Journal of Computer Security 31 (2023) 615–645
DOI 10.3233/JCS-230048
IOS Press

Flow-limited authorization for consensus, replication, and secret sharing[1]

Priyanka Mondal [a,*], Maximilian Algehed [b] and Owen Arden [c]
[a] *University of California, Santa Cruz, CA, USA*
E-mail: pmondal@ucsc.edu
[b] *Chalmers University of Technology, Gothenburg, Sweden*
E-mail: algehed@chalmers.se
[c] *University of California, Santa Cruz, CA, USA*
E-mail: owen@soe.ucsc.edu

AUTHOR COPY

Abstract. Availability is crucial to the security of distributed systems, but guaranteeing availability is hard, especially when participants in the system may act maliciously. Quorum replication protocols provide both integrity and availability: data and computation is replicated at multiple independent hosts, and a quorum of these hosts must agree on the output of all operations applied to the data. Unfortunately, these protocols have high overhead and can be difficult to calibrate for a specific application's needs. Ideally, developers could use high-level abstractions for consensus and replication to write fault-tolerant code that is secure by construction.

This paper presents Flow-Limited Authorization for Quorum Replication (FLAQR), a core calculus for building distributed applications with heterogeneous quorum replication protocols while enforcing end-to-end information security. Our type system ensures that well-typed FLAQR programs cannot *fail* (experience an unrecoverable error) in ways that violate their type-level specifications. We present noninterference theorems that characterize FLAQR's confidentiality, integrity, and availability in the presence of consensus, replication, and failures, as well as a liveness theorem for the class of majority quorum protocols under a bounded number of faults. Additionally, we present an extension to FLAQR that supports secret sharing as a form of declassification and prove it preserves integrity and availability security properties.

Keywords: Information flow control, fault tolerant systems, quorum replication, language-based security, distributed systems

1. Introduction

Failure is inevitable in distributed systems, but its consequences may vary. The consequences of failure are particularly severe in centralized system designs, where single points-of-failure can render the entire system inoperable. Even distributed systems are sometimes built using a single, centralized authority to execute security-critical tasks. If this trusted entity is compromised, the security of the entire system may be compromised as well.

Building reliable *decentralized systems*, which have no single point-of-failure, is a complex task. Quorum replication protocols such as Paxos [16] and PBFT [7], and blockchains such as Bitcoin [21] replicate state and computation at independent hosts and use consensus protocols to ensure the integrity and availability of operations on system state. In these protocols, there is neither centralization of function

[1]This paper is an extended and revised version of a paper presented at CSF'22.
[*]Corresponding author. E-mail: pmondal@ucsc.edu.

0926-227X/$35.00 © 2023 – IOS Press. All rights reserved.

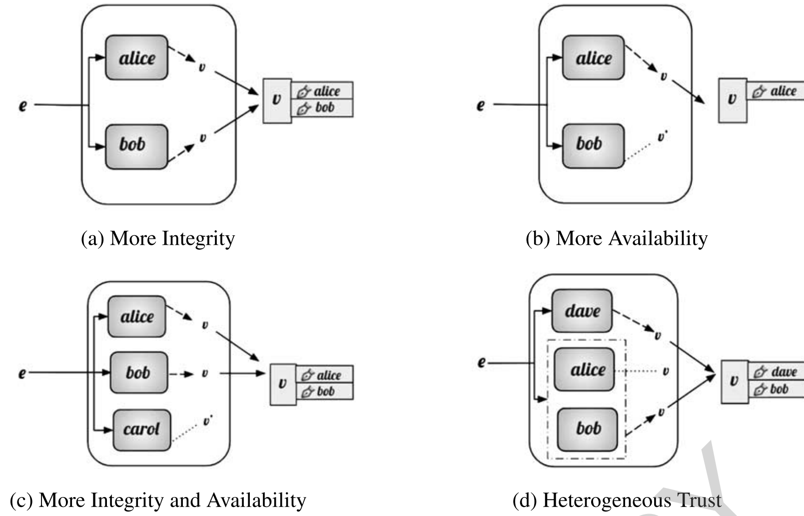

(a) More Integrity

(b) More Availability

(c) More Integrity and Availability

(d) Heterogeneous Trust

Fig. 1. Integrity-availability trade-off.

AUTHOR COPY

nor centralization of trust: all honest hosts work to replicate the same computation on the same data, and this redundancy helps the system tolerate a bounded number of node failures and corruptions.

Within a single trust domain such as a corporate data center, replicas likely have uniform trust relationships and may be treated interchangeably. However, many large-scale systems depend on services hosted by multiple external services. Even when a service's internal components are replicated, developers must take into account the failure properties of external dependencies when considering their own robustness.

Information flow control (IFC) has been used to enforce decentralized security in distributed systems for confidentiality and integrity (e.g., Fabric [18] and DStar [28]). Less attention has been paid to enforcing decentralized availability policies with IFC. In particular, no language (or protocol) we are aware of addresses systems that compose multiple quorums or consider quorum participants with arbitrary trust relationships.

To build a formal foundation for such languages, we present FLAQR, a core calculus for Flow-Limited Authorization [3] for Quorum Replication. FLAQR uses high-level abstractions for replication and consensus that help manage tradeoffs between the availability and integrity of computation and data.

Consider the scenarios in Fig. 1. Shaded boxes represent hosts in a distributed system. Dashed lines denote outputs that contribute to the final result, a value v. Dotted lines denote ignored outputs and solid lines indicate the flow of data from an initial expression e distributed to hosts to the collected result. Results are accompanied by labels that indicate which hosts influenced the final result.

In Fig. 1(b), e is distributed to hosts `alice` and `bob`. The hosts' results are compared and, if they match, the result is produced. Since a value is output only if the values match, we can treat the output of this protocol as having *more integrity* than just `alice` or `bob`. While both `alice` and `bob` technically influence the output, neither host can unilaterally control its value. However, either host can cause the protocol to fail.

By contrast, the protocol in Fig. 1(a) prioritizes availability over integrity: if either `alice` or `bob` produce a value, the protocol outputs a value – in this case `alice`'s. Here, neither host can unilaterally cause a failure; the protocol only fails if both `alice` and `bob` fail. Either `alice` or `bob` (but not both) has complete control over the result in the event of the other's failure, so we should treat the output as having *less integrity* than just `alice` or `bob`.

With an adequate number of hosts, we can combine these two techniques to form the essential components of a quorum system. In Fig. 1(c), e is replicated to `alice`, `bob`, and `carol`. This protocol outputs a value if any two hosts have matching outputs. Since `alice` and `bob` both output v, the protocol outputs v and attaches `alice` and `bob`'s signatures. The non-matching value v' from `carol` is ignored. Hence, this protocol prevents any single host from unilaterally controlling the failure of the protocol or its output.

Figure 1(c) is similar in spirit to consensus protocols such as Paxos or PBFT where quorums of independent replicas are used to tolerate a bounded number of failures. FLAQR also permits us to write protocols where principals have differing trust relationships. Figure 1(d) illustrates a protocol that tolerates failure (or corruption) of either `alice` or `bob`, but requires `dave`'s output to be part of any quorum. This protocol will fail if both `alice` and `bob` fail to produce matching outputs, but can also fail if `dave` fails to produce a matching output. This example illustrates the distributed systems where the hosts do not have homogeneous trust.

The main contributions of this paper are:

- An extension of the static fragment of the Flow Limited Authorization Model (FLAM) [3] with availability policies and algebraic operators representing the effective authority of consensus and replication protocols (Section 3–Section 5).
- A formalization of the FLAQR language (Section 4) and accompanying results:
  - ∗ A liveness theorem for majority-quorum FLAQR protocols (Section 7.1) which experience a bounded number of faults using a novel proof technique: a *blame semantics* that associates failing executions of a FLAQR program with a set of principals who may have caused the failure.
  - ∗ Noninterference theorems for confidentiality, integrity, and availability (Section 7.2).
  - ∗ An extension to FLAQR adding support for simple secret sharing, and results demonstrating it preserves integrity and availability noninterference as well as our liveness theorem, despite introducing an additional source of failure due to mismatched shares (Section 9).

This paper is an expanded and updated version of an article previously published in the proceedings of the 35th Computer Security Foundations Symposium [19]. This version adds support for secret sharing (Section 9) and extensions of our previous results that demonstrate these new terms neither impact integrity and availability noninterference, nor majority liveness, despite introducing an additional source of failure. In addition, we corrected a minor issue in the original blame semantics, and include complete rule sets and proofs for our formalization and theoretical results.

Non-goals. The design of FLAQR is motivated by application-agnostic consensus protocols such as Paxos [16] and PBFT [7], but our present goal is not to develop a framework for *verifying* implementations of such protocols (although it would be interesting future work). Rather, the goal is to develop security abstractions that make it easier to create components with application-specific integrity and availability guarantees, and compose them in a secure and principled way.

In particular, the FLAQR system model lacks some features that a protocol verification model would require, most notably a concurrent semantics, asynchronous message delivery, and arbitrary communication patterns. Although this simplifies some aspects of consensus protocols, our model retains many of the core challenges present in fault tolerance models. For example, perfect fault detection is impossible and faulty hosts can manipulate data to cause failures to manifest at other hosts. We argue that even in a synchronous, deterministic model with RPC-style communication, the challenges of specifying and enforcing policies remain quite difficult to solve, and are among the primary security concerns of high-level application developers.

