denied access will be approved. Second, there are situations in which it is appropriate for QCM to refuse to send any messages. For example, we might need to guard against an adversary who submits queries so as to make QCM send many messages. QCM can defend against such an attack while still functioning by operating in a mode where messages are not sent but certificates are accepted.

However, using lower bounds presents difficulties for revocation. The most obvious way to add revocation to QCM would be to add a non-membership test to the language. For example, suppose we define an access control list that assigns a unique serial number to each user on the list:

$$\text{ACL} = \{(K_L, 1), (K_R, 2), \dots\}$$

If we separately define CRL to be the set of serial numbers of users who should be revoked, then the query

$$\{x \mid (x, n) \in \text{ACL}, n \notin \text{CRL}\}$$

Table 1: QCM’s internal language

Constants	$c$	$\in$	Key $\cup$ Num $\cup$ Str $\cup$ Bool
Positive variables	$x^+$	$\in$	Var <sup>+</sup>
Negative variables	$x^-$	$\in$	Var <sup>-</sup>
Positive names	$u^+$	$\in$	Name <sup>+</sup>
Negative names	$u^-$	$\in$	Name <sup>-</sup>
Variables	$x$	$::=$	$x^+ \mid x^-$
Names	$u$	$::=$	$u^+ \mid u^-$
Polarities	$\gamma$	$::=$	$+ \mid -$
Expressions	$e$	$::=$	$x \mid c \mid$ $e\$u \mid$ $\{e, \dots, e\} \mid$ $(e, \dots, e) \mid$ $\bigcup e \mid$ $\{e \mid g, \dots, g\} \mid$ $(e@e) \mid$ $\text{cml}\{w, \dots, w\}$
Qualified names			$e\$u$
Enumerated sets			$\{e, \dots, e\}$
Tuples			$(e, \dots, e)$
Set union			$\bigcup e$
Comprehensions			$\{e \mid g, \dots, g\}$
Evaluate at			$(e@e)$
Co-finite sets			$\text{cml}\{w, \dots, w\}$
Values	$v$	$::=$	$c \mid (v, \dots, v) \mid$ $\{v, \dots, v\} \mid$ $\text{cml}\{w, \dots, w\}$
Comparable values	$w$	$::=$	$c \mid (w, \dots, w)$
	$g$	$::=$	
Generators			$p \in e \mid$
Guards			$e = e \mid$ $e \neq e \mid$ $e \notin e$
Patterns	$p$	$::=$	$x \mid (x, \dots, x)$
Definitions	$d$	$::=$	$u = e$
Programs	$P$	$::=$	$d_1, \dots, d_n$

is the set of users who should have access. If we wish to calculate a lower bound for the query, we need a lower bound for ACL, but we need an *upper* bound for CRL. Fortunately, these are dual notions, so most of the existing machinery of QCM can be adapted to calculate upper bounds. The main additions required are certificates that give upper bounds ( $K_{\text{Alice}} \notin \text{CRL}$ ) and the polarity discipline of Section 4 that will guarantee that we do not run into the semantic inconsistencies illustrated in Section 2.2.

### 3.2 The Internal Language

A grammar for the internal language is given in Table 1. It includes the usual constants (numbers, booleans, and strings) as well as keys. Names  $u$  and local variables  $x$  are tagged with positive or negative *polarities*,  $\gamma$ . These annotations are used in the polarity discipline given in the next section; they can be inferred for variables, but they are required for names.

We have already introduced most of the expressions  $e$  of the language. Most of them correspond to the usual set operations. Patterns are restricted so that a variable may not occur twice. Thus,  $(x, y)$  is a well-formed pattern, but  $(x, x)$  is not. Notice that every pattern is an expression. Patterns and tuples must contain at least two components, so an expression like  $(x)$  is not well-formed. On the other hand, an enumerated set may have one or no elements. An expression of the form  $\{e\}$  is a singleton, and  $\{\}$  is the empty set.

We use boldface metavariables for sequences, e.g.,  $\mathbf{g}$  denotes a sequence  $g_1, \dots, g_n$  of guards and generators. In the expression  $\{e_1 \mid x \in e_2, \mathbf{g}\}$ , the visible occurrence of  $x$  is a binding of that variable and binds free occurrences of  $x$  in  $e_1$  and  $\mathbf{g}$ , but not in  $e_2$ . Guards do not introduce bindings. For example, if  $e_2$  and  $e_3$  have no free variables, then the meaning of the expression  $\{x \mid x \in e_2, x \in e_3\}$  is the same as that of  $e_3$ . Expressions are considered syntactically identical modulo renaming of bound variables, reordering of enumerated sets, and elimination of duplicates in enumerated sets. Set difference is worth introducing as syntactic sugar:

$$e_1 - e_2 = \{x \mid x \in e_1, x \notin e_2\}$$

Note that  $x$  in the generator binds the occurrences of  $x$  in the guard and the range. Similar syntactic sugar, like a binary union operator, will be used without comment.

A QCM program is a set of definitions,

$$u_1 = e_1, \dots, u_n = e_n.$$

We require the definitions to be acyclic, and  $u_i \neq u_j$  if  $i \neq j$ .

QCM evaluation is based on a distributed family of QCM processes, each created from a QCM program. These processes act as servers to applications and to other QCM processes, both of which query QCM about the values of QCM expressions. QCM responds with a signed table created from the QCM processes acting as a distributed database over an ad hoc virtual private network. For brevity, these *response certificates* are omitted from this treatment. Dually, an application that has obtained one or more certificates previously may ‘push’ them to a QCM process in order to spare the process the trouble of obtaining them from other sources. These are called *push certificates*, and they take two forms:

$$‘u^+ \sqsupseteq e’ \text{ (signed } K),$$

$$‘u^- \sqsubseteq e’ \text{ (signed } K).$$

The first form expresses a lower bound for a positive name. The second form expresses an upper bound for a negative name, which might be used for revocation and key compromise. These certificates generalize the membership and non-membership certificates we have been using in our exposition.

We require the bounds to be sets, and we do not permit certificates with upper bounds for positive names, or lower bounds for negative names. Because of these restrictions, there is *always* a model for any set of push certificates; in fact, there is a best, or minimal model. For example, given the certificates

$$‘\text{ACL}^+ \sqsupseteq \{(K_L, 1), (K_R, 2)\}’ \text{ (signed } K),$$

$$‘\text{CRL}^- \sqsubseteq \text{cml}\{2\}’ \text{ (signed } K),$$

the minimal model assigns  $K\$ACL^+$  the set  $\{(K_L, 1), (K_R, 2)\}$ , and assigns  $K\$CRL^-$  the set  $\text{cml}\{2\}$ , the complement of the set  $\{2\}$ . (The condition  $\text{CRL}^- \sqsubseteq \text{cml}\{2\}$  is equivalent to  $2 \notin \text{CRL}^-$ .)

Existence of the minimal model ensures that we avoid the semantic inconsistencies of Section 2.2. It also provides a way to evaluate queries that takes advantage of any pushed certificates: we construct the minimal model and evaluate the query in the model. The crucial *monotonicity* result of the next section shows that this results in lower bounds for positive sets, and upper bounds for negative sets.

