

# A Formal Implementation of Value Commitment

Cédric Fournet<sup>2,1</sup>, Nataliya Guts<sup>1</sup>, and Francesco Zappa Nardelli<sup>3,1</sup>

<sup>1</sup> MSR-INRIA Joint Centre

<sup>2</sup> Microsoft Research

<sup>3</sup> INRIA

**Abstract.** In an optimistic approach to security, one can often simplify protocol design by relying on audit logs, which can be analyzed a posteriori. Such auditing is widely used in practice, but no formal studies guarantee that the log information suffices to reconstruct past runs of the protocol, in order to reliably detect (and provide evidence of) any cheating. We formalize audit logs for a sample optimistic scheme, the value commitment. It is specified in a pi calculus extended with committable locations, and compiled using standard cryptography to implement secure logs. We show that our distributed implementation either respects the abstract semantics of commitments or, using information stored in the logs, detects cheating by a hostile environment.

## 1 A Cautiously Optimistic Approach to Security

Mutual distrust in distributed computing makes enforcing system-wide security assurances particularly challenging. Common protocols perform an important number of mandatory runtime checks and allow only legal computations to progress: in session-establishment protocols, for instance, a strong security invariant is usually enforced at every step of the run of the protocol. These runtime checks have a cost, in terms of cryptographic and networking operations; they may also conflict with other goals of the protocol, such as confidentiality.

A different approach, which we call optimistic, presumes instead that all involved principals are honest and well-behaved, and thus omits some runtime checks. Traces of protocol runs are stored in a secure log and can be used a posteriori to verify the compliance of each principal to its role: principals who attempt non-compliant actions will be blamed using the logged evidence. The security invariant is weaker than those achieved by more conservative protocols, but adequate for many non-critical applications.

Some protocols inherently rely on logs to establish their security properties. These protocols are often based on a *commitment* scheme. A principal commits to a value kept hidden; other principals of a system cannot read this value, but have a procedure to detect any change to the value after the commitment. Distant coin flipping is a simple protocol that illustrates commitments: suppose that A and B are not physically at the same place and want to toss a coin. Both A and B flip their own coin, exchange commitments on their results, then reveal and compare these results; A wins the toss if the two results are the same. For fairness, A's commitment should neither reveal any information to B, nor enable A to change her committed result after receiving B's.

Commitment is a building block for many protocols such as mental poker [3], sealed bid auctions, e-voting [6,5], and online games [12]. For instance, mental poker relies

on commitment to build a fair shuffling of the deck, then gradually reveal cards as the game proceeds. At the end of the game, the deck permutations used by each player can be revealed for auditing purposes.

Secure logging is not only an essential component of optimistic schemes, but is also widely used in standard practice. Much research effort has been devoted to techniques for implementing logs so as to guarantee properties such as *correctness*, *forward integrity*, and *forward secrecy* [15,18,17]. Still, which data should be logged? and why? Between general recommendations such as “an audit trail should include sufficient information to establish what events occurred and who (or what) caused them” [14,11] and efficient implementation techniques, we are not aware of any formal studies that characterize and verify the security properties achieved by protocols relying on logs.

In this paper, we give a formal answer to this question for the commitment scheme. We extend a simple distributed language, the applied pi calculus [1], with commitment datatypes and primitives, and we illustrate this extension by programming an online game. To abstract away from the possible misbehaviors of the environment, we propose a trustful and strong operational semantics for our commitment primitives. We show that our language can be compiled to the applied pi calculus, using standard cryptographic primitives, with adequate protection against an arbitrary, possibly hostile environment. We obtain an important security property stating that, for any source systems, our distributed implementation either respects the semantics of commitments or, using information stored in the logs, detects (and proves) cheating by a hostile environment.

**Related work.** Value commitments appear in formal models of protocols (e.g. [13]) and implementations of language abstractions (e.g. [19]). More closely related to our work, Etalle et al. [10,4] advocate the usage of logs for optimistic security enforcement. They formalize audit-based discretionary access control in collaborative work environments, and develop a logical framework for user accountability; they also design cryptographic support for communication evidence in a decentralized setting [8].

**Contents.** Section 2 presents our source language with value commitment. Section 3 illustrates the use of commitment for programming online games. Section 4 describes the language implementation, as a cryptographic translation to the applied pi calculus. Section 5 develops a labeled semantics and an extended translation to keep track of source-program invariants. Section 6 states our main results. Section 7 reports on our prototype implementation. Section 8 discusses future work.

Additional details appear online, at <http://www.msr-inria.inria.fr/projects/sec/logs>, including complete definitions for the source and target semantics and all proofs.

## 2 A Language with Value Commitment

The applied pi calculus is a process language parametrized by an equational theory on terms, which provides flexible support for modeling symbolic cryptographic primitives and data structures. We refer to [1] for a general presentation of its semantics.

To express the value commitment scheme, we extend an instance of applied pi with *committable cells*. The grammar for terms ( $M, V$ ), processes ( $P$ ), and systems ( $A$ ) is given below. Our extensions to the standard syntax appear in grey boxes.

$M, V ::=$	$P ::=$	$A ::=$
$u$ $func(\widetilde{M})$ $u.\mathbf{Idu}$ $u.\mathbf{Idc}(p)$ $u.\mathbf{Rd}(p M)$	$0$ $P_1 \mid P_2$ $\nu c.P$ $u?(x).P$ $u!(M).P$ $\text{if } M = M' \text{ then } P \text{ else } P'$ $\text{repl } P$ $\text{newloc}(x, y).P$ $\text{commit } M u(x).P$	$0$ $A_1 \mid A_2$ $\nu u.A$ $\{M/x\}$ $p[P]$ $u.(p)$ $u.(p M)$

Terms are built from variables (denoted  $x, y, \dots$ ), names (denoted  $c, l, s, \dots$ ), function applications, and capabilities (described below). We assume that functions include at least a pairing function, denoted  $+$ , with associated projections  $+_1, +_2$  and equations  $+_i(x_1 + x_2) = x_i$  for  $i = 1, 2$ . (Our results extend to arbitrary data structures; we use integer constants in examples.) The metavariable  $u$  ranges over names and variables. Among names, we distinguish the set of *principals*, denoted  $p, a, e$ , and the set of *location names*, ranged over by  $l$ . Contrarily to standard applied pi, each process  $P$  runs under the control of a principal  $p$ , denoted  $p[P]$ .

**Committable cells and capabilities.** A cell is a memory location owned by a principal who can, once, commit its content to a value of its choice. In addition, the owner can pass capabilities to other principals, thereby granting these principals partial read access to the cell.