AUTHOR COPY

```
1  getBalance(acct):
2      bal_a = fetch bal(acct) @ alice;
3      bal_b = fetch bal(acct) @ bob;
4      bal_c = fetch bal(acct) @ carol;
5
6      if (bal_a==bal_b && bal_a != fail)
7         return bal_a;
8      else if (bal_b==bal_c && bal_b != fail)
9         return bal_b;
10     else if (bal_c==bal_a && bal_c != fail)
11        return bal_c;
12     else return fail;
```

Fig. 2. Majority quorum.

2. Motivating examples

In this section we present two motivating examples. The first example highlights the trade-off between integrity and availability. The second example highlights the need for availability policies in distributed systems.

2.1. Tolerating failure and corruption

If a bank's deposit records are stored in a single node, then customers will be unable to access their accounts if that node is unavailable or is compromised. To eliminate this single point-of-failure, banks can replicate their records on multiple hosts as illustrated in Fig. 1(c). If a majority of hosts agree on an account balance, then the system can tolerate the remaining minority of hosts failing or returning corrupted results.

Consider a quorum system with three hosts: `alice`, `bob`, and `carol`. To tolerate the failure of a single node, balance queries attempt to contact all three hosts and compare the responses. As long as the client receives two responses with the same balance, the client can be confident the balance is correct even if one node is compromised or has failed.

Figure 2 illustrates a pseudocode implementation of `getBalance` in this system. The code fetches balances from the three hosts (lines 2–4). The function returns the balance if each fetched value matches, otherwise the function returns `fail` (lines 6–12).

The downside of this approach is that it is quite verbose and repetitive compared to a single-line fetch without any fault tolerance. Small mistakes in any of these lines could have significant consequences. For example, suppose a programmer typed `bal_b` instead of `bal_c` on line 8. This small change gives `bob` (or an attacker in control of `bob`'s node) the ability to unilaterally choose the return value of the function, even when `alice` and `carol` agree on a different value.

2.2. Using best available services

Real world applications often consist of communication between entities with mutual distrust. The pseudocode in Fig. 3 communicates with two banks, represented by `b` and `b′`, during a distributed computation. A user has two accounts `acc_1`, and `acc_2` with `b` and `b′` respectively. The user has linked both accounts to a service and specifies the bill should be paid

(1) as long as at least one account is available
(2) using the highest-balance account, if available

AUTHOR COPY

```
1  highestBalance(acct_1, acct_2):
2    bal_1:= fetch getBalance(acct_1) @ b;
3    bal_2:= fetch getBalance(acct_2) @ b';
4
5    if (bal_1==fail) && (bal_2==fail) then
6       return fail;
7    else if (bal_1==fail) then
8       return acct_2;
9    else if (bal_2==fail) then
10      return acct_1;
11
12   if (bal_1 > bal_2) then
13       return acct_1;
14   else
15       return acct_2;
```

Fig. 3. Available largest balance.

Lines 7–10 take care of point (1), ensuring the comparison on line 12 does not get stuck if a `fetch` returns `fail`. Lines 12–15 cover point (2), returning the account with the highest balance when both balances are available.

This example shows how availability of data can effect the final result of an application and thus highlights the importance of enforcing availability in distributed computations. As in the previous example, the programmer must reason about failures due to unavailable hosts and make the correct comparisons to implement the (implicitly) desired policy. Furthermore, the programmer may be unaware of the availability guarantees offered by b and b′. For example, if b and b′ rely on the same replicas to implement `getBalance`, the availability of `highestBalance` may be lower than expected.

Finally, in both of the above examples, an attacker should not be able to read an account balance, or infer which account balance was greater. With the FLAQR type-system, programmers can not only specify and enforce availability and integrity, but also confidentiality – crucial for dealing with sensitive information. Moreover, FLAQR enables programmers to write fault-tolerant code concisely, with explicit primitives for consensus and replication operations that clarify the programmer's intentions.

AUTHOR COPY

3. Specifying FLAQR policies

FLAQR policies are specified using an extension of the FLAM [2,3] principal algebra that includes availability policies.[2] FLAM principals represent both the *authority* of entities in a system as well as bounds on the *information flow policies* that authority entails. For example, Alice's authority is represented by the principal `alice`. *Authority projections* allow us to refer to specific categories of Alice's authority. The principal $\mathtt{alice}^{c}$ refers to Alice's confidentiality authority: what Alice may read. Principal $\mathtt{alice}^{i}$ refers to Alice's integrity authority: what Alice may write or influence.[3] Principal $\mathtt{alice}^{a}$ refers to her availability authority: what Alice may cause to *fail*. A principal always acts for any projection of its authority, so for example $\mathtt{alice} \succcurlyeq \mathtt{alice}^{a}$. We refer to the set of all *primitive principals* such as `alice` and `bob` as $\mathcal{N}$.

We can write the conjunction of two principals with the Boolean connective $\wedge$ as $\mathtt{alice} \wedge \mathtt{bob}$, denoting the combined authority of Alice and Bob. Put another way, $\mathtt{alice} \wedge \mathtt{bob}$ is a principal both Alice and

[2] Specifically, we extend the static fragment of FLAM's principal algebra defined by FLAC [2].

[3] Prior FLAM-based formalizations have used $\rightarrow$ and $\leftarrow$ for confidentiality and integrity, respectively.

Bob *trust*. The disjunction of two principals' authority is written using the connective $\vee$ as $\texttt{alice}\vee\texttt{bob}$. This is a principal whose authority is less than both Alice and Bob; either Alice or Bob can act on behalf of the principal $\texttt{alice}\vee\texttt{bob}$. Put another way $\texttt{alice}\vee\texttt{bob}$ is a principal that trusts both Alice and Bob. Authority projections distribute over $\wedge$ and $\vee$, so for example $(\texttt{alice}\wedge\texttt{bob})^{\texttt{i}}\equiv\texttt{alice}^{\texttt{i}}\wedge\texttt{bob}^{\texttt{i}}$.

The confidentiality, integrity, and availability authorities make up the totality of a principal's authority, so writing $\texttt{alice}^{\texttt{c}}\wedge\texttt{alice}^{\texttt{i}}\wedge\texttt{alice}^{\texttt{a}}$ is equivalent to writing $\texttt{alice}$. For brevity, we sometimes write $\texttt{alice}^{\texttt{ci}}$ as a shorthand for $\texttt{alice}^{\texttt{c}}\wedge\texttt{alice}^{\texttt{i}}$ when we wish to include all but one kind of authority. Our complete formalization of the acts-for relation is presented in Figs 39 and 40 in Appendix A.

In addition to conjunctions and disjunctions of authority, FLAQR also introduce two new operators: *partial conjunction* ($\boxplus$), and *partial disjunction* ($\boxminus$). These operations are necessary to represent the tradeoffs between integrity and availability mediated by consensus and replication. Consider the "more integrity" protocol from Fig. 1(b). It is reasonable to think of the consensus value v as having more integrity than (or at least, "not less integrity than") Alice or Bob alone, but it turns out to be useful to distinguish between this authority and the combined integrity authority of Alice and Bob, $(\texttt{alice}^{\texttt{i}}\wedge\texttt{bob}^{\texttt{i}})$. A principal with integrity authority $(\texttt{alice}^{\texttt{i}}\wedge\texttt{bob}^{\texttt{i}})$ may act arbitrarily on behalf of both Alice and Bob since it is trusted by them. In contrast, the integrity of the value produced in Fig. 1(b) is *not* fully trusted by Alice and Bob. Instead, Alice and Bob only trust the value when Alice and Bob agree on it. If they do not agree, that trust is revoked and no value is produced. For this reason, we describe the integrity of consensus values such as v as the *partial conjunction* of Alice and Bob, written $\texttt{alice}^{\texttt{i}}\boxplus\texttt{bob}^{\texttt{i}}$.

Similarly, for replication protocols like that in Fig. 1(a), we want to distinguish the integrity of values that may have been received from either Alice or Bob due to failure, from the integrity of values that may have been influenced by both Alice and Bob: $\texttt{alice}^{\texttt{i}}\vee\texttt{bob}^{\texttt{i}}$. The integrity of a value produced by either Alice or Bob is written as the *partial disjunction* $\texttt{alice}^{\texttt{i}}\boxminus\texttt{bob}^{\texttt{i}}$. This principal does not have the same integrity authority as Alice or Bob alone since we cannot guarantee which host's value will be used in the event of a failure. However, the value does have more integrity than $\texttt{alice}^{\texttt{i}}\vee\texttt{bob}^{\texttt{i}}$, since *only* Alice or Bob (and not both) may have influenced it.

We compare the authority of principals using the *acts-for* relation $\succcurlyeq$, which partially orders (equivalence classes of) principals by increasing authority. We form the set of all principals $\mathcal{P}$ as the closure of the set $\mathcal{N}\cup\{\top,\bot\}$ over the operations $\wedge$, $\vee$, $\boxplus$, $\boxminus$, and authority projections $\texttt{c}$, $\texttt{i}$, and $\texttt{a}$. We say Alice acts for Bob (or equivalently, Bob trusts Alice) and write $\texttt{alice}\succcurlyeq\texttt{bob}$ when Alice has at least as much authority as Bob. The $\succcurlyeq$ relation forms a lattice with join $\wedge$, meet $\vee$, greatest element $\top$, and least element $\bot$.