We have already seen several examples of how QCM can express CRL’s. Now we will give an example that shows how

QCM can guard against key compromise. Key compromise occurs when a private, signing key is revealed to the adversary. Once the adversary has a signing key, they masquerade as the keyholder. To protect against this, we can maintain a set of compromised keys in QCM:

$$\text{Compromised}^- = \{K_1, K_2, \dots\}$$

Now, anyone who wants to protect themselves against the compromised keys can refer to  $\text{Compromised}^-$  in writing their policies. For example,

$$\text{ACL}^+ = \{x^+ \mid K_R \notin \text{Compromised}^-, x^+ \in K_R \$ \text{ACL}^+\}.$$

Here we check that the key  $K_R$  is not in the compromised set before adding any elements of  $K_R \$ \text{ACL}^+$  to the local ACL. In general, use of a compromised key can be prevented by transforming expressions of the form  $e \$ u$  into  $\{x \mid x \in e \notin \text{Compromised}, x \in e \$ u\}$ . This is a simple transformation that can be performed during the translation from the external language to the internal language.

**Relationship to the external language** Our aim is to show how to achieve an appropriate protocol using a general calculus for certificate checking and retrieval. We do this by: (1) compiling user programs into the internal language, checking polarities in the process; and (2) using a front end to control such operations as the creation of serial numbers for certificates. A separate mechanism controls issues like which certificates are revoked. The reader may be curious whether the program needs to be changed each time one of the sets in the definitions is changed, say by adding a new serial number to the revoked set. For the treatment in this paper the answer is yes, but in practice QCM is implemented as a middleware system on top of one or more data sources like flat files, XML expressions, or LDAP databases. When the data in these underlying data sources changes, the runtime system takes this into account without the need for recompilation.

### 3.3 Denotational Semantics

The denotational semantics is given using a recursive domain equation and a collection of semantic clauses that compositionally describe the meanings of QCM expressions.

Let  $\llbracket \text{Key} \rrbracket$ ,  $\llbracket \text{Num} \rrbracket$ ,  $\llbracket \text{Str} \rrbracket$ ,  $\llbracket \text{Bool} \rrbracket$  be the semantic spaces that interpret principals, numbers, strings, and booleans respectively. The semantic universe is defined via the following domain equation, where  $\uplus$  is the disjoint union:

$$U = \llbracket \text{Key} \rrbracket \uplus \llbracket \text{Num} \rrbracket \uplus \llbracket \text{Str} \rrbracket \uplus \llbracket \text{Bool} \rrbracket \\ \uplus \mathcal{P}_{\text{fin}}(U) \uplus \mathcal{P}_{\overline{\text{fin}}}(U) \\ \uplus \Pi(U, U) \uplus \Pi(U, U, U) \uplus \dots$$

We use the notation  $\mathcal{P}_{\text{fin}}(U)$  for the finite subsets of  $U$ , and  $\mathcal{P}_{\overline{\text{fin}}}(U)$  for the co-finite subsets of  $U$ . The expression  $\Pi(U, U)$  is for pairs of  $U$ 's, while  $\Pi(U, U, U)$  is for triples, and so on.

Notice that every set in the semantic universe  $U$  is either finite or co-finite with respect to  $U$ ; for example,  $\text{Num} \notin U$ . This means that every element of  $U$  has a finite representation, which is convenient for our operational semantics.

To achieve monotonicity, it is essential to restrict comparison by equality or inequality to elements that are in a distinguished subset,  $C$ , of the universe; we define  $C$  to be the least subset of  $U$  that includes the elements of base type (principals, integers, booleans, strings) and is closed under the tuple operation.

An *environment* is a function  $\rho : \text{Var} \uplus (\text{Key} \times \text{Name}) \rightarrow U$ . We write  $\rho(x)$  for the value on variables and  $\rho(K, u)$  for the value on principal / name pairs. We restrict ourselves to environments  $\rho$  such that  $\rho(K, u)$  is a set for any  $K$  and  $u$ .

The semantics of the language is a partial function on expressions  $e$  and environments  $\rho$  denoted by  $\llbracket e \rrbracket \rho \in U$ . In our definition, we use the convention that in a clause of the form

$$\llbracket e \rrbracket \rho = \dots \llbracket e' \rrbracket \rho' \dots,$$

the meaning of  $e$  relative to  $\rho$  has the value on the right hand side *provided*  $\llbracket e' \rrbracket \rho'$  is defined.

The denotational semantics of QCM is given by the semantic clauses of Table 2. Some expressions have no meaning; that is,  $\llbracket e \rrbracket \rho$  may be undefined. A subtle example of this comes up in the clauses defining the meaning of comprehensions:

$$\llbracket \{e_1 \mid x \in e_2, \mathbf{g}\} \rrbracket \rho = \bigcup_{\nu \in \llbracket e_2 \rrbracket \rho} \llbracket \{e_1 \mid \mathbf{g}\} \rrbracket \rho[x \mapsto \nu]$$

Here the right hand side may not be well-defined. For example, if  $\llbracket e_2 \rrbracket \rho$  is a co-finite set, the right hand side is an infinite union. The infinite union may result in a set which is neither finite nor co-finite, and hence, not in the semantic universe  $U$ . One example is

$$\{\{x\} \mid x \in \text{cml}\{\}\}.$$

Intuitively, this is the set of all singletons; it is neither finite nor co-finite, and hence, has no meaning in our semantics.

## 4 The Polarity Discipline

We define an ordering  $\sqsubseteq$  as the least relation on  $U$  satisfying the following properties:

- $\nu \sqsubseteq \nu$  for any  $\nu$
- $(\nu_1, \dots, \nu_n) \sqsubseteq (\nu'_1, \dots, \nu'_n)$  if  $\nu_i \sqsubseteq \nu'_i$  for each  $1 \leq i \leq n$ .
- $\nu \sqsubseteq \nu'$  if  $\nu$  and  $\nu'$  are sets, and for every  $\nu_0 \in \nu$ , there exists a  $\nu'_0 \in \nu'$  such that  $\nu_0 \sqsubseteq \nu'_0$ .

This is the ordering on  $U$  that would arise from treating it as the solution of a domain equation as a partial order, with coordinate-wise ordering on tuples and the lower powerdomain ordering on sets. Note in particular that  $\nu \sqsubseteq \nu'$  implies  $\nu \sqsubseteq \nu'$ . Also, elements of  $C$  are  $\sqsubseteq$ -comparable only to themselves: if  $\nu \in C$  and  $\nu \sqsubseteq \nu'$  or  $\nu' \sqsubseteq \nu$ , then  $\nu = \nu'$ .