Our language features three kinds of capabilities. The *read capability*  $l.\mathbf{Rd}(p M)$  is created by the owner  $p$  of the location  $l$  when it commits to a value  $M$ . Any principal can use a read capability to read the content of the location associated to the capability. The *identity capabilities* instead partially disclose the state of a cell without actually revealing the value possibly committed. So the *committed id capability*  $l.\mathbf{Idc}(p)$  proves that the location  $l$  is committed and reveals the owner  $p$  of the location. The *uncommitted id capability*  $l.\mathbf{Idu}$  just asserts the identity  $l$  of the location.

The language of terms is sorted: we distinguish *marshallable values*, that include all the terms except location and channel names, and *committable values*, that include all marshallable values except those that mention committed id and read capabilities.

The state of each committable cell is represented by a process:  $l.(p)$  denotes an uncommitted cell named  $l$  owned by  $p$ ;  $l.(p M)$  denotes the same cell once it has been committed to the committable value  $M$ . Two new kinds of processes manipulate cells. The *newloc* process creates a fresh, uncommitted location and binds both its unique identifier  $l$  (from  $\mathcal{L}$ ) and its uncommitted capability in its continuation:

$$a[\text{newloc}(x, y).P] \longrightarrow \nu l.(l.(a) \mid a[P\{l/x\}\{l.\mathbf{Idu}/y\}])$$

where  $l$  is fresh for  $P$ . The unique identifier  $l$  can then be used to commit an uncommitted cell to some committable value  $M$ :

$$l.(a) \mid a[\text{commit } M l(x).P] \longrightarrow l.(a M) \mid a[P\{l.\mathbf{Rd}(a M)/x\}]$$

The commit process yields a read capability for the newly-committed cell. The sort system does not allow to communicate or store in another location the cell name  $l$ : hence,

only the principal that created the cell can commit a value into it. The abbreviation `newcommit` creates a new committed location (where  $x', x''$  are fresh for  $P$ ):

$$p[\text{newcommit } M(x).P] \stackrel{\text{def}}{=} p[\text{newloc}(x', x'').\text{commit } M\ x'(x).P]$$

Capabilities can be communicated over channels; they can also be manipulated using special functions, according to the equational theory below.

$$\begin{aligned} \text{read}(x.\mathbf{Rd}(p\ v)) &= v & \text{get\_idc}(x.\mathbf{Rd}(p\ v)) &= x.\mathbf{Idc}(p) \\ \text{get\_idu}(x.\mathbf{Idc}(p)) &= x.\mathbf{Idu} & \text{get\_prin}(x.\mathbf{Idc}(p)) &= p \\ \text{is\_idu}(x.\mathbf{Idu}) &= \text{ok} & \text{is\_idc}(x.\mathbf{Idc}(p)) &= \text{ok} & \text{is\_rd}(x.\mathbf{Rd}(p\ v)) &= \text{ok} \end{aligned}$$

The read function yields the value from read capabilities. Since the read capability is generated when committing the cell, the semantics of the source language guarantees that all reads for a given cell always return the same value. The `get_prin` function yields the principal that owns the cell from committed capabilities. (We could also provide `get_prin` from uncommitted capabilities, at some additional cost in the cryptographic implementation.) The `get_idu` and `get_idc` functions downgrade capabilities, yielding a more restrictive capability for the same cell. Hence, `get_idu` yields an uncommitted capability, which can be used only to identify the cell, whereas `get_idc` takes a read capability and hides its committed value. The language finally has functions that support dynamic typechecking of capabilities. In particular,  $\text{is\_idc}(x) = \text{ok}$  or  $\text{is\_rd}(x) = \text{ok}$  implies that the cell associated with  $x$  is committed.

### 3 Example: An Online Game

Our example describes a game run by a server  $a_0$ , between  $n$  players  $a_1, \dots, a_n$ . The game is played in one turn, with all players revealing their moves simultaneously. (A simple instance of the game with  $n = 2$  is Rock, Paper, Scissors.) The players and the server are willing to cooperate, but with minimal trust assumptions between them; however, it is deemed sufficient to detect any dishonest principal at the end of the game. Similar examples include multiparty protocols for online auctions, voting, or partial-information games [16,3,6].

The protocol has three exchange rounds between the server and each player, using channels  $c_i$  for  $i = 1..n$ : (1) the server sets up the game, distributes the details to the players, and collects their sealed moves; (2) the server distributes all the players' sealed moves and collects their actual moves; (3) the server distributes the result of the game.

We begin with the server code, given below. For simplicity, the code does not provide any error handling—execution stops when a test fails.

$$\begin{aligned} A_0 &= a_0[\text{newloc}(l, \text{result}_{id}).\text{newcommit } \text{result}_{id} + \text{details}(challenge). \\ &\quad (c_i!\langle challenge \rangle.c_i?(promise_i).\text{if } \text{get\_prin}(promise_i) = a_i \text{ then})_{i=1..n} \\ &\quad \text{newcommit } challenge + \widetilde{promise}(game). \\ &\quad (c_i!\langle game \rangle.c_i?(move_i).\text{if } \text{get\_idc}(move_i) = promise_i \text{ then})_{i=1..n} \\ &\quad \text{commit winner}(\widetilde{move}, challenge)\ l(\text{result}).(c_i!\langle result \rangle.0)_{i=1..n}] \end{aligned}$$

In round (1), the server creates an uncommitted cell  $l$  for storing the outcome of the game, and a readable cell *challenge* that provides the identifier for  $l$  and the (unspecified) details of the game. Upon receiving each player's response, the server authenticates it as a committed capability from that player. In round (2), the server creates a second committed cell that binds the challenge to the received commitments from all players. Upon receiving each player's second response, the server correlates it as the read capability associated with their first response. In round (3), the server has all the players' information: it resolves the game and finally commits the cell  $l$  to the published result of the game (which may include, for instance, selected information from the players' moves). We omit the code for the function *winner* that computes this result.

The code for the players performs symmetric operations:

$$A_i = a_i[c_i?(challenge).if\ get\_prin(get\_idc(challenge)) = a_0\ then \\ newcommit\ z_i(move_i).c_i!\{get\_idc(move_i)\}. \\ c_i?(game).if\ valid\_game(game, challenge, move_i)\ then \\ c_i!\{move_i\}.c_i?(result_i).if\ no\_cheat(result_i, read(game))\ then\ P_i]$$

In round (1), after receiving the challenge, each player  $l$  confirms its validity, for instance by checking that it is a genuine readable capability from  $a_0$ , then it selects a move and sends back its commitment. In round (2), after receiving all commitments, the player correlates them to the challenge and verifies that its own commitment is recorded (using for instance *valid\_game*) then it releases its move in clear. In round (3), the player checks the outcome of the game and verifies a posteriori that the server followed the rules (using for instance *no\_cheat*). The tests are defined as follows:

$$valid\_game(x_1, x_2, x_3) \stackrel{def}{=} +_1(read(x_1)) = x_2\ and\ get\_idc(x_3) \in +_2(read(x_1)) \\ no\_cheat(x, y) \stackrel{def}{=} get\_idu(get\_idc(x)) = +_1(y)\ and\ get\_idc(x) \in +_2(y)$$

**Guarantees offered to the players.** We distinguish *language level* guarantees, enforced by the abstract semantics of locations, and *application level* guarantees, relying on high-level, application-specific checks on top of the language semantics. For each kind of guarantees, we also distinguish between immediate (conservative) and deferred (optimistic) enforcement. For instance, enforcement may be deferred until the content of a cell becomes readable.