In addition to the trust relationships such as $p\wedge q\succcurlyeq p$ and $p\succcurlyeq p^{\texttt{i}}$ implied by the principal algebra, explicit delegations of trust such as $p\succcurlyeq q$ (for any p, q in $\mathcal{P}$) may expressed using a *delegation context* Π. An acts-for judgment has the form $\Pi\Vdash p\succcurlyeq q$ and means that p acts for q in context Π. While FLAC has a feature that allows dynamic extensions of Π, for simplicity we fix Π to a static set of delegations in FLAQR.

We extend the acts-for relation defined by Arden et al. [2] with new rules for availability authority and partial conjunction and disjunction. Figure 4 presents a selection of these rules – we have omitted the distributivity rules for brevity. In Fig. 4, an *acts for* judgement of form $\Pi\Vdash p\succcurlyeq q$, states that p has at least as much authority as q in delegation context Π. Figures 39 and 40 in Appendix A together present the complete formalization of the $\succcurlyeq$ relation. In Fig. 40 we present only the extensions to this relation introduced by FLAQR.

AUTHOR COPY

$$[\textsc{PAndL}]\ \frac{\Pi \Vdash p_i \succcurlyeq p \qquad k \in \{1,2\}}{\Pi \Vdash p_1 \boxplus p_2 \succcurlyeq p} \qquad [\textsc{PAndR}]\ \frac{\Pi \Vdash p \succcurlyeq p_1 \qquad \Pi \Vdash p \succcurlyeq p_2}{\Pi \Vdash p \succcurlyeq p_1 \boxplus p_2} \qquad [\textsc{AndPAnd}]\ \Pi \Vdash p \wedge q \succcurlyeq p \boxplus q$$

$$[\textsc{PAndPOr}]\ \Pi \Vdash p \boxplus q \succcurlyeq p \boxminus q \qquad [\textsc{ProjPAndL}]\ \Pi \Vdash p^{\pi} \boxplus q^{\pi} \succcurlyeq (p \boxplus q)^{\pi}$$

$$[\textsc{ProjPAndR}]\ \Pi \Vdash (p \boxplus q)^{\pi} \succcurlyeq p^{\pi} \boxplus q^{\pi} \qquad [\textsc{ProjPOrL}]\ \Pi \Vdash p^{\pi} \boxminus q^{\pi} \succcurlyeq (p \boxminus q)^{\pi}$$

$$[\textsc{ProjPOrR}]\ \Pi \Vdash (p \boxminus q)^{\pi} \succcurlyeq p^{\pi} \boxminus q^{\pi} \qquad [\textsc{POrOr}]\ \Pi \Vdash p \boxminus q \succcurlyeq p \vee q$$

Fig. 4. Selected acts-for rules for partial conjunction and disjunction.

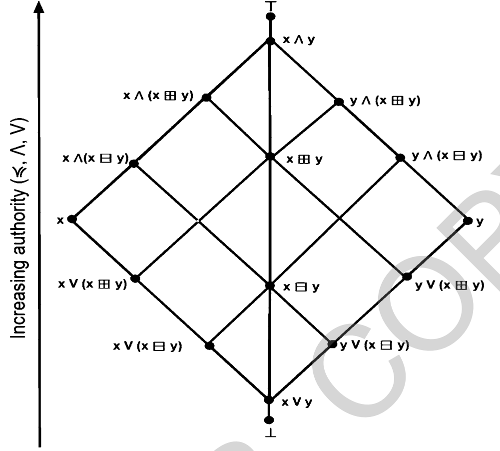

Fig. 5. The FLAQR authority lattice for the principal set $\{\bot, x, y, \top\}$.

AUTHOR COPY

As a consequence of these new acts-for rules we have additional distinct points in the authority lattice. Figure 5 illustrates the authority sublattice over elements $\{\bot, x, y, \top\}$. Figure 5 shows the trust ordering of all possible distinct combinations of elements that can be formed on the set $\{\bot, x, y, \top\}$ with operations $\wedge$, $\vee$, $\boxplus$ and $\boxminus$ over them. The relationship between principals $\bot, x, y, x \wedge y, x \vee y$, and $\top$ is the same as in FLAM, but Fig. 5 also includes principals constructed using partial conjunctions and disjunctions. For example, $x \wedge (x \boxplus y)$ is the least upper bound of $x \wedge (x \boxminus y)$ and $x \boxplus y$. This is due to rule PAndPOr in Fig. 4, which lets us simplify $x \wedge (x \boxminus y) \wedge x \boxplus y$ to $x \wedge (x \boxplus y)$.

To compare the restrictiveness of information flow policies, we use the *flows-to* relation $\sqsubseteq$, which partially orders principals by increasing policy restrictiveness, rather than by authority. For example, we say Alice's integrity flows to Bob's integrity and write $\texttt{alice}^{\texttt{i}} \sqsubseteq \texttt{bob}^{\texttt{i}}$ if Bob trusts information influenced by Alice at least as much as information he influenced himself. Likewise, we write $\texttt{alice}^{\texttt{c}} \sqsubseteq \texttt{bob}^{\texttt{c}}$ if Alice trusts Bob to protect the confidentiality of her information, and $\texttt{alice}^{\texttt{a}} \sqsubseteq \texttt{bob}^{\texttt{a}}$ if Bob is trusted to keep Alice's data available. The flows-to relation behaves similarly to a sub-typing relation. Treating information labeled $\texttt{alice}^{\texttt{cia}}$ (i.e. $\texttt{alice}$) as though it was labeled $\texttt{bob}^{\texttt{cia}}$ (i.e. $\texttt{bob}$) is only safe (doesn't violate anyone's policies) if $\texttt{alice}^{\texttt{cia}} \sqsubseteq \texttt{bob}^{\texttt{cia}}$ (i.e. $\texttt{alice} \sqsubseteq \texttt{bob}$).

One advantage of the FLAM principal algebra is that we can define the flows-to relation, as well as the upper and lower bounds of information flow policies, in terms of the acts-for relation, simplifying our formalism.

$$p \sqsubseteq q \Leftrightarrow q^{\texttt{c}} \succcurlyeq p^{\texttt{c}} \quad \text{and} \quad p^{\texttt{i}} \succcurlyeq q^{\texttt{i}} \quad \text{and} \quad p^{\texttt{a}} \succcurlyeq q^{\texttt{a}}$$

$\pi \in \{\mathtt{c}, \mathtt{i}, \mathtt{a}\}$ (projections)
$n \in \mathcal{N}$ (primitive principals)
$x \in \mathcal{V}$ (variable names)

$$
\begin{aligned}
p, \ell, pc &::= n \mid \top \mid \bot \mid p^{\pi} \mid p \wedge p \mid p \vee p \\
&\quad \mid p \sqcup p \mid p \sqcap p \mid p \boxminus p \mid p \boxplus p \\
\tau &::= \mathtt{unit} \mid X \mid (\tau + \tau) \mid (\tau \times \tau) \\
&\quad \mid \tau \xrightarrow{pc} \tau \mid \forall X[pc].\, \tau \mid \ell \; \mathtt{says} \; \tau \\
v &::= () \mid \underline{(\overline{\eta}_{\ell}\; v)} \mid \mathtt{inj}_i^{(\tau+\tau)}\; v \mid \langle v, v \rangle^{\tau} \\
&\quad \mid \lambda(x{:}\tau)[pc].\, e \mid \Lambda X[pc].\, e \\
f &::= v \mid \underline{\mathtt{fail}^{\tau}} \\
e &::= f \mid x \mid e\; e \mid e\; \tau \mid \eta_{\ell}\; e \mid \langle e, e \rangle^{\tau} \\
&\quad \mid \mathtt{proj}_i\; e \mid \mathtt{inj}_i^{(\tau+\tau)}\; e \mid \mathtt{bind}\; x = e \;\mathtt{in}\; e \\
&\quad \mid \mathtt{case}^{\tau}\; e \;\mathtt{of}\; \mathtt{inj}_1^{\tau}(x).e \mid \mathtt{inj}_2^{\tau}(x).e \\
&\quad \mid \mathtt{run}^{\tau}\; e@p \mid \underline{\mathtt{ret}\; e@p} \mid \underline{\mathtt{expect}^{\tau}} \\
&\quad \mid \mathtt{select}^{\tau}\; e \;\mathtt{or}\; e \mid \mathtt{compare}^{\tau}\; e \;\mathtt{and}\; e
\end{aligned}
$$

Fig. 6. FLAQR syntax. Shaded terms are new to FLAQR. Underlined terms are used during evaluation and not available at the source level.

$$p \sqcup q \triangleq (p^{\mathtt{c}} \wedge q^{\mathtt{c}}) \wedge (p^{\mathtt{i}} \vee q^{\mathtt{i}}) \wedge (p^{\mathtt{a}} \vee q^{\mathtt{a}})$$

$$p \sqcap q \triangleq (p^{\mathtt{c}} \vee q^{\mathtt{c}}) \wedge (p^{\mathtt{i}} \wedge q^{\mathtt{i}}) \wedge (p^{\mathtt{a}} \wedge q^{\mathtt{a}})$$

Based on this, the equivalence classes of $\succcurlyeq$ and $\sqsubseteq$ are identical, meaning that the lattice formed by $\sqsubseteq$ with joins $\sqcup$ and meets $\sqcap$ has the same elements as the acts-for lattice. A flow from p to q is secure only when $q^{\mathtt{c}}$ is at least as confidential as $p^{\mathtt{c}}$, $q^{\mathtt{i}}$ trusts information influenced by $p^{\mathtt{i}}$, and $q^{\mathtt{a}}$ cannot cause failures that $p^{\mathtt{a}}$ cannot.