For a polarity  $\gamma$  we define

$$\sqsubseteq_{\gamma} = \begin{cases} \sqsubseteq & \text{if } \gamma = +, \\ \supseteq & \text{if } \gamma = -, \end{cases}$$

and we extend the ordering  $\sqsubseteq$  to environments:  $\rho \sqsubseteq \rho'$  iff

- $\rho(x^{\gamma}) \sqsubseteq_{\gamma} \rho'(x^{\gamma})$  for every  $x^{\gamma}$ , and
- $\rho(K, u^{\gamma}) \sqsubseteq_{\gamma} \rho'(K, u^{\gamma})$  for every  $K \$ u^{\gamma}$ .

We will adopt the convention that  $\llbracket e \rrbracket \rho \sqsubseteq \llbracket e' \rrbracket \rho'$  iff  $\llbracket e \rrbracket \rho$  and  $\llbracket e' \rrbracket \rho'$  are both undefined, or they are both defined and ordered by  $\sqsubseteq$ .

The rules for assigning polarities to QCM expressions are given in Table 3.

Table 2: Clauses of QCM's denotational semantics

**Variables:**  $\llbracket x \rrbracket \rho = \rho(x)$

**Constants:** Constants are given their usual interpretation. For instance  $\llbracket 2 \rrbracket \rho$  is the number 2 viewed as an element of the  $\llbracket \text{Num} \rrbracket$  part of  $U$ . Principals  $K$  take their meanings in the  $\llbracket \text{Key} \rrbracket$  part of the domain.

**Qualified Names:**  $\llbracket e\$u \rrbracket \rho = \rho(\llbracket e \rrbracket \rho, u)$ . Note that the conventions about definition mean that the meaning of  $e\$u$  with respect to environment  $\rho$  is given by the expression on the right if  $\llbracket e \rrbracket \rho$  is defined and is a principal; it is undefined otherwise. Also, our restriction on environments implies that the meaning of  $e\$u$  is a set if it is defined.

**Evaluate At:**  $\llbracket e_1 @ e_2 \rrbracket \rho = \llbracket e_1 \rrbracket \rho$  provided  $\llbracket e_2 \rrbracket \rho \in \llbracket \text{Key} \rrbracket$ ; otherwise  $\llbracket e_1 @ e_2 \rrbracket \rho$  is undefined.

**Enumerated Sets:**  $\llbracket \{e_1, \dots, e_n\} \rrbracket \rho = \{\llbracket e_1 \rrbracket \rho, \dots, \llbracket e_n \rrbracket \rho\}$ .

**Co-finite Sets:**  $\llbracket \text{cml}\{w_1, \dots, w_n\} \rrbracket \rho = U - \{\llbracket w_1 \rrbracket \rho, \dots, \llbracket w_n \rrbracket \rho\}$ .

Notice here that the complement is taken with respect to  $U$ , not  $C$ .

**Tuples:**  $\llbracket (e_1, \dots, e_n) \rrbracket \rho = (\llbracket e_1 \rrbracket \rho, \dots, \llbracket e_n \rrbracket \rho)$ .

**Set Unions:**  $\llbracket \bigcup e \rrbracket \rho = \{\nu \mid \nu \in \nu' \text{ for some } \nu' \in \llbracket e \rrbracket \rho\}$ .

**Is an Element of:**  $\llbracket \{e_1 \mid x \in e_2, \mathbf{g}\} \rrbracket \rho = \bigcup_{\nu \in \llbracket e_2 \rrbracket \rho} \llbracket \{e_1 \mid \mathbf{g}\} \rrbracket \rho[x \mapsto \nu]$

**Is a Pattern in:**  $\llbracket \{e_1 \mid (\mathbf{x}) \in e_2, \mathbf{g}\} \rrbracket \rho = \bigcup_{(\nu) \in \llbracket e_2 \rrbracket \rho} \llbracket \{e_1 \mid \mathbf{g}\} \rrbracket \rho'$   
where  $\rho' = \rho[\mathbf{x} \mapsto \nu]$ .

**Is Equal to:**  $\llbracket \{e_1 \mid e_2 = e_3, \mathbf{g}\} \rrbracket \rho = \llbracket \{e_1 \mid \mathbf{g}\} \rrbracket \rho$  provided  $\llbracket e_2 \rrbracket \rho = \nu_2$  and  $\llbracket e_3 \rrbracket \rho = \nu_3$  are equal and both are elements of  $C$ . If they are both in  $C$  but they are not equal then the meaning of the expression is the empty set  $\{\}$ . If either of them is undefined or not in  $C$  then the meaning is undefined.

**Is not Equal to:**  $\llbracket \{e_1 \mid e_2 \neq e_3, \mathbf{g}\} \rrbracket \rho = \llbracket \{e_1 \mid \mathbf{g}\} \rrbracket \rho$  provided  $\llbracket e_2 \rrbracket \rho = \nu_2$  and  $\llbracket e_3 \rrbracket \rho = \nu_3$  are unequal and both are elements of  $C$ . If both are elements of  $C$  but they are not equal then the meaning of the expression is the empty set  $\{\}$ . If either of them is undefined or not in  $C$  then the meaning is undefined.

**Is not an Element of:**  $\llbracket \{e_1 \mid e_2 \notin e_3, \mathbf{g}\} \rrbracket \rho = \llbracket \{e_1 \mid \mathbf{g}\} \rrbracket \rho$  provided  $\llbracket e_2 \rrbracket \rho = \nu_2$  is in  $C$ , and  $\llbracket e_3 \rrbracket \rho$  is a set  $\nu_3$  and  $\nu_2 \notin \nu_3$ . If  $\nu_2$  is in  $C$  and  $\nu_3$  is a set with  $\nu_2 \in \nu_3$ , then the meaning is the empty set  $\{\}$ . If  $\nu_2$  is not in  $C$ , or  $\nu_3$  is not a set, then the meaning is undefined. Note that  $\nu_3$  may contain elements not in  $C$ .

**Base Case:**  $\llbracket \{e \mid \} \rrbracket \rho = \llbracket \{e\} \rrbracket \rho$ .

Table 3: QCM's polarity rules

$c : \gamma$	$\frac{e_1 : \gamma \quad \dots \quad e_n : \gamma}{(e_1, \dots, e_n) : \gamma}$	$\frac{e_1 : \gamma \quad \dots \quad e_n : \gamma}{\{e_1, \dots, e_n\} : \gamma}$	$x^\gamma : \gamma$	$e\$u^\gamma : \gamma$	$\bigcup_{e : \gamma} \frac{e : \gamma}{e : \gamma}$
$\frac{e_1 : \gamma}{(e_1 @ e_2) : \gamma}$	$\text{cml}\{\mathbf{w}\} : \gamma$	$\frac{e : \gamma}{\{e \mid \} : \gamma}$	$\frac{\{e_1 \mid \mathbf{g}\} : \gamma \quad x : \gamma \quad e_2 : \gamma}{\{e_1 \mid x \in e_2, \mathbf{g}\} : \gamma}$	$\frac{\{e_1 \mid \mathbf{g}\} : \gamma \quad p : \gamma \quad e_2 : \gamma}{\{e_1 \mid p \in e_2, \mathbf{g}\} : \gamma}$	
	$\frac{\{e_1 \mid \mathbf{g}\} : \gamma}{\{e_1 \mid e_2 = e_3, \mathbf{g}\} : \gamma}$	$\frac{\{e_1 \mid \mathbf{g}\} : \gamma}{\{e_1 \mid e_2 \neq e_3, \mathbf{g}\} : \gamma}$	$\frac{\{e_1 \mid \mathbf{g}\} : \gamma \quad e_3 : -\gamma}{\{e_1 \mid e_2 \notin e_3, \mathbf{g}\} : \gamma}$		

The last, most interesting rule shows how non-membership guards affect polarity:

$$\frac{\{e_1 \mid \mathbf{g}\} : \gamma \quad e_3 : -\gamma}{\{e_1 \mid e_2 \notin e_3, \mathbf{g}\} : \gamma}$$

Here,  $-+ = -$  and  $-- = +$ , so the rule is contravariant in  $e_3$ . Every other rule treats polarities covariantly.