As an illustration of immediate language-level checks, committed values offer basic authentication guarantees to the participants. For instance, each player has the privilege to choose its moves, and the move is securely attributed to the player even if the communication channels  $c_i$  are unprotected; participants can also check this attribution later.

To protect application integrity, the code must perform sufficient checks before proceeding with the game. Systematic testing of the owner identities for the received capabilities avoids unauthorized, possibly non-accountable, participants. Some checks are immediate, e.g. testing if two capabilities are associated to the same location; other checks that depend on the commitment semantics are delayed. In the example, players are guaranteed that they all get the same result (if any) for any given game, since they must get the same location read capability, but it is up to the application code to correlate the received read capability to the initial uncommitted capability.

At the same time, the applicative logic of our protocol guarantees that, even if the server is willing to leak information to the other players, those players cannot get that information before committing to their own moves.

## 4 Distributed Cryptography Implementation

The target language is an instance of applied pi, with standard (symbolic) cryptographic primitives and data structures but without ad-hoc rules or constructs for locations.

We rely on a cryptographic hash function, denoted  $h$ , and a public-key signature mechanism satisfying the equation  $\text{verify}(v, \text{sign}(v, \text{sk}(m)), \text{pk}(m)) = \text{ok}$ . The functions  $\text{sk}(m)$  and  $\text{pk}(m)$  generate a pair of secret/public keys from a nonce  $m$ . All other data constructors admit a projection function  $\text{func}_i(\text{func}(x_1, \dots, x_n)) = x_i$ .

To every principal  $p$ , we associate a keypair and export its public key tagged with constructor  $\text{prin}$  using an active substitution of the form  $\{\text{prin}(\text{pk}(m_p)) / p\}$ .

**Cryptographic implementation of capabilities.** We compile the capabilities associated to a location  $l.(p \ V)$  as follows:

$$\begin{array}{ll} l. \mathbf{Rd} \ (p \ V) & \text{rd}(p, s, \llbracket V \rrbracket, w) \\ l. \mathbf{Idc} \ (p) & \text{idc}(p, h(s) + h(s + \llbracket V \rrbracket), w) \\ l. \mathbf{Idu} & \text{idu}(h(p + h(s))) \end{array}$$

where  $p = \text{prin}(\text{pk}(m_p))$  is the owner's public key,  $s$  is a fresh value used as a seed, and  $w = \text{sign}(h(s) + h(s + \llbracket V \rrbracket), \text{sk}(m_p))$  signs the committed value  $\llbracket V \rrbracket$ .

A read capability is a tagged tuple that includes these elements. A committed id capability is a tagged tuple that provides  $p$  and verifiable evidence of the commitment without actually revealing  $\llbracket V \rrbracket$ . To this end, it includes both a hash of the committed value, first concatenated with the seed  $s$ , to protect against brute force attacks, yielding  $h(s + \llbracket V \rrbracket)$ , and the hash  $h(s)$ , to enable the receiver to correlate the owner and signature with a previously-received uncommitted id capability by recomputing the identifier  $h(p + h(s))$ . An uncommitted id capability just includes this unique location identifier, which may be compared to other capabilities and, later, associated with  $p$  and  $s$ . The receiver can compute committed capabilities from read capabilities, and uncommitted capabilities from committed capabilities, but not the converse.

The signature  $w$  authenticates read and committed id capabilities, binding their content to the owner's key  $\text{sk}(m_p)$ . Their receiver can extract  $p$  and  $h(s) + h(s + \llbracket V \rrbracket)$  from these tagged tuples and use them to verify  $w$ . When the signature is valid, the public key identifies the owner of the location associated to the capability.

**Detection of multiple commitments.** In a typical run, an honest principal receives a commitment to some value from the principal  $p$ , say  $\text{idc}(p, v_1 + v_2, w)$ , and later the value itself, say  $\text{rd}(p, s, z, w')$ . The receiver can easily check that the two capabilities refer to the same location, by testing  $h(s) = v_1$ , and verify the two signatures  $w = \text{sign}(v_1 + v_2, \text{sk}(m_p))$  and  $w' = \text{sign}(h(s) + h(s + z), \text{sk}(m_p))$ . If these tests succeed, then the receiver can check whether  $v_2 = h(s + M)$ : if the test fails, the principal  $p$  can be convicted of multiply committing the location identified by  $h(p + h(s))$ .

In preparation for the translation, we introduce functions that operate on tuples representing capabilities in the target language. For instance, the function `read` implements source-language reads as a projection, and `check_idc` verifies the seal of committed ids.

$$\begin{aligned} \text{read}(x) &\stackrel{\text{def}}{=} \text{rd}_3(x) \\ \text{get\_idc}(x) &\stackrel{\text{def}}{=} \text{idc}(\text{rd}_1(x), \text{h}(\text{rd}_2(x)) + \text{h}(\text{rd}_2(x) + \text{rd}_3(x)), \text{rd}_4(x)) \\ \text{check\_idc}(x) &\stackrel{\text{def}}{=} \text{verify}(\text{idc}_2(x), \text{idc}_3(x), \text{prin}_1(\text{idc}_1(x))) = \text{ok} \\ \text{get\_idu}(x) &\stackrel{\text{def}}{=} \text{idu}(\text{h}(\text{idc}_1(x) + (+_1 \text{idc}_2(x)))) \end{aligned}$$

In general, inconsistent capabilities may be scattered in the whole system. To detect such inconsistencies and reliably blame cheating principals, a compiled system logs all the committed capabilities generated or received by honest principals by sending them over the channel `log` to the following resolution process  $R$ :

$$\begin{aligned} R = & \text{repl } \text{log?}(y_1). \text{log?}(y_2). \\ & \text{if } \text{check\_idc}(y_1) \text{ and } \text{check\_idc}(y_2) \text{ then} \\ & \text{if } \text{get\_idu}(y_1) = \text{get\_idu}(y_2) \text{ and } \text{idc}_2(y_1) \neq \text{idc}_2(y_2) \text{ then } \text{bad!}(\text{get\_prin}(y_1)) \end{aligned}$$

This resolution process repeatedly reads pairs of `ldc` capabilities over the `log` channel and tests them for inconsistencies, as described above. If cheating is detected, the principal is blamed on channel `bad`. The resolution process acts as an external judge auditing the compiled system, and the data sent over the channel `log` as a secure audit trail. Since all messages on `log` are replicated, log entries cannot be erased or modified by a malicious principal, and every principal may run its own copy of process  $R$ . At the same time, a malicious principal cannot forge capabilities that would accuse an honest principal, as it cannot produce a valid seal associated with the honest principal.