4. FLAQR syntax and semantics

Figures 6 and 7 present the FLAQR syntax and selected evaluation rules. Figure 8 presents the evaluation context. For space and exposition purposes, we omit some term annotations and standard lambda calculus rules in order to focus on FLAQR's contributions, but the complete, annotated FLAQR syntax and semantics can be found in Figs 29 and 31 in the Appendix.

FLAQR is based on FLAC [2,5], a monadic calculus in the style of Abadi's Polymorphic DCC [1]. In addition to standard extensions to System F [12,13,23] such as pairs and tagged unions, an Abadi-style calculus supports monadic operations on values in a monad indexed by a lattice of security labels. Such a value has a type of the form $\ell \; \mathtt{says} \; \tau$, meaning that it is a value of type τ, *protected* at level ℓ, where ℓ is an element of the security lattice. Here we focus on FLAQR's additions to FLAC and DCC, and refer the readers to Figs 33 and 34 in the Appendix for our complete formalization.

FLAQR builds on FLAC's expressive principal algebra and type system to model distributed security policies for applications that use replication and consensus. FLAC supports arbitrary policy downgrades through dynamic delegations of authority, but for simplicity we omit these features in FLAQR.

AUTHOR COPY

[E-SEALED] $\eta_\ell\ v \longrightarrow (\overline{\eta}_\ell\ v)$ [E-BINDM] $\mathtt{bind}\ x = (\overline{\eta}_\ell\ v)\ \mathtt{in}\ e \longrightarrow e[x \mapsto v]$

[E-COMPARE] $\mathtt{compare}\ (\overline{\eta}_{\ell_1}\ v)\ \mathtt{and}\ (\overline{\eta}_{\ell_2}\ v) \longrightarrow (\overline{\eta}_{\ell_1 \oplus \ell_2}\ v)$

[E-COMPAREFAIL] $$\frac{v_1 \neq v_2 \qquad \tau = (\ell_1 \oplus \ell_2)\ \mathtt{says}\ \tau'}{\mathtt{compare}\ (\overline{\eta}_{\ell_1}\ v_1)\ \mathtt{and}\ (\overline{\eta}_{\ell_2}\ v_2) \longrightarrow \mathtt{fail}^{\tau}}$$

[E-COMPAREFAILL] $$\frac{\tau_1 = \ell_1\ \mathtt{says}\ \tau \qquad \tau' = (\ell_1 \oplus \ell_2)\ \mathtt{says}\ \tau \qquad f_2 = \begin{cases} (\overline{\eta}_{\ell_2}\ v) \\ \mathtt{fail}^{\ell_2\ \mathtt{says}\ \tau} \end{cases}}{\mathtt{compare}\ (\mathtt{fail}^{\tau_1})\ \mathtt{and}\ f_2 \longrightarrow \mathtt{fail}^{\tau'}}$$

[E-COMPAREFAILR] $$\frac{\tau_2 = \ell_2\ \mathtt{says}\ \tau \qquad \tau' = (\ell_1 \oplus \ell_2)\ \mathtt{says}\ \tau}{\mathtt{compare}\ (\overline{\eta}_{\ell_1}\ v)\ \mathtt{and}\ \mathtt{fail}^{\tau_2} \longrightarrow \mathtt{fail}^{\tau'}}$$

[E-SELECT] $\mathtt{select}\ (\overline{\eta}_{\ell_1}\ v_1)\ \mathtt{or}\ (\overline{\eta}_{\ell_2}\ v_2) \longrightarrow (\overline{\eta}_{\ell_1 \ominus \ell_2}\ v_1)$

[E-SELECTL] $\mathtt{select}\ (\overline{\eta}_{\ell_1}\ v)\ \mathtt{or}\ (\mathtt{fail}^{\ell_1\ \mathtt{says}\ \tau}) \longrightarrow (\overline{\eta}_{\ell_1 \ominus \ell_2}\ v)$

[E-SELECTFAIL] $$\frac{\forall i \in \{1,2\} \qquad \tau_i = \ell_i\ \mathtt{says}\ \tau \qquad \tau' = (\ell_1 \ominus \ell_2)\ \mathtt{says}\ \tau}{\mathtt{select}\ (\mathtt{fail}^{\tau_1})\ \mathtt{or}\ (\mathtt{fail}^{\tau_2}) \longrightarrow \mathtt{fail}^{\tau'}}$$

[E-RETSTEP] $$\frac{e \longrightarrow e'}{\mathtt{ret}\ e@c \longrightarrow \mathtt{ret}\ e'@c}$$ [E-STEP] $$\frac{e \longrightarrow e'}{E[e] \longrightarrow E[e']}$$

Fig. 7. FLAQR local semantics.

$$\begin{aligned} E ::= {} & [\cdot] \mid E\ e \mid v\ E \mid \eta_\ell\ E \mid \mathtt{bind}\ x = E\ \mathtt{in}\ e \\ & \mid \mathtt{ret}\ E@p \mid \mathtt{select}\ E\ \mathtt{or}\ e \mid \mathtt{select}\ f\ \mathtt{or}\ E \\ & \mid \mathtt{compare}\ E\ \mathtt{and}\ e \mid \mathtt{compare}\ f\ \mathtt{and}\ E \end{aligned}$$

Fig. 8. FLAQR evaluation context.

AUTHOR COPY

The monadic unit or return term $\eta_\ell\ e$ protects the value that e evaluates to at level ℓ (E-SEALED).[4] Protected values, $(\overline{\eta}_\ell\ v)$ cannot be operated on directly. Instead, a bind expression must be used to bind the protected value to a variable whose scope is limited to the body of the bind term (E-BINDM). The body performs the desired computation and "returns" the result to the monad, ensuring the result is protected. These rules (E-SEALED and E-BINDM) FLAQR inherits from FLAC. The remaining rules of Fig. 7 are specific to FLAQR.

The primary novelty in the FLAQR calculus is the introduction of `compare` and `select` terms for expressing consensus and replication operations. We represent the consensus problem as a comparison of two values with the same underlying type but distinct outer security labels. In other words, we want to check the equality of values produced by two different principals. If the values match, we can treat them as having the (partially) combined integrity of the principals. If not, then the principals failed to reach consensus.

Rule E-COMPARE defines the former case: two syntactically equal values protected at different labels evaluate to a value that combines labels using the *compare action* on labels $\oplus$. Intuitively, $\ell_1 \oplus \ell_2$

[4]Polymorphic DCC does not define a term similar to $(\overline{\eta}_\ell\ v)$ and thus does not have an rule equivalent to E-SEALED. This approach enables us to distinguish where a value may be created (e.g., on a host authorized to read and create values protected at ℓ) and use more permissive rules to control where a sealed value may flow.

$$[\text{E-AppFail}] \quad \lambda(x{:}\tau)[pc].\, e\ \mathtt{fail}^{\tau} \longrightarrow e[x \mapsto \mathtt{fail}^{\tau}] \qquad [\text{E-SealedFail}] \quad \eta_\ell\ \mathtt{fail}^{\tau} \longrightarrow \mathtt{fail}^{\ell\ \mathtt{says}\ \tau}$$

$$[\text{E-InjFail}] \quad \mathtt{inj}_i^{(\tau_1+\tau_2)}\ \mathtt{fail}^{\tau_i} \longrightarrow \mathtt{fail}^{(\tau_1+\tau_2)} \qquad [\text{E-ProjFail}] \quad \mathtt{proj}_i\ \mathtt{fail}^{(\tau_1\times\tau_2)} \longrightarrow \mathtt{fail}^{\tau_i}$$

Fig. 9. `fail` propagation rules.

determines the increase in integrity and the corresponding decrease in availability inherent in requiring a consensus. We define $\oplus$ formally in Definition 1.

Definition 1 (Compare action on principals)**.**

$$\ell_1 \oplus \ell_2 \triangleq (\ell_1^{\mathsf{c}} \wedge \ell_2^{\mathsf{c}}) \wedge (\ell_1^{\mathsf{i}} \boxplus \ell_2^{\mathsf{i}}) \wedge (\ell_1^{\mathsf{a}} \vee \ell_2^{\mathsf{a}})$$

We also lift this notation to `says` types by defining

$$\ell_1\ \mathtt{says}\ \tau \oplus \ell_2\ \mathtt{says}\ \tau \triangleq (\ell_1 \oplus \ell_2)\ \mathtt{says}\ \tau$$

As discussed in Section 3, the integrity authority of `compare` is not as trusted as the conjunction of ℓ_1 and ℓ_2's integrity. Instead, we represent the limited "increase" in integrity authority[5] using a partial conjunction in Definition 1. In contrast, the decrease in availability is represented by a (full) $\ell_1^{\mathsf{a}} \vee \ell_2^{\mathsf{a}}$ since either ℓ_1 or ℓ_2 could unilaterally cause the compare expression to fail.