Another point to note is that the rules do not take the polarities of operands of other guards into consideration (e.g., the polarities of  $e_2$  and  $e_3$  in the guard  $e_2 = e_3$  are ignored). This is because the denotational semantics forces their meanings to be in  $C$ , and every element of  $C$  is  $\sqsubseteq$ -comparable only to itself. Similarly, the polarity of the qualifier  $e$  in the expression  $e\$u^\gamma$  is also irrelevant.

Some expressions (including all values) have both positive and negative polarity; some have only one or the other; and some have neither polarity. For example, set difference has the following derived polarity rule:

$$\frac{e_1 : \gamma \quad e_2 : -\gamma}{e_1 - e_2 : \gamma}$$

Here the subtrahend is treated contravariantly. Consequently, the expression  $x^+ - x^+$  has no polarity. Nevertheless, it is possible to use positive variables in ‘negative’ positions, as in the following example.

$$\{y^+ \mid x^+ \in \{K, K'\}, \\ y^+ \in x^+\$ACL^+ - x^+\$CRL^-\} : +$$

The expression denotes elements of  $U$  that appear on the ACL’s of  $K$  and  $K'$ , but not on their respective CRL’s. The positive variable  $x^+$  is used in the negative expression  $x^+\$CRL^-$ , but in a position where its polarity does not matter.

It is possible to translate QCM expressions without polarity annotations on variables into internal QCM expressions in time proportional to the size of the term. Thus QCM as a programming language does not need to expose the complexity of tagged variables to the programmer: the existence of a polarity is inferred automatically. However, even for external QCM we insist on annotating the polarities of names: a programmer must indicate for each name whether it is to be viewed as ‘ACL-like’ or ‘CRL-like.’

**Theorem 1 (Monotonicity)** *Let  $e$  be an expression and let  $\rho, \rho'$  be environments such that  $\rho \sqsubseteq \rho'$ . Suppose  $\nu = \llbracket e \rrbracket \rho$  and  $\nu' = \llbracket e \rrbracket \rho'$  are defined.*

1. If  $e : +$ , then  $\nu \sqsubseteq \nu'$ .
2. If  $e : -$ , then  $\nu \supseteq \nu'$ .
3. If  $\nu \in C$  or  $\nu' \in C$ , then  $\nu = \nu'$ .

The Monotonicity Theorem justifies our strategy of evaluating push certificates using the minimal model described in the previous section.

## 5 Operational Semantics

We now present a formal operational semantics for the internal language of QCM. For simplicity, the semantics omits details of the QCM implementation that are not relevant to our main Soundness Theorem. Some of the important omissions are:

- The semantics assumes nodes can communicate instantaneously and securely. In our implementation, nodes must exchange explicit, signed messages to communicate.
- The semantics does not mention expiration times. Our implementation keeps track of a valid time interval for each job executed, and this interval is adjusted each time a certificate is used or message is exchanged.
- The semantics does not describe how certificates are ‘pushed’ at a node, or how certificates are verified, or what happens if a signature is found to be invalid.
- The semantics ignores efficiency; in practice we perform meaning-preserving source-to-source transformations to optimize evaluation. For example, we reorder computations and insert @-annotations to reduce the size of messages and intermediate results.

Another paper describes how our implementation handles these details [9].

The rules of Table 4 define a rewriting relation,  $\rightarrow$ , for computations that can be performed locally, on a single node in the network. The rules use an auxiliary function,  $\text{eq}$ , for testing the equality of values. We define  $\text{eq}$  to be the least partial function from values to booleans satisfying the following conditions.

- $\text{eq}(c, c) = \text{true}$
- $\text{eq}(c, c') = \text{false}$  if  $c \neq c'$
- $\text{eq}((v_1, \dots, v_n), (v'_1, \dots, v'_n)) = \text{true}$  if  $\text{eq}(v_i, v'_i) = \text{true}$  for all  $i \in \{1, \dots, n\}$
- $\text{eq}((v_1, \dots, v_n), (v'_1, \dots, v'_n)) = \text{false}$  if  $\text{eq}(v_j, v'_j) = \text{false}$  for some  $j \in \{1, \dots, n\}$

In other words,  $\text{eq}(w, w')$  is true iff  $w = w'$ , and  $\text{eq}(v, v')$  is undefined iff  $v$  or  $v'$  is not a comparable value. In particular,  $\text{eq}$  is not defined on sets, e.g.,  $\text{eq}(\{\}, \{\})$  is undefined. We do not permit equality on sets because this would make the query language nonmonotonic.

Notice that the rule for non-membership may require applying the function  $\text{eq}$  to some values on which it is undefined. However, it is decidable whether  $\text{eq}$  is defined. When  $\text{eq}$  is undefined, an evaluation can become *stuck*.

Other examples of stuck expressions arise from co-finite sets. For example,

- $\{e \mid p \in \text{cml}\{\mathbf{w}\}\}$
- $\bigcup \{\text{cml}\{\mathbf{w}\}, \{\mathbf{v}\}\}$

It seems hard to imagine a rule for effectively evaluating the first expression here; in fact, in our semantics, such expressions may not have a meaning, because they would denote sets which are neither finite nor co-finite. On the other hand, it is possible to extend our rules for  $\bigcup$  to handle co-finite sets. One design issue is that we have used comprehensions to define sugared versions of set intersection and difference; if we cannot use comprehensions to range over co-finite sets, we cannot perform intersection between and difference from co-finite sets. This could be addressed by adding intersection and difference as primitives to the language and adding operational rules for them.