**Translation of initial configurations.** Protocol descriptions can be expressed as initial configurations of a source system that do not contain, or refer to, locations and capabilities; these are created later, during the run of the protocol. We describe the translation of such configurations; a full treatment of capabilities and locations is deferred to Section 5. Our translation is a homomorphism over terms and over most systems.

$$\begin{aligned} \llbracket x \rrbracket &= x & \llbracket c \rrbracket &= c & \llbracket \text{func}(M_1, \dots, M_n) \rrbracket &= \text{func}(\llbracket M_1 \rrbracket, \dots, \llbracket M_n \rrbracket) \\ \llbracket A \rrbracket &= \llbracket A \rrbracket \mid R \mid E & \llbracket a[P] \rrbracket &= \nu m_a . (\llbracket P \rrbracket_a \mid \{ \text{prin}(\text{pk}(m_a)) / a \}) \\ \llbracket A_1 \mid A_2 \rrbracket &= \llbracket A_1 \rrbracket \mid \llbracket A_2 \rrbracket & \llbracket \nu u . A \rrbracket &= \nu u . \llbracket A \rrbracket & \llbracket \{ M / x \} \rrbracket &= \{ \llbracket M \rrbracket / x \} \end{aligned}$$

Let  $\mathcal{A}$  the set of principals running a process in the system and  $\mathcal{E}$  the set of other (possibly dishonest) principals whose names occur in the system ( $\mathcal{E} = \mathcal{P} \cap \text{fn}(A) \setminus \mathcal{A}$ ).

For each principal  $a \in \mathcal{A}$ , the translation creates a secret seed  $m_a$  used to generate the pair of secret/public keys of the principal. The public key is published using an active substitution, while the process run by the principal is compiled within the scope of the private seed  $m_a$  used for signing. Similarly, the translation includes active substitutions  $E = \prod_{e \in \mathcal{E}} (\{ \text{prin}(\text{pk}(m_e)) / e \} \mid \{ H_e / m_e \})$  that records, for each principal  $e \in \mathcal{E}$ , a public key  $\text{pk}(m_e)$  and an associated secret  $H_e$ . The translation also spawns a replicated resolution server  $R$ .

The translation of processes is given next. (We omit the homomorphic clauses for  $0$ ,  $P_1 \mid P_2$ ,  $\text{repl } P$ , and  $\nu c . P$ ).

$$\begin{aligned}
\llbracket \text{newloc } (x, y).P \rrbracket_a &= \nu s'_l . \nu c_l . (c_l! \langle \text{None} \rangle \mid \llbracket P \rrbracket_a \{c_l / c_x\} \{s'_l / s_x\} \{\text{idu}(\text{h}(a + \text{h}(s'_l))) / y\}) \\
\llbracket \text{commit } V x (x').P \rrbracket_a &= c_x?(y).(\llbracket P \rrbracket_a \mid \text{repl } \log! \langle \text{idc}(a, v_x, w_x) \rangle \\
&\quad \{ \text{h}(s_x) + \text{h}(s_x + \llbracket V \rrbracket) / v_x \} \{ \text{sign}(v_x, \text{sk}(m_a)) / w_x \} \{ \text{rd}(a, s_x, \llbracket V \rrbracket, w_x) / x' \} \\
\text{parse } x P &= \\
\text{if is\_rd}(x) = \text{ok then} & \\
\quad \text{if check\_idc}(\text{get\_idc}(x)) \text{ then parse read}(x) (P \mid \text{repl } \log! \langle \text{get\_idc}(x) \rangle) \text{ else } r! \langle \text{None} \rangle & \\
\quad \text{else if is\_idc}(x) = \text{ok then if check\_idc}(x) \text{ then } P \mid \text{repl } \log! \langle x \rangle \text{ else } r! \langle \text{None} \rangle & \\
\quad \text{else if is\_prin}(x) = \text{ok or is\_idu}(x) = \text{ok then } P & \\
\quad \text{else if is\_pair}(x) = \text{ok then parse } (+_1 x) (\text{parse } (+_2 x) P) \text{ else } r! \langle \text{None} \rangle & \\
\llbracket c! \langle M \rangle . P \rrbracket_a &= c! \langle \llbracket M \rrbracket \rangle . \llbracket P \rrbracket_a \\
\llbracket c?(x).P \rrbracket_a &= \nu r . (c?(x). \text{parse } x \llbracket P \rrbracket_a \mid \text{repl } (r?(\_).c?(x). \text{parse } x \llbracket P \rrbracket_a)) \\
\llbracket \text{if } M = M' \text{ then } P_1 \text{ else } P_2 \rrbracket_a &= \text{if } \llbracket M \rrbracket = \llbracket M' \rrbracket \text{ then } \llbracket P_1 \rrbracket_a \text{ else } \llbracket P_2 \rrbracket_a
\end{aligned}$$

The translation of `newloc` creates a fresh location seed  $s'_l$  and a local channel  $c_l$  (with a message `None`, recording that the location is uncommitted), and substitutes  $c_l$  for  $c_x$ ,  $s'_l$  for  $s_x$  and the `idu` capability for  $y$  in the continuation.

The translation of `commit` can proceed only if the location has not been previously committed (the message on  $c_x$  provides mutual exclusion); it then substitutes the `rd` capability for  $x'$  in the continuation code. It also generates the corresponding `idc` capability for the location and logs it by sending it to the resolution protocol.

The `parse` function filters any received value received over channels. If the value is tagged with `rd` or `idc`, then it might (or not) be a valid capability, depending on the validity of its embedded signature: valid capabilities are passed to the continuation, while the associated `idc` is sent to the resolution protocol. If the value is tagged as a principal or an uncommitted capability, it is always passed to the continuation. For compound data, here pairs, each element is separately parsed. Other values, as well as non-valid committed capabilities, are silently discarded. In the translation of an input, we assume that the channel  $r$  is fresh for  $\llbracket P \rrbracket_a$ , and use this channel to loop after discarding such values.