The decreased availability resulting from applying `compare` is more apparent in rules E-CompareFail, E-CompareFailL and E-CompareFailR. In E-CompareFail, two unequal values are compared, resulting in a failure. Failure is represented syntactically using a $\mathtt{fail}^{\tau}$ term. We use a type annotation τ on many terms in our formal definitions so that our semantics is well defined with respect to failure terms, but we omit most of these annotations in Fig. 7. These annotations are only necessary for our formalization and would be unnecessary in a FLAQR implementation.

A `compare` term may also result in failure if either subexpression fails. Rule E-CompareFailL and E-CompareFailR, defines how failure of an input propagates to the output. In fact, most FLAQR terms result in failure when a subexpression fails. Figure 9 presents selected failure propagation rules (complete failure propagation rules are presented in Fig. 32). Note that `fail` terms are treated similarly to values, but are distinct from them. For example, in E-AppFail, applying a lambda term to a fail term substitutes the failure as it would a value, but in E-SealedFail the failure is propagated beyond the monadic unit term. This latter behavior captures the idea that failures cannot be hidden or isolated in the same way as secrets or untrusted data.

Failures are tolerated using replication. A `select` term will evaluate to a value as long as at least one of its subexpressions does not fail. For example, rule E-SelectL returns its left subexpression when the right subexpression fails. In contrast to `compare`, applying `select` increases availability since either subexpression can be used, but reduces integrity since influencing only one of the subexpressions is potentially sufficient to influence the result of evaluating `select`. The effect of a `select` statement on the labels of its sub-expressions is captured with the *select action* $\ominus$.

Definition 2 (Select action on principals)**.**

$$\ell_1 \ominus \ell_2 \triangleq (\ell_1^{\mathsf{c}} \wedge \ell_2^{\mathsf{c}}) \wedge (\ell_1^{\mathsf{i}} \boxminus \ell_2^{\mathsf{i}}) \wedge (\ell_1^{\mathsf{a}} \wedge \ell_2^{\mathsf{a}})$$

[5]Strictly speaking, $x \boxplus y$ is not an increase in integrity over x (or y); $x \boxplus y$ and x are incomparable.

AUTHOR COPY

$$[\text{E-DSTEP}]\ \frac{e \longrightarrow e'}{\langle E[e], c\rangle \,\&\, t \Longrightarrow \langle E[e'], c\rangle \,\&\, t}$$

$$[\text{E-RUN}]\ \langle E[\mathtt{run}^{\tau}\ e@c'], c\rangle \,\&\, t \Longrightarrow \langle \mathtt{ret}\ e@c, c'\rangle \,\&\, \langle E[\mathtt{expect}^{\tau}], c\rangle :: t$$

$$[\text{E-RETV}]\ \langle \mathtt{ret}\ v@c, c'\rangle \,\&\, \langle E[\mathtt{expect}^{pc^{ia}}\ \mathtt{says}\ \tau'], c\rangle :: t \Longrightarrow \langle E[(\overline{\eta}_{pc^{ia}}\ v)], c\rangle \,\&\, t$$

$$[\text{E-RETFAIL}]\ \langle \mathtt{ret}\ (\mathtt{fail}^{\tau'})@c, c'\rangle \,\&\, \langle E[\mathtt{expect}^{pc^{ia}}\ \mathtt{says}\ \tau'], c\rangle :: t \Longrightarrow \langle E[\mathtt{fail}^{pc^{ia}}\ \mathtt{says}\ \tau'], c\rangle \,\&\, t$$

Fig. 10. Global semantics.

$$s ::= \langle e, c\rangle \,\&\, t$$
$$t ::= \mathtt{empty} \mid \langle E[\mathtt{expect}^{\tau}], c\rangle :: t$$

Fig. 11. Global configuration stack.

We define the select action on types similarly to compare:

$$\ell_1\ \mathtt{says}\ \tau \ominus \ell_2\ \mathtt{says}\ \tau = (\ell_1 \ominus \ell_2)\ \mathtt{says}\ \tau$$

The end result of a *select* statement, $\mathtt{select}\ (\overline{\eta}_{\ell_1}\ v)\ \mathtt{or}\ (\overline{\eta}_{\ell_2}\ v)$, will have integrity of either $\ell_1^{\mathtt{i}}$ or $\ell_2^{\mathtt{i}}$ since only one of the two possible values will be used. We use a partial disjunction to represent this integrity since the result does not have the same integrity as ℓ_1 or ℓ_2, but does have more integrity than $\ell_1 \vee \ell_2$ since it is never the case that *both* principals influence the output.

4.1. Global semantics

We capture the distributed nature of quorum replication by embedding the local semantic rules within a global distributed semantics defined in Fig. 10. This semantics uses a *configuration stack* $s = \langle e, c\rangle \,\&\, t$ (Fig. 11) to keep track of the currently executing expression e, the host on which it is executing c, and the remainder of the stack t. We also make explicit use of the evaluation contexts from Fig. 7 to identify the reducible subterms across stack elements.

The core operation for distributed computation is $\mathtt{run}^{\tau}\ e@p$ which runs the computation e of type τ on node p. Local evaluation steps are captured in the global semantics via rule E-DSTEP. This rule says that if e steps to e' locally, then $E[e]$ steps to $E[e']$ globally.

Rule E-RUN takes an expression e at host c, pushes a new configuration on the stack containing e at host c' and places an `expect` term at c as a place holder for the return value.

Once the remote expression is fully evaluated, rule E-RETV pops the top configuration off the stack and replaces the `expect` term at c with the protected value $(\overline{\eta}_{pc^{\mathtt{ia}}}\ v)$. Rule E-RETFAIL serves the same purpose for `fail` terms, but is necessary since `fail` terms are not considered values (see Fig. 6). The label $pc^{\mathtt{ia}}$ reflects both the integrity and availability context of the caller (c) as well as the integrity and availability of the remote host (c'). We discuss this aspect of remote execution in more detail in Section 5.

5. FLAQR typing rules

As we have a local and global semantics, we have two corresponding forms of typing judgements: local typing judgments for expressions and global typing judgments for the stack. Local typing judgments

AUTHOR COPY

$$\boxed{\Pi; \Gamma; pc; c \vdash e : \tau}$$

$$[\textsc{Unit}]\ \frac{\Pi \Vdash c \succcurlyeq pc}{\Pi; \Gamma; pc; c \vdash () : \mathtt{unit}} \qquad [\textsc{Fail}]\ \frac{\Pi \Vdash c \succcurlyeq pc}{\Pi; \Gamma; pc; c \vdash \mathtt{fail}^{\tau} : \tau} \qquad [\textsc{Expect}]\ \frac{\Pi \Vdash c \succcurlyeq pc}{\Pi; \Gamma; pc; c \vdash \mathtt{expect}^{\tau} : \tau}$$

$$[\textsc{Lam}]\ \frac{\begin{array}{c}\Pi; \Gamma, x{:}\tau_1; pc'; u \vdash e : \tau_2 \qquad \Pi \Vdash c \succcurlyeq pc \\ u = \mathtt{C}(\tau_1 \xrightarrow{pc'} \tau_2) \qquad \Pi \Vdash c \succcurlyeq u\end{array}}{\Pi; \Gamma; pc; c \vdash \lambda(x{:}\tau_1)[pc'].\, e : \tau_1 \xrightarrow{pc'} \tau_2} \qquad [\textsc{App}]\ \frac{\begin{array}{c}\Pi; \Gamma; pc; c \vdash e_1 : \tau' \xrightarrow{pc'} \tau \\ \Pi; \Gamma; pc; c \vdash e_2 : \tau' \qquad \Pi \Vdash pc \sqsubseteq pc' \\ \Pi \Vdash c \succcurlyeq pc\end{array}}{\Pi; \Gamma; pc; c \vdash e_1\ e_2 : \tau}$$

$$[\textsc{UnitM}]\ \frac{\begin{array}{c}\Pi; \Gamma; pc; c \vdash e : \tau \qquad \Pi \Vdash pc \sqsubseteq \ell \\ \Pi \Vdash c \succcurlyeq pc\end{array}}{\Pi; \Gamma; pc; c \vdash \eta_\ell\ e : \ell\ \mathtt{says}\ \tau} \qquad [\textsc{Sealed}]\ \frac{\Pi; \Gamma; pc; c \vdash v : \tau \qquad \Pi \Vdash c \succcurlyeq pc}{\Pi; \Gamma; pc; c \vdash (\overline{\eta}_\ell\ v) : \ell\ \mathtt{says}\ \tau}$$