**Lemma 2 (Local Soundness)** *If  $e \rightarrow e'$ , then  $\llbracket e \rrbracket = \llbracket e' \rrbracket$ .*

Table 4: Rules defining the local evaluation relation,  $e \rightarrow e'$ 


---


$$\begin{aligned}
\bigcup(\{\mathbf{v}_1\}, \dots, \{\mathbf{v}_n\}) &\rightarrow \{\mathbf{v}_1, \dots, \mathbf{v}_n\} \\
\{e \mid \} &\rightarrow \{e\} \\
\{e \mid v_1 = v_2, \mathbf{g}\} &\rightarrow \begin{cases} \{\} & \text{if } \text{eq}(v_1, v_2) = \text{false} \\ \{e \mid \mathbf{g}\} & \text{if } \text{eq}(v_1, v_2) = \text{true} \end{cases} \\
\{e \mid v_1 \neq v_2, \mathbf{g}\} &\rightarrow \begin{cases} \{\} & \text{if } \text{eq}(v_1, v_2) = \text{true} \\ \{e \mid \mathbf{g}\} & \text{if } \text{eq}(v_1, v_2) = \text{false} \end{cases} \\
\{e \mid v_1 \notin \{\mathbf{v}\}, \mathbf{g}\} &\rightarrow \begin{cases} \{\} & \text{if } \text{eq}(v_1, v_i) = \text{true for some } v_i \in \mathbf{v} \\ \{e \mid \mathbf{g}\} & \text{if } \text{eq}(v_1, v_i) = \text{false for every } v_i \in \mathbf{v} \end{cases} \\
\{e \mid v \notin \text{cmp}\{\mathbf{w}\}, \mathbf{g}\} &\rightarrow \begin{cases} \{e \mid \mathbf{g}\} & \text{if } \text{eq}(v, w_i) = \text{true for some } w_i \in \mathbf{w} \\ \{\} & \text{if } \text{eq}(v, w_i) = \text{false for every } w_i \in \mathbf{w} \end{cases} \\
\{e \mid p \in \{\}, \mathbf{g}\} &\rightarrow \{\} \\
\{e \mid x \in \{v, \mathbf{v}'\}, \mathbf{g}\} &\rightarrow \bigcup\{ \{e \mid \mathbf{g}\}[x \mapsto v], \{e \mid p \in \{\mathbf{v}'\}, \mathbf{g}\} \} \\
\{e \mid (x_1, \dots, x_n) \in \{(v_1, \dots, v_n), \mathbf{v}'\}, \mathbf{g}\} &\rightarrow \bigcup\{ \{e \mid \mathbf{g}\}[x_1, \dots, x_n \mapsto v_1, \dots, v_n], \{e \mid p \in \{\mathbf{v}'\}, \mathbf{g}\} \}
\end{aligned}$$


---

Table 5: Jobs and evaluation contexts

---

Job numbers	$n$	$\in$	Num
Job identifiers	$i, j$	$::=$	$(K, n)$
Jobs	$J$	$::=$	$w_j^i(e) \mid a_j^i(E[e@K])$
Evaluation contexts	$E$	$::=$	$[\cdot] \mid E\$u \mid e@E \mid (\mathbf{v}, E, \mathbf{e}) \mid \{\mathbf{v}, E, \mathbf{e}\} \mid \bigcup E \mid \{e \mid p \in E, \mathbf{g}\} \mid \{e \mid E = e, \mathbf{g}\} \mid \{e \mid v = E, \mathbf{g}\} \mid \{e \mid E \neq e, \mathbf{g}\} \mid \{e \mid v \neq E, \mathbf{g}\} \mid \{e \mid E \notin e, \mathbf{g}\} \mid \{e \mid v \notin E, \mathbf{g}\}$

---

Next we define a left-to-right order of evaluation, using the machinery of redexes and evaluation contexts.

**Definition:** We say an expression  $e$  is a *local redex* if it matches the left hand side of a rule of Table 4. We say  $e$  is a *global redex* if it has the form  $e@K$  or  $K\$u$ .

Evaluation contexts are defined in Table 5. If  $E$  is an evaluation context and  $e$  is an expression, then  $E[e]$  is the expression obtained by replacing the hole,  $[\cdot]$ , of  $E$  with  $e$ . Note that the grammar for  $E$  ensures that the hole does not appear in the scope of a variable binding, so capture of  $e$ 's free variables is not an issue. We say  $e$  *has parse*  $E[e']$  if  $e = E[e']$  and  $e'$  is a redex.

Redexes and evaluation contexts establish a notion of “next step,” formalized as follows.

**Lemma 3 (Parsing)** *For any expression  $e$ , exactly one of the following holds:*

1.  $e$  is a value;
2. there is a unique parse  $E[e']$  of  $e$ ; or
3.  $e$  is not a value and does not have a parse.

In other words, either evaluation is completed, or there is a unique redex to work on, or the expression is stuck.

Global computation in QCM involves multiple *jobs* running at locations distributed throughout a network. To simplify our formal treatment, we identify locations with principals; that is, each principal is a distinct location. To distinguish the jobs running at a location we use *job numbers*, ranged over by  $n$ . Each location will ensure that it assigns a unique job number to each of its jobs, but because a location has no control over the job numbers assigned by other locations, unique global *job identifiers* consist of the job number and location together. We use  $i$  and  $j$  to range over job identifiers, and we write  $\text{loc}(i)$  for the location given by identifier  $i$ : if  $i = (K, n)$ , then  $\text{loc}(i) = K$ .

Each job  $J$  runs at the request of a client, so jobs are associated with two job identifiers, the identifier of the job itself, and the identifier of the client job. Jobs can be in one of two states, given by the following constructors:

- $w_j^i(e)$ : A job running at  $i$  for client  $j$ , working on the expression  $e$ .
- $a_j^i(E[e@K])$ : A job running at  $i$  for client  $j$ , awaiting a value for  $e$  from location  $K$ .

A *global state* consists of a set of jobs and a set of job identifiers, and is written  $\mathbf{J}; \mathbf{i}$ . We write  $\vdash \mathbf{J}; \mathbf{i}$  and say  $\mathbf{J}; \mathbf{i}$  is *well-formed* if every identifier in  $\mathbf{J}$  appears in  $\mathbf{i}$ , and every job in  $\mathbf{J}$  has a distinct identifier and client identifier (for simplicity, a client can have only one outstanding request). The set  $\mathbf{i}$  will be used to ensure that job numbers are not re-used.

A network of QCM programs is described by a partial function  $D : \text{Key} \times \text{Key} \times \text{Name} \rightarrow \text{Exp}$ . The expression  $D(K_1, K_2, u)$  is what  $K_1$  believes  $K_2$  to be, and may be completely unrelated to  $D(K_2, K_2, u)$ . We require  $D(K_1, K_2, u)$  to be closed. We write  $\rho \models D$  if for all  $K$ , we have  $\rho(K\$u) = \llbracket D(K, K, u) \rrbracket$ , and for all  $K_1$  and  $K_2$ ,  $\llbracket D(K_1, K_2, u^\gamma) \rrbracket \sqsubseteq_\gamma \rho(K_2\$u^\gamma)$ . We allow  $\llbracket D(K_1, K_2, u) \rrbracket$  differ from  $\rho(K_2\$u)$  to model the effect of push certificates in the system.

We define a rewriting relation,  $\Rightarrow_D$ , on global states by the rules of Table 6. We write  $\vdash \mathbf{J}; \mathbf{i} \Rightarrow_D^* \mathbf{J}'; \mathbf{i}'$  if  $\mathbf{J}'; \mathbf{i}'$  is obtained by zero or more rewrites from a well-formed global state  $\mathbf{J}; \mathbf{i}$ .