## 5 Model and Translation of Environment Interactions

We define a labeled source semantics that explicitly captures all possible interactions between a system composed of honest principals and an abstract environment composed of potentially hostile principals. To maintain the committable-cell invariants, this semantics keeps track of the capabilities exported to the environment and of the partial knowledge acquired when receiving capabilities from the environment. We then extend our translation from initial configurations to any such reachable configuration.

**Extended location states and capabilities.** We use overlapping syntaxes for capabilities appearing in values, in transition labels, and in the processes representing the state of the cells. Their general form is  $l . \text{Cap}([p][H][V])$ , where  $l$  is the location identifier;  $\text{Cap} \in \{0, \text{Idu}, \text{Idc}, \text{Rd}\}$  is a capability tag;  $p$  is a principal name;  $H$  ranges over terms of the target language; and  $V$  is a value of the source language.



(This syntaxes extend those given in Section 2 for capabilities and location states, with  $l.(a M) = l.0(a M)$ ). The fields  $p$ ,  $H$ , and  $V$  are optional. The presence of a value  $V$  indicates that the location is committed to this value. The term  $H$  plays no role in the source language, but is technically convenient in its translation: it enables us to represent any reachable state of our implementation as the translation of a source system.

The interpretation of  $Cap$  depends on the principal  $p$  that owns the location. If a location is owned by  $a \in \mathcal{A}$ , then  $Cap$  represents the most permissive capability *sent to* the environment (and  $H$  is omitted), with  $Cap = 0$  when no capabilities have been exported so far. If a location is owned by  $e \notin \mathcal{A}$ , then  $Cap$  represents the most permissive capability *received from* the environment (and  $H$  records some opaque cryptographic value in its received representation).

**Ordering capabilities.** We formalize the notion of “more permissive capability” by defining a preorder  $\preceq$  on capabilities. Intuitively,  $C \preceq C'$  holds if  $C$  and  $C'$  have compatible contents and  $C$  can be derived from  $C'$  using the equational theory. We also introduce a special capability  $\perp$  that represents the absence of knowledge on a location. The order is defined by the axioms below:

$$\begin{aligned} \perp \preceq 0 \quad ct \quad 0 \quad ct \preceq \mathbf{Id}u \quad ct \quad \mathbf{Id}u \quad f_u \quad (ct) \preceq \mathbf{Id}c \quad ct \quad \mathbf{Id}c \quad f_c \quad (ct) \preceq \mathbf{R}d \quad ct \\ Cap(p H) \preceq Cap(p H V) \end{aligned}$$

where  $ct$  is any fixed contents and  $f_u$  and  $f_c$  are fixed functions that rewrite  $H$  in  $ct$ . We write  $C \vee C'$  for the sup of  $C$  and  $C'$  with respect to  $\preceq$ , when it exists.

**Normal form.** We say that a system is in *normal form* when it is of the form

$$S = \nu \mathcal{N} (\prod_{l \in \mathcal{L}} l . C_l \mid \prod_{a \in \mathcal{A}} a[P_a] \mid \phi)$$

for some finite sets of names  $\mathcal{N}$ ,  $\mathcal{L}$ , and  $\mathcal{A}$  and active substitutions  $\phi$ . Every initial configuration can be written in normal form (with  $\mathcal{L} = \emptyset$ ) using structural equivalence. A system  $S$  is *well-formed* when it is structurally equivalent to a normal form such that if  $l$  is a location name within  $S$  then  $l \in \mathcal{L}$  and  $l$  occurs only

1. in terms  $l.C$  such that: (a) if  $\text{get\_prin}(l . C_l) \in \mathcal{A}$ , then  $C$  and  $C_l$  are owned by the same principal and if  $C$  has a value, then  $C_l$  has the same value; and (b) if  $\text{get\_prin}(l . C_l) \notin \mathcal{A}$ , then  $C \preceq C_l$  (informally, for a cell owned by the environment, the system cannot have capabilities more permissive than those received);
2. in subprocesses  $\text{commit } M l(x).P$  of  $P_a$  when  $a = \text{get\_prin}(l . C_l)$ ;
3. in  $\mathcal{N}$  when  $\text{get\_prin}(l . C_l) \in \mathcal{A}$  and  $C_l = 0 \quad ct$ .

In the labeled semantics below, we require that the initial and final systems and the label be well-formed. We define labeled transitions  $A \xrightarrow{\alpha} A'$  between source systems on top of an auxiliary relation  $C \xrightarrow{\gamma} C'$  between capabilities.

**Labeled transitions on capabilities.** Input/output actions with the environment can affect the state of memory cells. To model these updates compositionally we define a labeled transition semantics between capabilities.

$$\frac{}{C \xrightarrow{!C'} C \curlywedge C'} \quad \frac{C' \preceq C \quad \text{prin}(C') \in \mathcal{A}}{C \xrightarrow{?C'} C} \quad \frac{\text{prin}(C') \notin \mathcal{A}}{C \xrightarrow{?C'} C \curlywedge C'}$$

The label  $!C'$  records that the capability  $C'$  is exported to the environment: the outcome of the transition  $C \curlywedge C'$  is an updated record of the most permissive exported capability. The label  $?C'$  records that the capability  $C'$  is imported from the environment. There are two import rules, depending on the owner of  $C'$ . If the owner is in  $\mathcal{A}$ , then the capability refers to a location which is part of the system, so the environment can send back at most capabilities that can be derived from those exported by the system, hence the  $C' \preceq C$  condition. On the contrary, if the owner is not in  $\mathcal{A}$ , the environment can send any capability, provided that the capability is compatible with the partial knowledge that the system already has, i.e. that  $C \curlywedge C'$  exists.

**Labeled transitions on systems.** The labeled semantics for systems is adapted from the one for the applied pi calculus. We point out the novelties, and refer to the companion paper for the full semantics.

The labeled semantics has silent steps for all system reductions, including the location-specific reductions described in Section 2. The axioms for input and output are recalled below. (We refer to [1] for a discussion of admissible output values when the equational theory includes cryptographic primitives.)

$$a[c!\langle M \rangle.P] \xrightarrow{c!M} a[P] \quad a[c?(x).P] \xrightarrow{c?M} a[P\{M^\# / x\}]$$

When a capability is received, the rule substitutes in a capability value  $M^\#$  obtained from the capability label  $M$  by erasing information used only to update the cell state.