$$[\textsc{BindM}]\ \frac{\begin{array}{c}\Pi; \Gamma; pc; c \vdash e' : \ell\ \mathtt{says}\ \tau' \qquad \Pi \Vdash \ell \sqcup pc \sqsubseteq \tau \\ \Pi; \Gamma, x{:}\tau'; \ell \sqcup pc; c \vdash e : \tau \qquad \Pi \Vdash c \succcurlyeq pc\end{array}}{\Pi; \Gamma; pc; c \vdash \mathtt{bind}\ x = e'\ \mathtt{in}\ e : \tau} \qquad [\textsc{Run}]\ \frac{\begin{array}{c}\Pi; \Gamma; pc'; c' \vdash e : \tau' \qquad \Pi \Vdash pc \sqsubseteq pc' \\ \Pi \Vdash c \succcurlyeq pc \qquad \Pi \Vdash c \succcurlyeq \mathtt{C}(\tau') \\ \tau = pc'^{\mathrm{ia}}\ \mathtt{says}\ \tau'\end{array}}{\Pi; \Gamma; pc; c \vdash \mathtt{run}^{\tau}\ e@c' : \tau}$$

$$[\textsc{Ret}]\ \frac{\begin{array}{c}\Pi; \Gamma; pc; c \vdash e : \tau \qquad \Pi \Vdash c' \succcurlyeq \mathtt{C}(\tau) \\ \Pi \Vdash c \succcurlyeq pc\end{array}}{\Pi; \Gamma; pc; c \vdash \mathtt{ret}\ e@c' : pc^{\mathrm{ia}}\ \mathtt{says}\ \tau}$$

$$[\textsc{Compare}]\ \frac{\begin{array}{c}\forall i \in \{1,2\}. \Pi; \Gamma; pc; c \vdash e_i : \ell_i\ \mathtt{says}\ \tau \\ \Pi \Vdash c \rhd \ell_i\ \mathtt{says}\ \tau \qquad \Pi \Vdash c \succcurlyeq pc\end{array}}{\Pi; \Gamma; pc; c \vdash \mathtt{compare}\ e_1\ \mathtt{and}\ e_2 : (\ell_1 \oplus \ell_2)\ \mathtt{says}\ \tau}$$

$$[\textsc{Select}]\ \frac{\begin{array}{c}\forall i \in \{1,2\}. \Pi; \Gamma; pc; c \vdash e_i : \ell_i\ \mathtt{says}\ \tau \\ \Pi \Vdash c \succcurlyeq pc\end{array}}{\Pi; \Gamma; pc; c \vdash \mathtt{select}\ e_1\ \mathtt{or}\ e_2 : (\ell_1 \ominus \ell_2)\ \mathtt{says}\ \tau}$$

Fig. 12. Typing rules for expressions.

AUTHOR COPY

have the form $\Pi; \Gamma; pc; c \vdash e : \tau$. Π is the program's delegation context and is used to derive acts-for relationships with the rules in Figs 4 and 40. Γ is the typing context containing in-scope variable names and their types. The *pc* label tracks the information flow policy on the program counter (due to control flow) and on unsealed protected values such as in the body of a `bind`.

Figure 12 presents a selection of local typing rules. Each typing rule includes an acts-for premise of the form $\Pi \Vdash c \succcurlyeq pc$. This enforces the invariant that each host principal c has control of the program it executes locally. Thus for any judgment $\Pi; \Gamma; pc; c \vdash e : \tau$ *pc* should never exceed the authority of c, the principal executing the expression. Rules FAIL and EXPECT type `fail` and `expect` terms according to their type annotation τ. Rule LAM types lambda abstractions. Since functions are first-class values, we have to ensure that the *pc* annotation on the lambda term preserves the invariant $\Pi \Vdash c \succcurlyeq pc$. The *clearance* of a type τ, written $\mathtt{C}(\tau)$, is an upper bound on the *pc* annotations of the function types in τ. By checking that $\Pi \Vdash c \succcurlyeq \mathtt{C}(\tau_1 \xrightarrow{pc'} \tau_2)$ holds (along with similar checks in RUN and RET), we ensure the contents of the lambda term is protected when sending or receiving lambda expressions, and that hosts never receive a function they cannot securely execute. Due to space constraints, the definition of $\mathtt{C}(\cdot)$ is presented in Appendix A, in Fig. 35. The APP rule requires the *pc* label at any function application to flow to the function's *pc* label annotation. Hence the premise $\Pi \Vdash pc \sqsubseteq pc'$.

Protected terms $\eta_\ell\ e$ are typed by rule UNITM as ℓ `says` τ where τ is the type of e. Additionally, it requires that $\Pi \Vdash pc \sqsubseteq \ell$. This ensures that any unsealed values in the context are adequately protected

$\boxed{\Pi; \Gamma; pc \vdash \langle e, c\rangle \,\&\, t : \tau}$

$$[\textsc{Head}]\ \frac{\Pi; \Gamma; pc'; c \vdash e : \tau' \qquad \Pi; \Gamma; pc \vdash t : [\tau']\tau \qquad \Pi \Vdash pc \sqsubseteq pc' \qquad \Pi \Vdash c \succcurlyeq pc}{\Pi; \Gamma; pc \vdash \langle e, c\rangle \,\&\, s : \tau}$$

$\boxed{\Pi; \Gamma; pc \vdash \langle e, c\rangle :: t : [\tau']\tau}$

$$[\textsc{Tail}]\ \frac{\Pi; \Gamma; pc'; c \vdash E[\texttt{expect}^{\tau'}] : \hat{\tau} \qquad \Pi; \Gamma; pc \vdash t : [\hat{\tau}]\tau \qquad \Pi \Vdash pc \sqsubseteq pc' \qquad \Pi \Vdash c \succcurlyeq pc}{\Pi; \Gamma; pc \vdash \langle E[\texttt{expect}^{\tau'}], c\rangle :: t : [\tau']\tau}$$

$$[\textsc{Emp}]\ \Pi; \Gamma; pc \vdash \texttt{empty} : [\tau]\tau$$

Fig. 13. Typing rules for configuration stack.

by policy ℓ if they are used by e. The SEALED rule types protected values ($\overline{\eta}_\ell\ v$). These values are well-typed at any host, and does not require $\Pi \Vdash pc \sqsubseteq \ell$ since no unsealed values in the context could be captured by the (closed) value v.

Computation on protected values occurs in bind terms $\texttt{bind}\ x = e'\ \texttt{in}\ e$. The policy protecting e must be at least as restrictive as the policy on e' so that the occurrences of x in e are adequately protected. Thus, rule BINDM requires $\Pi \Vdash \ell \sqcup pc \sqsubseteq \tau$, and furthermore e is typed at a more restrictive program counter label $\ell \sqcup pc$ to reflect the dependency of e on the value bound to x.

Rule RUN requires that the pc at the local host flow to the pc' of the remote host, and that e be well-typed at c', which implies that c' acts for pc'. Additionally, c must act for the clearance of the remote return type τ' to ensure c is authorized to receive the return value. The type of the run expression is $pc'^{\texttt{ia}}\ \texttt{says}\ \tau'$, which reflects the fact that c' controls the availability of the return value and also has some influence on which value of type τ' is returned. Although c' may not be able to *create* a value of type τ' unless $pc'^{\texttt{ia}}$ flows to τ', if c' has *access* to more than one value of type τ', it could choose which one to return. Rule RET requires that expression e is welltyped at c and that c' is authorized to receive the return value based on the clearance of τ.

The COMPARE rule gives type $(\ell_1 \oplus \ell_2)\ \texttt{says}\ \tau$ to the expression $\texttt{compare}\ e_1\ \texttt{and}\ e_2$ where e_1 and e_2 have types $\ell_1\ \texttt{says}\ \tau$ and $\ell_2\ \texttt{says}\ \tau$ respectively. Additionally, it requires that c, the host executing the $\texttt{compare}$, is authorized to fully examine the results of evaluating e_1 and e_2 so that they may be checked for equality.[6] This requirement is captured by the premise $\Pi \Vdash c \rhd \ell_i\ \texttt{says}\ \tau$, pronounced "$c$ *reads* $\ell_i\ \texttt{says}\ \tau$". The inference rules for the reads judgment are found in Fig. 37 in Appendix A.

Finally, the SELECT rule gives type $(\ell_1 \ominus \ell_2)\ \texttt{says}\ \tau$ to the expression $\texttt{select}\ e_1\ \texttt{or}\ e_2$ where e_1 and e_2 have types $\ell_1\ \texttt{says}\ \tau$ and $\ell_2\ \texttt{says}\ \tau$ respectively.

The typing judgment for the global configuration is presented in Fig. 13 and consists of three rules. Rule HEAD shows that the global configuration $\langle e, c\rangle \,\&\, t$, is well-typed if the expression e is well-typed at host c with program counter pc' where $\Pi \Vdash pc \sqsubseteq pc'$ and the tail t is well-typed. $[\tau']\tau$ means that the tail of the stack is of type τ while the expression in the head of the configuration is of type τ'. We introduced rules TAIL(when $t \neq \texttt{empty}$) and EMP(when $t = \texttt{empty}$) to typecheck the tail t.

$\langle E[\texttt{expect}^{\tau'}], c\rangle :: t$ is well-typed with type $[\tau']\tau$, if expression $E[\texttt{expect}^{\tau'}]$ is well-typed with type $\hat{\tau}$ at host c. And, the rest of the stack t needs to be well-typed with type $[\hat{\tau}]\tau$. Rule EMP says the tail is empty and the type of the expression in the head of the configuration is τ, in which case the type of the whole stack is $[\tau]\tau$.