**Theorem 4 (Soundness)** *If  $\vdash w_j^i(e), \mathbf{J}; \mathbf{i} \Rightarrow_D^* w_j^i(e'), \mathbf{J}'; \mathbf{i}'$ ,  $\rho \models D$ , and  $e : \gamma$ , then  $\llbracket e' \rrbracket \rho \sqsubseteq_\gamma \llbracket e \rrbracket \rho$ .*

## 6 Conclusion

Revocation of certificates presents a variety of fundamental problems for the semantics of certificates and pose significant problems about practical distribution and the appropriate assignment of effort and liability. We have given a language and the first full theoretical semantics for certificate revocation, an important step toward a tractable analysis of these issues. We have also demonstrated a general mechanism for realizing a range of strategies for CRL distribution.

## A Proof Sketches

### A.1 Proof of the Monotonicity Theorem

Theorem 1 is proved by induction on the structure of  $e$ . The interesting case is given below.

- $e = \{e_1 \mid e_2 \notin e_3, \mathbf{g}\}$ . Let  $\nu_2 = \llbracket e_2 \rrbracket \rho$ ,  $\nu'_2 = \llbracket e_2 \rrbracket \rho'$ ,  $\nu_3 = \llbracket e_3 \rrbracket \rho$ , and  $\nu'_3 = \llbracket e_3 \rrbracket \rho'$ .

If  $e : +$  and  $\nu_2 \in \nu_3$ , then  $\nu = \{\}$ . Then  $\nu \sqsubseteq \nu'$  because  $\nu'$  is a set by the definition of meaning, and  $\{\}$  is the least set in the ordering  $\sqsubseteq$ .

If  $e : +$  and  $\nu_2 \notin \nu_3$ , then  $\nu = \llbracket \{e_1 \mid \mathbf{g}\} \rrbracket \rho$ . By induction and the polarity rules,  $\nu_2 \sqsubseteq \nu'_2$  and  $\nu_3 \supseteq \nu'_3$ . By the definition of meaning,  $\nu_2, \nu'_2 \in C$ , so by Lemma 5,  $\nu_2 = \nu'_2$ . Also by the definition of meaning,  $\nu_3, \nu'_3$  are sets, so by Lemma 6,  $\nu_3 \cap C \supseteq \nu'_3 \cap C$ . In particular,  $\nu_2 = \nu'_2 \notin \nu'_3$ . By the definition of meaning,  $\nu' = \llbracket \{e_1 \mid \mathbf{g}\} \rrbracket \rho'$ . Then by induction,  $\nu \sqsubseteq \nu'$ , proving (1).

If  $e : -$  and  $\nu'_2 \in \nu'_3$ , then  $\nu' = \{\}$ . Then  $\nu \supseteq \nu'$  because  $\nu$  is a set by the definition of meaning, and  $\{\}$  is the least set in the ordering  $\sqsubseteq$ .

If  $e : -$  and  $\nu'_2 \notin \nu'_3$ , then  $\nu' = \llbracket \{e_1 \mid \mathbf{g}\} \rrbracket \rho'$ . By induction and the polarity rules,  $\nu_2 \supseteq \nu'_2$  and  $\nu_3 \sqsubseteq \nu'_3$ . By the definition of meaning,  $\nu_2, \nu'_2 \in C$ , so by Lemma 5,  $\nu_2 = \nu'_2$ . Also by the definition of meaning,  $\nu_3, \nu'_3$  are sets, so by Lemma 6,  $\nu_3 \cap C \subseteq \nu'_3 \cap C$ . In particular,  $\nu'_2 = \nu_2 \notin \nu_3$ . By the definition of meaning,  $\nu = \llbracket \{e_1 \mid \mathbf{g}\} \rrbracket \rho$ . Then by induction,  $\nu \supseteq \nu'$ , proving (2).

Both  $\nu$  and  $\nu'$  are sets, so  $\nu, \nu' \notin C$ , proving (3).  $\square$

**Lemma 5** *If  $\nu \sqsubseteq \nu'$ , and  $\nu \in C$  or  $\nu' \in C$ , then  $\nu = \nu'$ .  $\square$*

**Lemma 6** *For any sets  $\nu, \nu' \in U$ , if  $\nu \sqsubseteq \nu'$  then  $\nu \cap C \subseteq \nu' \cap C$ .  $\square$*

### A.2 Proof of the Local Soundness Lemma

The proof of Lemma 2 is trivial for every rule but the last two rules, which require the following lemma.

**Lemma 7 (Substitutivity)**  $\llbracket e \rrbracket (\rho[x \mapsto \llbracket e' \rrbracket \rho]) = \llbracket e[x \mapsto e'] \rrbracket \rho$

**Proof:** By induction on the structure of  $e$ .

- $e = x$ . Then  $\llbracket x \rrbracket (\rho[x \mapsto \llbracket e' \rrbracket \rho]) = \llbracket e' \rrbracket \rho = \llbracket x[x \mapsto e'] \rrbracket \rho$ .
- $e = y \neq x$ . Then  $\llbracket y \rrbracket (\rho[x \mapsto \llbracket e' \rrbracket \rho]) = \llbracket y \rrbracket \rho = \llbracket y[x \mapsto e'] \rrbracket \rho$ .
- $e = \{e_1 \mid y \in e_2, \mathbf{g}\}$ . Then

$$\begin{aligned} & \llbracket e \rrbracket (\rho[x \mapsto \llbracket e' \rrbracket \rho]) \\ &= \bigcup_{\nu \in \llbracket e_2 \rrbracket (\rho[x \mapsto \llbracket e' \rrbracket \rho])} \llbracket e_1 \rrbracket (\rho[x \mapsto \llbracket e' \rrbracket \rho][y \mapsto \nu]). \end{aligned}$$

By induction, this is equal to

$$\bigcup_{\nu \in \llbracket e_2[x \mapsto e'] \rrbracket \rho} \llbracket e_1 \rrbracket (\rho[x \mapsto \llbracket e' \rrbracket \rho][y \mapsto \nu]). \quad (*)$$

We can assume that  $y$  is not  $x$  and is not free in  $e'$  (otherwise, rename  $y$  in  $e$ ). Then

$$\rho[x \mapsto \llbracket e' \rrbracket \rho][y \mapsto \nu] = (\rho[y \mapsto \nu])[x \mapsto \llbracket e' \rrbracket (\rho[y \mapsto \nu])]$$

Then by induction,

$$\llbracket e_1 \rrbracket (\rho[x \mapsto \llbracket e' \rrbracket \rho][y \mapsto \nu]) = \llbracket e_1[x \mapsto e'] \rrbracket (\rho[y \mapsto \nu])$$

This shows that (\*) is equal to

$$\bigcup_{\nu \in \llbracket e_2[x \mapsto e'] \rrbracket \rho} \llbracket e_1[x \mapsto e'] \rrbracket (\rho[y \mapsto \nu]) = \llbracket e[x \mapsto e'] \rrbracket \rho,$$

as desired.