The context rules below ensure that the communication of capabilities is reflected in the state of the cells of the system; the condition  $l.C$  in  $M$  checks whether the cell  $l.C$  occurs in the transmitted capability (possibly within another capability).

$$\frac{A \xrightarrow{c!M} A' \quad C_0 \xrightarrow{!C} C_1 \quad l.C \text{ in } M}{l.C_0 \mid A \xrightarrow{c!M} l.C_1 \mid A'} \quad \frac{A \xrightarrow{c?M} A' \quad C_0 \xrightarrow{?C} C_1 \quad l.C \text{ in } M}{l.C_0 \mid A \xrightarrow{c?M} l.C_1 \mid A'}$$

$$\frac{A \xrightarrow{\alpha} A' \quad l.C \text{ not in } \alpha}{l.C_0 \mid A \xrightarrow{\alpha} l.C_0 \mid A'}$$

We equate  $l.\perp \mid A$  to  $A$ , so that the input rule covers the case of an input carrying fresh, unknown locations from the environment. (The resulting configuration must be well-formed, which excludes the introduction of a fresh location state for  $l$  if one already exists in the system.) We impose the following well-formedness conditions on labels: (1) in every label, a location name occurs at most in a single, well-formed capability, plus possibly in the label restriction—this excludes e.g. pairs of simultaneous, incompatible commitments; and (2) the target term  $H$ , the principal in uncommitted capabilities, and the value in committed capabilities, appear iff the transition is an input and the capability is owned by  $e \notin \mathcal{A}$ .

**Example of transitions in the source language.** Consider the third round of the game of Section 3, with two honest players  $a_1$  and  $a_2$  and an external, untrusted principal  $e_0 \notin \mathcal{A}$  running the server. A simplified configuration of this system can be written

$$A' = l. \mathbf{Idu} (e_0 H) \mid a_1[c_1?(x_1).P_1] \mid a_2[c_2?(x_2).P_2]$$

where  $l$  is the uncommitted cell pre-allocated by  $e_0$  to store the winning move. (Here  $H = h(e_0 + h(s))$  for some secret  $s$  created by  $e_0$ .) We have possible input transitions on channels  $c_1$  and  $c_2$ , to notify the winning move to each of the players. The first transition may be:

$$A' \xrightarrow{c_1?l. \mathbf{Rd} (e_0 s \mathbf{11})} l. \mathbf{Rd} (e_0 s \mathbf{11}) \mid a_1[P_1\{^l. \mathbf{Rd} (e_0 \mathbf{11})/x_1\}] \mid a_2[c_2?(x_2).P_2]$$

which triggers the final process  $P_1$  with a read capability for  $l$  substituted for  $x_1$ , carrying the game result (here  $\mathbf{11}$ ). At the same time, the state for  $l$  is updated by the third capability-transition rule, since  $\mathbf{Idu} (e_0 H) \vee \mathbf{Rd} (e_0 s \mathbf{11}) = \mathbf{Rd} (e_0 s \mathbf{11})$ . Conversely, for instance, transitions with a label that attributes  $l$  to  $a_1$  instead of  $e_0$  are disabled. At this stage, the configuration records the commitment on  $l$ , so the only subsequent input transition  $A'' \xrightarrow{c_2?l. C'} A'''$  carrying a read capability  $C'$  for  $l$  must be such that  $\mathbf{Rd} (e_0 s \mathbf{11}) \preceq C'$  (by the third capability-transition rule), that is,  $C' = \mathbf{Rd} (e_0 s \mathbf{11})$ . This guarantees that the second player gets exactly the same result as the first one.

**Relating the reduction-based and labeled semantics for the source language.** The labeled semantics precisely characterizes the interactions between a system and an arbitrary environment. Given two systems  $A$  and  $E$  consisting of principals in  $\mathcal{A}$  and  $\mathcal{E}$ , respectively, if  $E \mid A \longrightarrow^* S$  then there exist two such systems  $A'$  and  $E'$  and transitions  $A \xrightarrow{\phi} A'$  such that  $S \equiv \nu \mathcal{N}.(E' \mid A')$ , where  $\mathcal{N}$  is the set of names exported in the labels of  $\phi$ . Conversely, for all systems  $A$  and transitions  $A \xrightarrow{\phi} A'$ , there exists a system  $E'$  and reductions  $E \mid A \longrightarrow^* \nu \mathcal{N}.(E' \mid A')$ .

**Translation of extended location states and capabilities.** We extend the translation of Section 4 to cover all configurations reachable by transitions from initial configurations. This extended translation is inductively defined for all well-formed configurations in normal form, using the clauses of Section 4 plus the rules below for location states and capabilities.

We extensively rely on active substitutions [1] with the following naming conventions: for a location  $l$ ,  $c_l$  denotes the local channel that contains the state of the location,  $s_l$  the secret seed,  $v_l$  the hidden value, and  $w_l$  the seal. We define two extended processes that compute and log identifiers, commitment values, and seals for a location owned by a given principal  $p$  using active substitutions.

$$\begin{aligned} \varphi(M_1, M_2)_p &= \{ h(p + M_1) / l \} \mid \varsigma(M_1, M_2)_p \\ \varsigma(M_1, M_2)_p &= \{ M_1 + M_2 / v_l \} \mid \{ \text{sign}(v_l, \text{sk}(m_p)) / w_l \} \mid \text{repl log!}(\text{idc}(p, v_l, w_l)) \end{aligned}$$

We first translate locations owned by honest principals  $a \in \mathcal{A}$ . The translation implements these locations by sending the location state on the local channel  $c_l$ , activating the relevant substitutions, creating a fresh secret and, for committed locations only, running a replicated log entry:

$$\begin{aligned} [l. 0(a)] &= [[l. \mathbf{Idu}(a)]] = c_l!\langle \text{None} \rangle \mid \{ h(a + h(s_l)) / l \} \mid \nu s. \{ s / s_l \} \\ [l. 0(a \ V)] &= [[l. \mathbf{Idc}(a \ V)]] = [[l. \mathbf{Rd}(a \ V)]] = \varphi(h(s_l), h(s_l + [V]))_a \mid \nu s. \{ s / s_l \} \end{aligned}$$

We also translate locations owned by principals  $e \notin \mathcal{A}$  whose capabilities have been previously received by some principals in  $\mathcal{A}$ . The translation records partial knowledge of these locations, in the form of active substitutions plus, for committed locations only, a replicated log entry. The form of the terms in these substitutions reflect the test that processes in  $\mathcal{A}$  have successfully performed before accepting these values, e.g. that the seal is well-formed signature from  $e$ .

$$\begin{aligned} \llbracket l . \mathbf{Idu} (e H) \rrbracket &= \{ H / l \} \\ \llbracket l . \mathbf{Idc} (e (M' + M'') V) \rrbracket &= \varphi(M', M'')_e \\ \llbracket l . \mathbf{Rd} (e M V) \rrbracket &= \{ M / s_l \} \mid \varphi(\mathbf{h}(M), \mathbf{h}(M + \llbracket V \rrbracket))_e \end{aligned}$$