[6]Assuming a more sophisticated mechanism for checking equality that reveals less information to the host such as zero-knowledge proofs or a trusted execution environment could justify relaxing this constraint.

AUTHOR COPY

6. Availability attackers

Availability attackers are different from traditional integrity and confidentiality attackers. While an integrity attacker's goal is to manipulate data and a confidentiality attacker's goal is to learn secrets, an availability attacker's goal is to cause failures. In our model, an availability attacker can substitute a value only with a `fail` term. Integrity attackers may also cause failures in consensus based protocols when consensus is not reached because of data manipulation. In FLAQR this scenario is relevant during executing a `compare` statement: if one of the values in the `compare` statement is substituted with a wrong (mismatching) value then a `fail` term is returned. Thus we need to consider an availability attacker's integrity authority when reasoning about its power to `fail` a program. Specifically, the authority of principal ℓ as an availability attacker is $\ell^{\mathtt{ia}}$.

We consider a static but active attacker model similar to those used in Byzantine consensus protocols. By static we mean which principal or collection of principals can act maliciously is fixed prior to program execution. By active we mean that the attackers may manipulate inputs (including higher-order functions) during run time. We formally define the power of an availability attacker with respect to quorum systems.

Availability attackers in FLAQR are somewhat different than integrity and confidentiality attackers because we want to represent multiple possible attackers but limit which attackers are active for a particular execution. This goal supports the bounded fault assumptions found in consensus protocols where system configurations assume an upper bound on the number of faults possible.

A quorum system $\mathcal{Q}$ is represented as set of sets of hosts (or principals) e.g. $\mathcal{Q} = \{q_1, q_2, \dots, q_n\}$. Here each q_i represents a set of principals whose consensus is adequate for the system to make progress. We define availability attackers in terms of the *toleration set* $[\![\mathcal{Q}]\!]$ of a quorum system $\mathcal{Q}$. The toleration set is a set of principals where each principal represents an upper bound on the authority of an attacker the quorum can tolerate without failing.

Example.

(1) The toleration set for quorum $\mathcal{Q}_1 = \{q_1 := \{a, b\}; q_2 := \{b, c\}; q_3 := \{a, c\}\}$ is $[\![\mathcal{Q}_1]\!] = \{a^{\mathtt{ia}}, b^{\mathtt{ia}}, c^{\mathtt{ia}}\}$,
(2) For heterogeneous quorum system $\mathcal{Q}_2 = \{q_1 := \{p, q\}; q_2 := \{r\}\}$ the toleration set is $[\![\mathcal{Q}_2]\!] = \{p^{\mathtt{ia}} \wedge q^{\mathtt{ia}}, r^{\mathtt{ia}}\}$
(3) For $\mathcal{Q}_3 = \{q := \{alice\}\}$ the toleration set is $[\![\mathcal{Q}_3]\!] = \{\}$, i.e. $\mathcal{Q}_3$ can not tolerate any fault.

An availability attacker's authority is at most equivalent to a (single) principal's authority in the toleration set. We define the set of all such attackers for a quorum $\mathcal{Q}$ as

$$\mathcal{A}_{[\![\mathcal{Q}]\!]} = \{\ell \mid \exists \ell' \in [\![\mathcal{Q}]\!].\Pi \Vdash \ell' \succcurlyeq \ell\}.$$

which includes weaker attackers who a principal in the toleration set may act on behalf of.

The fails relation ($\gtrdot$) determines whether a principal can cause a program of a particular type to evaluate to `fail`. Similar to the reads judgment, the fails judgment not only considers the outermost `says` principal, but also any nested `says` principals whose propagated failures could cause the whole term to fail. Figure 14 defines the fails judgment, written $\Pi \Vdash l \gtrdot \tau$, which describes when a principal l can fail an expression of type τ in delegation context Π.

AUTHOR COPY

$$[\text{A-PAIR}]\ \frac{\Pi \Vdash \ell \gtrdot \tau_i \quad i \in \{1,2\}}{\Pi \Vdash \ell \gtrdot (\tau_1 \times \tau_2)} \qquad [\text{A-SUM}]\ \frac{\Pi \Vdash \ell \gtrdot \tau_i \quad i \in \{1,2\}}{\Pi \Vdash \ell \gtrdot (\tau_1 + \tau_2)} \qquad [\text{A-FUN}]\ \frac{\Pi \Vdash \ell \gtrdot \tau_2}{\Pi \Vdash \ell \gtrdot \tau_1 \xrightarrow{pc'} \tau_2}$$

$$[\text{A-TYPE}]\ \frac{\Pi \Vdash \ell \gtrdot \tau}{\Pi \Vdash \ell \gtrdot \ell' \ \mathsf{says}\ \tau} \qquad [\text{A-AVAIL}]\ \frac{\Pi \Vdash \ell^{\mathsf{a}} \succcurlyeq \ell'^{\mathsf{a}}}{\Pi \Vdash \ell \gtrdot \ell' \ \mathsf{says}\ \tau} \qquad [\text{A-INTEGCOM}]\ \frac{\Pi \Vdash \ell^{\mathsf{i}} \succcurlyeq \ell_j{}^{\mathsf{i}}, j \in \{1,2\}}{\Pi \Vdash \ell \gtrdot (\ell_1 \oplus \ell_2)\ \mathsf{says}\ \tau}$$

Fig. 14. Fails judgments.

Consider an expression $\eta_\ell\ (\eta_{\ell'}\ e)$ and an attacker principal l_a. If $\Pi \Vdash l_a^{\mathsf{c}} \succcurlyeq \ell'^{\mathsf{c}}$, and $\Pi \nVdash l_a^{\mathsf{c}} \succcurlyeq \ell^{\mathsf{c}}$, then the attacker learns nothing by evaluating $\eta_\ell\ (\eta_{\ell'}\ e)$. Similarly, if $\Pi \Vdash l_a^{\mathsf{i}} \succcurlyeq \ell'^{\mathsf{i}}$ and $\Pi \nVdash l_a^{\mathsf{i}} \succcurlyeq \ell^{\mathsf{i}}$, then the attacker cannot influence the value $\eta_\ell\ (\eta_{\ell'}\ e)$.

In contrast, if $\Pi \Vdash l_a{}^{\mathsf{a}} \succcurlyeq \ell'^{\mathsf{a}}$, and $\Pi \nVdash l_a^{\mathsf{a}} \succcurlyeq \ell^{\mathsf{a}}$, an availability attacker may cause $\eta_{\ell'}\ e$ to evaluate to $\mathtt{fail}^{\ell'\ \mathsf{says}\ \tau}$, which steps to $\mathtt{fail}^{\ell\ \mathsf{says}\ (\ell'\ \mathsf{says}\ \tau)}$ by E-SEALEDFAIL. The fails relation reflects this possibility. Using A-TYPE and A-AVAIL (or A-INTEGCOM if ℓ' was of form $(\ell_1 \oplus \ell_2)$) we get $\Pi \Vdash l_a \gtrdot \ell\ \mathsf{says}\ (\ell'\ \mathsf{says}\ \tau)$.

We use the fails relation and the attacker set to define which availability policies a particular quorum system is capable of enforcing. We say $\mathcal{Q}$ *guards* τ if the following rule applies:

$$[\text{Q-GUARD}]\ \frac{\forall \ell \in \mathcal{A}_{[\![\mathcal{Q}]\!]}. \Pi \nVdash \ell \gtrdot \tau}{\Pi \Vdash \mathcal{Q}\ \mathit{guards}\ \tau}$$

Definition 3 (Valid quorum type)**.** A type τ is a *valid quorum type* with respect to quorum system $\mathcal{Q}$ and delegation set Π if the condition $\Pi \Vdash \mathcal{Q}\ \mathit{guards}\ \tau$ is satisfied.

Example. If $\mathcal{Q} = \{q_1 := \{a, b\}; q_2 := \{b, c\}; q_3 := \{a, c\}\}$ and $\ell_{\mathcal{Q}} = (a \oplus b) \ominus (b \oplus c) \ominus (a \oplus c)$ then $\ell_{\mathcal{Q}}\ \mathsf{says}\ (a\ \mathsf{says}\ \tau)$ is not a valid quorum type because $\Pi \nVdash \mathcal{Q}\ \mathit{guards}\ (\ell_{\mathcal{Q}}\ \mathsf{says}\ (a\ \mathsf{says}\ \tau))$ as $\Pi \Vdash a^{\mathsf{ia}} \gtrdot \ell_{\mathcal{Q}}\ \mathsf{says}\ (a\ \mathsf{says}\ \tau)$ and $a^{\mathsf{ia}} \in \mathcal{A}_{[\![\mathcal{Q}]\!]}$. But it is a valid quorum type for heterogeneous quorum system $\mathcal{Q}' = \{q_1 := \{a, b\}; q_2 := \{a, c\}\}$ as $a^{\mathsf{ia}} \notin \mathcal{A}_{[\![\mathcal{Q}']\!]}$.

7. Security properties

To evaluate the formal properties of FLAQR, we prove that FLAQR preserves noninterference for confidentiality, integrity, and availability (Section 7.2). These theorems state that attackers cannot learn secret inputs, influence trusted outputs, or control the failure behavior of well-typed FLAQR programs. In addition, we also prove additional theorems that formalize the soundness of our type system with respect to a program's failure behavior.