The remaining cases are proved similarly.  $\square$

### A.3 Proof of the Parsing Lemma

**Proof:** First, note that no value contains a redex, so no value has a parse. This shows that (1), (2), and (3) are mutually exclusive. For the remainder, we show that every  $e$  falls into at least one of these cases, by induction on the structure of  $e$ .

- $e = x$ . Then  $e$  is not a value, and it contains no redex, so we have case (3).
- $e = c$ . Then  $e$  is a value, so we have case (1).
- $e = e'\$u$ . By induction, it is sufficient to consider the following cases.
  - $e'$  is a value. If  $e'$  is a principal  $K$ , then  $e$  is a redex and we have (2). Otherwise,  $e$  is not a redex and we have (3).
  - $e'$  has a unique parse  $E[e'']$ . Then there is a unique context  $E' = E\$u$  and redex  $e''$  such that  $e = E[e'']$ , so we have (2).
  - $e'$  is not a value and has no parse. Then we have (3).
- The other cases are proved similarly.  $\square$

Table 6: Rules defining the global evaluation relation,  $\mathbf{J}; \mathbf{i} \Rightarrow_D \mathbf{J}'; \mathbf{i}'$ 

**Local Evaluation** The local evaluation rules are applied at a given node.

$$w_j^i(E[e]), \mathbf{J}; \mathbf{i} \Rightarrow_D w_j^i(E[e']), \mathbf{J}; \mathbf{i} \quad \text{if } e \rightarrow e'$$

**Evaluation of names** When a qualified name  $K\$u$  is encountered at a node with a definition for  $K\$u$ , it is replaced by the definition. If no definition is available locally, then a query will be sent to  $K$ .

$$w_j^i(E[K\$u]), \mathbf{J}; \mathbf{i} \Rightarrow_D w_j^i(E[e]), \mathbf{J}; \mathbf{i} \quad \text{if } D(\text{loc}(i), K, u) = e$$

$$w_j^i(E[K\$u]), \mathbf{J}; \mathbf{i} \Rightarrow_D w_j^i(E[(K\$u)@K]), \mathbf{J}; \mathbf{i} \quad \text{if } D(\text{loc}(i), K, u) \text{ is undefined and } K \neq \text{loc}(i).$$

**Queries** To evaluate an expression at a different location  $K$ , we start a job at  $K$  marked with the requesting location and a fresh job number.

$$w_j^i(E[e@K]), \mathbf{J}; \mathbf{i} \Rightarrow_D w_j^i(E[e]), \mathbf{J}; \mathbf{i} \quad \text{if } K = \text{loc}(i)$$

$$w_j^i(E[e@K]), \mathbf{J}; \mathbf{i} \Rightarrow_D a_j^i(E[e@K]), w_i^{(K,n)}(e), \mathbf{J}; (K, n), \mathbf{i} \quad \text{if } K \neq \text{loc}(i) \text{ and } (K, n) \notin \mathbf{i}$$

**Responses** When  $K_1$  has completed the evaluation of an expression on behalf of  $K$ , then it is provided to  $K$  and substituted into the awaiting context.

$$a_j^i(E[e@K]), w_i^{i'}(v), \mathbf{J}; \mathbf{i} \Rightarrow_D w_j^i(E[v]), \mathbf{J}; \mathbf{i}$$

#### A.4 Proof of the Soundness Theorem

Theorem 4 is proved by induction on the definition of  $\Rightarrow_D^*$ .

- If the reduction has length 0, then  $e = e'$  and the result follows immediately.
- If the first step of the reduction does not involve job  $i$ , then the result is immediate by induction. Note, job  $i$  cannot be the response to a query in this reduction because it still appears at the end of the reduction; the set  $\mathbf{i}$  would prevent it from being re-introduced.
- The reduction has the form

$$\begin{aligned} w_j^i(E[e_1]), \mathbf{J}; \mathbf{i} \\ \Rightarrow_D w_j^i(E[e_1']), \mathbf{J}; \mathbf{i} \\ \Rightarrow_D^* w_j^i(e'), \mathbf{J}'; \mathbf{i}' \end{aligned}$$

where  $e = E[e_1]$  and  $e_1 \rightarrow e_1'$ .

Then  $\llbracket e_1 \rrbracket \rho = \llbracket e_1' \rrbracket \rho$  by the Local Soundness Lemma. It follows immediately that  $\llbracket e \rrbracket \rho = \llbracket E[e_1] \rrbracket \rho = \llbracket E[e_1'] \rrbracket \rho$ ; and by induction,  $\llbracket e' \rrbracket \rho \sqsubseteq_\gamma \llbracket E[e_1'] \rrbracket \rho$ , so that  $\llbracket e' \rrbracket \rho \sqsubseteq_\gamma \llbracket e \rrbracket \rho$  as desired.

- The reduction has the form

$$\begin{aligned} w_j^i(E[K\$u^{\gamma_1}]), \mathbf{J}; \mathbf{i} \\ \Rightarrow_D w_j^i(E[e_1]), \mathbf{J}; \mathbf{i} \\ \Rightarrow_D^* w_j^i(e'), \mathbf{J}'; \mathbf{i}' \end{aligned}$$

where  $e = E[K\$u^{\gamma_1}]$  and  $D(\text{loc}(i), K, u^{\gamma_1}) = e_1$ .

Since  $\rho \models D$ , we have  $\llbracket e_1 \rrbracket \rho \sqsubseteq_{\gamma_1} \llbracket K\$u^{\gamma_1} \rrbracket \rho$ . Then by Lemma 8, we have  $\llbracket E[e_1] \rrbracket \rho \sqsubseteq_\gamma \llbracket E[K\$u^{\gamma_1}] \rrbracket \rho$ . By induction, we have  $\llbracket e' \rrbracket \rho \sqsubseteq_\gamma \llbracket E[e_1] \rrbracket \rho$ , so  $\llbracket e' \rrbracket \rho \sqsubseteq_\gamma \llbracket e \rrbracket \rho$  as desired.

- The reduction has the form

$$\begin{aligned} w_j^i(E[K\$u]), \mathbf{J}; \mathbf{i} \\ \Rightarrow_D w_j^i(E[(K\$u)@K]), \mathbf{J}; \mathbf{i} \\ \Rightarrow_D^* w_j^i(e'), \mathbf{J}'; \mathbf{i}' \end{aligned}$$

where  $e = E[K\$u]$ .

This follows because  $\llbracket K\$u \rrbracket = \llbracket (K\$u)@K \rrbracket$ .