In a well-formed system, there is a location state for every capability that occurs in the system. Accordingly, the translation of capabilities relies on the active substitutions introduced by the translation of location states, as follows:

$$\llbracket l . \mathbf{Idu} \rrbracket = \text{idu}(l) \quad \llbracket l . \mathbf{Idc} (p) \rrbracket = \text{idc}(p, v_l, w_l) \quad \llbracket l . \mathbf{Rd} (p V) \rrbracket = \text{rd}(p, s_l, \llbracket V \rrbracket, w_l)$$

The compilation of each location state  $l . C$  introduces name  $c_l$  and variables  $s_l, v_l, w_l, l$  whose visibility from the environment depend on the exported capability recorded in  $C$ . Thus, our translation finally introduces the following top-level restrictions: for every location, if no capability have been exported, all these names and variables are restricted; if  $C$  has tag **Idu**, the identifier  $l$  is unrestricted. for locations owned by principals in  $\mathcal{A}$ ; if  $C$  has tag **Idc**, the variables  $w_l$  and  $v_l$  are also unrestricted; if  $C$  has tag **Rd**, only the channel  $c_l$  is restricted.

**Example of transitions in the target language.** Let us consider how our translation operates on the following transition, which represents player  $a_1$  receiving the result of the game from server  $e_0$  (with  $H = h(e_0 + h(s))$ ).

$$l . \mathbf{Idu} (e_0 H) \mid a_1[c_1?(x).P_1] \xrightarrow{c_1?! . \mathbf{Rd} (e_0 s \mathbf{11})} l . \mathbf{Rd} (e_0 s \mathbf{11}) \mid a_1[P_1\{^l . \mathbf{Rd} (e_0 \mathbf{11})/x\}]$$

The translated system  $\{ H / l \} \mid \llbracket a_1[c_1?(x).P_1] \rrbracket$  simulates the source transition by an input with label  $c_1 ? (\text{rd}(e_0, s, \mathbf{11}, \text{sign}(\mathbf{h}(s) + \mathbf{h}(s + \mathbf{11}), \text{sk}(m_{e_0})))$ , followed by a series of reductions through the code of parse, including dynamic checks on `is_rd` and `check_idc`. In 6 silent steps (including 3 steps for recursive processing of value  $\mathbf{11}$ ), this yields the process

$$\begin{aligned} &\{ H / l \} \mid \llbracket a_1[P_1] \rrbracket \{ \text{rd}(e_0, s, \mathbf{11}, \text{sign}(\mathbf{h}(s) + \mathbf{h}(s + \mathbf{11}), \text{sk}(m_{e_0}))) / x \} \\ &\mid \text{repl } \log! \{ \text{get\_idc}(x) \} \mid \nu r . (\text{repl } r?(\_).c_1?(x). \text{parse } x \llbracket P \rrbracket_a) \mid R \mid E. \end{aligned}$$

After applying structural equivalence with active substitutions and eliminating the dead loop on channel  $r$ , we obtain a system

$$\nu s_l . \nu v_l . \nu w_l . (\{ s / s_l \} \mid \varphi(\mathbf{h}(s_l), \mathbf{h}(s_l + \mathbf{11}))_{e_0} \mid \llbracket a_1[P_1] \rrbracket \{ \text{rd}(e_0, s_l, \mathbf{11}, w_l) / x \}) \mid R \mid E$$

that matches the translation of the resulting source system above.

## 6 Correctness Results

The first proposition states that the behavior of every source system can be simulated by its translation. That is, for any labeled trace of all source systems, there is a labeled trace of the process resulting from its translation. This shows the correctness (or functional adequacy) of our translation. We let  $\xrightarrow{\phi}$  (resp.  $\xrightarrow{\psi}$ ) range over series of transitions in the labeled semantics of the source (resp. target) language.

**Theorem 1 (Functional adequacy).** *Let  $A$  be a well-formed source system.*

*For all series of transitions  $A \xrightarrow{\phi^*} A'$ , there exist transitions  $\llbracket A \rrbracket \xrightarrow{\psi^*} \llbracket A' \rrbracket$ .*

The proof of the theorem is by induction on a series of source transitions between systems in normal forms. For each source transition, we exhibit target transitions that commute with the translation.

The “upwards” direction is more challenging: the trace produced by the translation of a source process  $A$  can be related to a trace produced by  $A$  *unless* its translation emits the name of a cheating principal on the special channel *bad*. This property uniformly guarantees the security of the translation of all systems with respect to the source semantics, provided that a proof that a principal cheated is a reasonable exceptional outcome for the other principals.

We let  $S \xrightarrow{*}_D S'$  denote that a target system  $S$  goes to  $S'$  with a (possibly empty) series of silent deterministic transitions, and let  $S \Downarrow M$  abbreviate  $S \xrightarrow{*}_D \xrightarrow{bad!M} S'$  for some  $S'$ ; we then say that  $M$  is blamed.

**Theorem 2 (Security).** *For all transitions  $\llbracket A \rrbracket \xrightarrow{\psi^*} S$  starting from a well-formed source system  $A$ , we have*

1. *either there are source transitions  $A \xrightarrow{\phi^*} A'$  leading to a well-formed source system  $A'$  such that  $S \xrightarrow{*}_D \llbracket A' \rrbracket$ ; or  $S \Downarrow e$  for some  $e \notin \mathcal{A}$ ;*
2. *if  $S \Downarrow M$ , then  $M \notin \mathcal{A}$ .*

The proof is by induction on the series of transitions in the target language that do not trigger a blame. The first part of the theorem states that either the source semantics is respected, or the implementation at least provides the honest participants with the name of one dishonest principal to blame. Said otherwise, its statement excludes the possibility of cheating without eventual detection. The second part of the theorem expresses that honest participants are never blamed (even in the case some dishonest participants cheat), a necessary property for any optimistic implementation.

The form of our theorem differs from security properties for other programming abstractions (e.g. [7,2]), where any run or labeled trace of the cryptographic implementation of a source program is related to a run or labeled trace of the program on the source level. Reflecting a more flexible approach to security, it enables bad runs as long as malicious principals are reliably detected and blamed.

We illustrate how the Resolution protocol and the verifications made by the translation of receive suffice to detect write-after-commit attacks. Consider the online game example and suppose that  $a_1, a_2 \in \mathcal{A}$  and  $e_0 \notin \mathcal{A}$ , that is, the server implementation is malicious. In particular, the server implementation may commit location  $l$  twice, to

convince  $a_1$  that he is the winner with his bid 11 and  $a_2$  that he is the winner with his bid 8. The system composed by the translation of the two clients  $\llbracket A_1 \mid A_2 \rrbracket$  generates a trace

$$\llbracket A_1 \mid A_2 \rrbracket \rightarrow \dots \rightarrow \llbracket A' \rrbracket \xrightarrow{c_1 ? (\text{rd}(e_0, s, \mathbf{11}, w))} c_2 ? (\text{rd}(e_0, s, 8, w')) \mathcal{S}$$

where the seals  $w$  and  $w'$  sign commitments of  $l$  to 11 and 8, respectively.

For the first input transition, there exists a matching source transition, with a resulting source system  $A''$  that includes the location state  $l$ .  $\mathbf{Rd}(e_0 \ s \ \mathbf{11})$ . Moreover, the translation of  $A''$  emits the corresponding  $\text{idc}$  on  $\text{log}$ .

For the second input transition, however, there is no matching source transition. This would require a capability transition from  $\mathbf{Rd}(e_0 \ s \ \mathbf{11})$  to  $\mathbf{Rd}(e_0 \ s \ 8)$ , which is excluded by our definition of the  $\preceq$  preorder. Instead, the resulting system sends a second  $\mathbf{Idc}$  on  $\text{log}$ . As soon as the Resolution process reads both commitments, it detects that they are inconsistent, and blames  $e_0$  on *bad*.

## 7 Prototype Implementation

We have implemented committable cells as a library for OCaml [9]. We have also coded a series of examples, including simple online games and sealed-bid auctions.

The library provides abstract datatypes and access functions that closely follow those of our source language. Its implementation relies on standard cryptographic libraries and on a public-key (X.509) infrastructure for processing capabilities; it uses pseudo-random number generation for creating fresh secret seeds. Programs that use our library may communicate with one another using OCaml marshalling and network socket interfaces—cryptographic validation of received capabilities then occurs during unmarshalling.

The main difference between the implementation and its formal semantics is the handling of resolution. We refine the idealized resolution mechanism of Section 4 as follows: instead of relying on a central resolution process, our implementation keeps track of all principals and cells involved in a run of the system, and eventually implements the exchange and local resolution for all shared commitments.

## 8 Conclusions and Future Work

We presented a simple language for specifying systems based on optimistic commitments, and we compiled this language into a realistic concurrent framework modeled in the applied pi calculus. We established two security properties relating the labeled traces of a source semantics with commitment primitives to those of their implementation, with a target semantics that uses only ordinary communications and cryptographic functions. We only consider authenticity for now, but we believe it would also be possible to guarantee some properties of formal secrecy.

Although committable cells provide a reasonably useful (and formally challenging) block for building protocols, we focused on one particular usage of secure logs, rather than proposing a comprehensive language design for optimistic protocols. Our formal approach could be extended to other, more involved datatypes—as long as we can represent their live cycles using a preorder on exported capabilities, as detailed in Section 5.

It would be interesting, for instance, to design compilers for such datatypes with incremental commitment properties.

More generally, audit logs constitute an important tool for designing protocols and applications. Although their efficient implementation has been thoroughly studied, we believe ours is the first work to address their reliable, principled usage from a programming-language viewpoint.

## References

1. Abadi, M., Fournet, C.: Mobile values, new names, and secure communication. In: 28th ACM Symposium on Principles of Programming Languages (POPL 2001) (2001)
2. Abadi, M., Fournet, C., Gonthier, G.: Secure Implementation of Channel Abstractions. *Information and Computation* 174(1), 37–83 (2002)
3. Castellà-Roca, J., Domingo-Ferrer, J., Riera, A., Borrell, J.: Practical Mental Poker Without a TTP Based on Homomorphic Encryption. In: Johansson, T., Maitra, S. (eds.) *INDOCRYPT 2003*. LNCS, vol. 2904, pp. 280–294. Springer, Heidelberg (2003)
4. Cederquist, J.G., Corin, R., Dekker, M.A.C., Etalle, S., den Hartog, J.I., Lenzini, G.: Audit-based compliance control. *Int'l Journal of Information Security* 6(2), 133–151 (2007)
5. Chaum, D.: Secret-ballot receipts: True voter-verifiable elections. *IEEE Security and Privacy* 2(1), 38–47 (2004)
6. Chaum, D., Ryan, P.Y.A., Schneider, S.: A practical, voter-verifiable election scheme. Technical Report CS-TR-880 (2004)
7. Corin, R., Denielou, P.-M., Fournet, C., Bhargavan, K., Leifer, J.: Secure implementations for typed session abstractions. In: *IEEE Computer Security Foundations Symposium* (2007)
8. Corin, R., Galindo, D., Hoepman, J.H.: Securing data accountability in decentralized systems. In: Meersman, R., Tari, Z., Herrero, P. (eds.) *OTM 2006 Workshops*. LNCS, vol. 4277, pp. 626–635. Springer, Heidelberg (2006)
9. Leroy, X., et al.: Objective caml, <http://caml.inria.fr>
10. Etalle, S., Winsborough, W.H.: A posteriori compliance control. In: 12th ACM Symposium on Access Control Models and Technologies (2007)
11. ISO/IEC. Common criteria for information technology security evaluation (2004), <http://www.commoncriteriaportal.org/public/expert/index.php?menu=3>
12. Jha, S., Katzenbeisser, S., Schallhart, C., Veith, H., Cheney, S.: Enforcing semantic integrity on untrusted clients in networked virtual environments. In: *IEEE Symposium on Security and Privacy* (2007)
13. Kremer, S., Ryan, M.D.: Analysing the vulnerability of protocols to produce known-pair and chosen-text attacks. In: 2nd Int'l Workshop on Security Issues in Coordination Models, Languages and Systems (SecCo 2004). ENTCS (2005)
14. NIST Special Publications. Generally accepted principles and practices for securing information technology systems (1996)
15. Schneier, B., Kelsey, J.: Secure audit logs to support computer forensics. *ACM Transactions on Information and System Security* 2(2), 159–176 (1999)
16. Shamir, A., Rivest, R., Adleman, L.: Mental poker. *Mathematical Gardener* (1981)
17. Waters, B.R., Balfanz, D., Durfee, G., Smetters, D.K.: Building an encrypted and searchable audit log. In: *Network and Distributed System Security Symposium (NDSS)* (2004)
18. Xu, W., Chadwick, D., Otenko, S.: A PKI Based Secure Audit Web Server. In: *IASTED Communications, Network and Information and CNIS* (2005)
19. Zheng, L., Chong, S., Myers, A.C., Zdancewic, S.: Using replication and partitioning to build secure distributed systems. In: *IEEE Symposium on Security and Privacy* (2003)