- The reduction has the form

$$\begin{aligned} w_j^i(E[e_1@K]), \mathbf{J}; \mathbf{i} \\ \Rightarrow_D a_j^i(E[e_1@K]), w_i^{i'}(e_1), \mathbf{J}; i', \mathbf{i} \\ \Rightarrow_D^* a_j^i(E[e_1@K]), w_i^{i'}(v), \mathbf{J}_1; \mathbf{i}_1 \\ \Rightarrow_D w_j^i(E[v]), \mathbf{J}_1; \mathbf{i}_1 \\ \Rightarrow_D^* w_j^i(e'), \mathbf{J}'; \mathbf{i}' \end{aligned}$$

By induction,

$$\llbracket v \rrbracket \rho \sqsubseteq_{\gamma_1} \llbracket e_1 \rrbracket \rho = \llbracket e_1@K \rrbracket \rho.$$

Since  $E[e_1@K] : \gamma$ , there must be a polarity  $\gamma_1$  such that  $e_1 : \gamma_1$  and  $E[x^{\gamma_1}] : \gamma$ , where  $x^{\gamma_1}$  is a fresh variable. Then by Lemma 8,

$$\llbracket E[v] \rrbracket \rho \sqsubseteq_\gamma \llbracket E[e_1@K] \rrbracket \rho.$$

By induction,  $\llbracket e' \rrbracket \rho \sqsubseteq_\gamma \llbracket E[v] \rrbracket \rho$ , so  $\llbracket e' \rrbracket \rho \sqsubseteq_\gamma \llbracket e \rrbracket \rho$ , as desired.

- The remaining cases are proved similarly.  $\square$

**Lemma 8** *If  $\llbracket e \rrbracket \rho \sqsubseteq_\gamma \llbracket e' \rrbracket \rho$  and  $E[x^\gamma] : \gamma'$ , where  $x^\gamma$  is not free in  $E$ , then  $\llbracket E[e] \rrbracket \rho \sqsubseteq_{\gamma'} \llbracket E[e'] \rrbracket \rho$ .*

**Proof:** Let  $\rho_1 = \rho[x^\gamma \mapsto \llbracket e \rrbracket \rho]$  and  $\rho_2 = \rho[x^\gamma \mapsto \llbracket e' \rrbracket \rho]$ . Then  $\rho_1 \sqsubseteq \rho_2$ . By the Monotonicity Theorem,

$$\llbracket E[x^\gamma] \rrbracket \rho_1 \sqsubseteq_{\gamma'} \llbracket E[x^\gamma] \rrbracket \rho_2.$$

And by Substitutivity,

$$\llbracket E[x^\gamma] \rrbracket \rho_1 = \llbracket (E[x^\gamma])[x^\gamma \mapsto e] \rrbracket \rho = \llbracket E[e] \rrbracket \rho,$$

and similarly,

$$\llbracket E[x^\gamma] \rrbracket \rho_2 = \llbracket E[e'] \rrbracket \rho.$$

Then we have

$$\llbracket E[e] \rrbracket \rho \sqsubseteq_{\gamma'} \llbracket E[e'] \rrbracket \rho,$$

as desired.  $\square$

## References

- [1] Shimshon Berkovits, Santosh Chokhani, Judit A. Furlong, Jisoo A. Geiter, and Jonathan C. Guild. Public key infrastructure study final report. Technical report, MITRE, on behalf of the National Institute of Standards and Technology, 1994.
- [2] Matt Blaze, Joan Feigenbaum, and Jack Lacy. Decentralized trust management. In *Proceedings of the 17th Symposium on Security and Privacy*, pages 164–173. IEEE Computer Society Press, 1996.
- [3] Peter Buneman, Leonid Libkin, Dan Suciu, Val Tannen, and Limsoon Wong. Comprehension syntax. *SIGMOD Record*, 23(1):87–96, March 1994.
- [4] Peter Buneman, Shamim Naqvi, Val Tannen, and Limsoon Wong. Principles of programming with complex objects and collection types. *Theoretical Computer Science*, 149(1):3–48, September 1995.
- [5] Carl M. Ellison, Bill Frantz, Butler Lampson, Ron Rivest, Brian M. Thomas, and Tatu Ylonen. SPKI certificate theory. Internet Draft, March 1998.
- [6] Carl M. Ellison, Bill Frantz, Butler Lampson, Ron Rivest, Brian M. Thomas, and Tatu Ylonen. SPKI examples. Internet Draft, March 1998.
- [7] Warwick Ford and Michael S. Baum. *Secure Electronic Commerce*. Prentice Hall, 1999.
- [8] Barbara Fox and Brian LaMacchia. Certificate revocation: Mechanics and meaning. In R. Hirschfeld, editor, *Advances in Cryptology: Proceedings of Financial Cryptography '98*, number 1465 in Lecture Notes in Computer Science. Springer-Verlag, 1998.
- [9] Carl A. Gunter and Trevor Jim. Policy directed certificate retrieval, July 1998. [www.cis.upenn.edu/~qcm/papers/qcm-abstract.html](http://www.cis.upenn.edu/~qcm/papers/qcm-abstract.html).
- [10] R. Housley, W. Ford, W. Polk, and D. Solo. *Internet X.509 Public Key Infrastructure: Certificate and CRL Profile*. IETF RFC 2459, January 1999.
- [11] ISO/IEC 9594-1. *Information technology—Open Systems Interconnection—The Directory: Overview of concepts, models and services*, 1997. Equivalent to ITU-T Rec. X.500, 1997.
- [12] ISO/IEC 9794-8. *Information technology—Open Systems Interconnection—The Directory: Authentication framework*, 1997. Equivalent to ITU-T Rec. X.509, 1997.
- [13] Pankaj Kakkar, Michael McDougall, Carl A. Gunter, and Trevor Jim. Credential distribution with local autonomy. [www.cis.upenn.edu/~qcm/papers/autonomy-abstract.html](http://www.cis.upenn.edu/~qcm/papers/autonomy-abstract.html), March 1999.
- [14] S. Kent. *Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management*. IETF RFC 1422, February 1993.
- [15] Patrick D. McDaniel and Sugih Jamin. Windowed key revocation in public key infrastructures, 1999.
- [16] Michael Myers, Rich Ankney, Ambarish Malpani, Slava Galperin, and Carlisle Adams. *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol—OCSP*, March 1999.
- [17] Ronald L. Rivest. Can we eliminate certificate revocation lists? In *Financial Cryptography: Second International Conference*, Anguilla, British West Indies, 23–25 February 1998.
- [18] Ronald L. Rivest and Butler Lampson. SDSI—a simple distributed security infrastructure. [theory.lcs.mit.edu/~cis/sdsi.html](http://theory.lcs.mit.edu/~cis/sdsi.html), 1996.
- [19] Stuart Stubblebine. Recent-secure authentication: Enforcing revocation in distributed systems. In *Proceedings of the 1995 IEEE Symposium on Research in Security and Privacy*, pages 224–235, 1995.
- [20] Sean Turner and Alfred Arsenault. *Internet X.509 Public Key Infrastructure: PKIX Roadmap*. IETF, 1999. [www.ietf.org/internet-drafts/draft-ietf-pkix-roadmap-01.txt](http://www.ietf.org/internet-drafts/draft-ietf-pkix-roadmap-01.txt).
- [21] W. Yeong, T. Howes, and S. Kille. *Lightweight Directory Access Protocol*. IETF RFC 1777, 1995.