7.1. Soundness of failure

FLAQR's semantics uses the `compare` and `select` security abstractions and the failure propagation rules to model failure and failure-tolerance in distributed programs, and FLAQR's type system lets us reason statically about this failure behavior. To verify that such reasoning is *sound*, we prove two related theorems regarding the type of a program and the causes of potential failures.

In pursuit of this goal, this section introduces our *blame semantics* which reasons about failure-causing (faulty) principals during program execution. The goal is to record the set of principals which may cause run-time failures as a constraint on the set of faulty hosts $\mathcal{F}$. Figure 15 presents the syntax of *blame constraints*, which are boolean formulas representing a lower bound on the contents of $\mathcal{F}$. Atomic

AUTHOR COPY

$$\mathcal{C} ::= \mathcal{F} = \varnothing \mid \mathcal{B}$$
$$\mathcal{B} ::= \ell \in \mathcal{F} \mid \mathcal{B}_1 \ \texttt{OR} \ \mathcal{B}_2 \mid \mathcal{B}_1 \ \texttt{AND} \ \mathcal{B}_2$$

Fig. 15. Blame constraint syntax.

$$[\text{C-Conj}]\ \frac{\mathcal{C} \models \ell_1 \in \mathcal{F} \quad \mathcal{C} \models \ell_2 \in \mathcal{F}}{\mathcal{C} \models \ell_1 \wedge \ell_2 \in \mathcal{F}} \qquad [\text{C-Disj}]\ \frac{\exists i \in \{1,2\}.\, \mathcal{C} \models \ell_i \in \mathcal{F}}{\mathcal{C} \models \ell_1 \vee \ell_2 \in \mathcal{F}} \qquad [\text{C-ParAnd}]\ \frac{\exists i \in \{1,2\}.\, \mathcal{C} \models \ell_i \in \mathcal{F}}{\mathcal{C} \models \ell_1 \boxplus \ell_2 \in \mathcal{F}}$$

$$[\text{C-ParOr}]\ \frac{\mathcal{C} \models \ell_1 \in \mathcal{F} \quad \mathcal{C} \models \ell_2 \in \mathcal{F}}{\mathcal{C} \models \ell_1 \boxminus \ell_2 \in \mathcal{F}} \qquad [\text{C-In}]\ \frac{\Pi \Vdash \ell' \succcurlyeq p^{\pi} \quad \pi \in \{\mathtt{i}, \mathtt{a}\}}{\ell' \in \mathcal{F} \models p^{\pi} \in \mathcal{F}}$$

$$[\text{C-Or}]\ \frac{\mathcal{C}_1 \models p^{\pi} \in \mathcal{F} \quad \mathcal{C}_2 \models p^{\pi} \in \mathcal{F} \quad \pi \in \{\mathtt{i}, \mathtt{a}\}}{\mathcal{C}_1 \ \texttt{OR} \ \mathcal{C}_2 \models p^{\pi} \in \mathcal{F}} \qquad [\text{C-And}]\ \frac{\exists i \in \{1,2\}.\, \mathcal{C}_i \models p^{\pi} \in \mathcal{F} \quad \pi \in \{\mathtt{i}, \mathtt{a}\}}{\mathcal{C}_1 \ \texttt{AND} \ \mathcal{C}_2 \models p^{\pi} \in \mathcal{F}}$$

Fig. 16. Blame membership: to apply C-In, C-Or and C-And the label p needs to be a primitive principal in $\mathcal{N} \cup \{\perp, \top\}$. The blame semantics rules ensure all statements added to the blame set only refer to primitive principals. This rule set differs from the originally published one [19], which didn't correctly handle compound principals such as $p \wedge q$.

$$[\text{C-CompareFail}]\ \frac{v_1 \neq v_2 \quad \mathcal{C}' := \mathcal{L}(v_1, v_2, \mathcal{C}, \ell_1, \ell_2)}{\langle\langle \texttt{compare}\ (\overline{\eta}_{\ell_1}\ v_1)\ \texttt{and}\ (\overline{\eta}_{\ell_2}\ v_2), c\rangle \,\&\, s\rangle^{\mathcal{C}} \Longrightarrow \langle\langle \texttt{fail}^{(\ell_1 \oplus \ell_2)}\ \texttt{says}\ \tau, c\rangle \,\&\, s\rangle^{\mathcal{C}'}}$$

Fig. 17. E-CompareFail with blame semantics.

constraints $\ell \in \mathcal{F}$ denote that label ℓ is in faulty set $\mathcal{F}$. This initial blame constraint ($\mathcal{C}_{\text{init}}$) is represented using the toleration set of the implied quorum system.

Definition 4 (Initial blame constraint)**.** For toleration set $[\![\mathcal{Q}]\!]$ of the form $\{(p_1^1 \wedge \ldots \wedge p_{m_1}^1)^{\mathtt{ia}}, \ldots, (p_1^k \wedge \ldots \wedge p_{m_k}^k)^{\mathtt{ia}}\}$ the initial blame constraint $\mathcal{C}_{\text{init}}$ is defined as a (logical) disjunction of conjunctions:

$$\mathcal{C}_{\text{init}} \triangleq (p_1^1 \in \mathcal{F}\ \texttt{AND} \ldots \texttt{AND}\ p_{m_1}^1 \in \mathcal{F}) \quad \texttt{OR} \quad \ldots \quad \texttt{OR} \quad (p_1^k \in \mathcal{F}\ \texttt{AND} \ldots \texttt{AND}\ p_{m_k}^k \in \mathcal{F})$$

Each disjunction represents a minimal subset of a possible satisfying assignment for the faulty set $\mathcal{F}$. For brevity, we will refer to these subsets as the *possible faulty sets* implied by a particular blame constraint. Observe that for quorum system $\mathcal{Q}$, there is a one-to-one correspondence between every $t_{\mathrm{i}} \in [\![\mathcal{Q}]\!]$ and every possible faulty set $\mathcal{F}_1, \ldots, \mathcal{F}_k$ in $\mathcal{C}_{\text{init}}$ where $\mathcal{F}_i$ is the set implied by the ith disjunction in $\mathcal{C}_{\text{init}}$ such that $t_{\mathrm{i}} = {b_{\mathrm{i}}}^{\mathtt{ia}}$, where $b_{\mathrm{i}} = \bigwedge_{p \in \mathcal{F}_{\mathrm{i}}} p$.

Evaluation rule C-CompareFail, in Fig. 17, shows how function $\mathcal{L}$ (discussed below) updates the blame constraint from $\mathcal{C}$ to $\mathcal{C}'$. We omit the blame-enabled versions of other evaluation rules since they simply propagate the blame constraint without modification.

Example.

(1) Quorum system $\mathcal{Q}_1 = \{q_1 = \{a, b\}; q_2 = \{b, c\}; q_3 = \{a, c\}\}$ has toleration set $[\![\mathcal{Q}_1]\!] = \{a^{\mathtt{ia}}, b^{\mathtt{ia}}, c^{\mathtt{ia}}\}$ and three possible faulty sets in $\mathcal{C}_{\text{init}}$: $\mathcal{F} = \{a\}$ or $\mathcal{F} = \{b\}$ or $\mathcal{F} = \{c\}$

(2) Quorum system $\mathcal{Q}_2 = \{q_1 := \{p, q\}; q_2 := \{r\}\}$ has toleration set $[\![\mathcal{Q}_2]\!] = \{p^{\mathtt{ia}} \wedge q^{\mathtt{ia}}, r^{\mathtt{ia}}\}$ and two possible faulty sets in $\mathcal{C}_{\text{init}}$: $\mathcal{F} = \{p, q\}$ or $\mathcal{F} = \{r\}$.

While $\mathcal{C}_{\text{init}}$ is defined statically according to the type of the program, rule C-CompareFail updates these constraints according to actual failures that occur during the program's execution. This approach identifies "unexpected" failures not implied by $\mathcal{C}_{\text{init}}$.

AUTHOR COPY

For example, $\mathcal{Q}_2 = \{q_1 := \{p, q\}; q_2 := \{r\}\}$ has two possible faulty sets $\mathcal{F} = \{p, q\}$ or $\mathcal{F} = \{r\}$. The initial blame constraint is $\mathcal{C}_{\text{init}} ::= (p \in \mathcal{F} \texttt{ AND } q \in \mathcal{F}) \texttt{ OR } (r \in \mathcal{F})$

Placing blame for a specific failure in a distributed system is challenging, (and often impossible!). For example, when a comparison of values signed by ℓ_1 and ℓ_2 fails, it is unclear who to blame since either principal (or a principal acting on their behalf) could have influenced the values that led to the failure. We do know, however, that at least one of them is faulty; recording this information helps constrain the contents of possible faulty sets.

We can reason about principals that *must* be in $\mathcal{F}$ by considering all possible faulty sets implied by the blame constraints. We write $\mathcal{C} \vDash \ell \in \mathcal{F}$ (read as $\mathcal{C}$ *entails* $\ell \in \mathcal{F}$), when every possible faulty set in $\mathcal{C}$, has the $\ell \in \mathcal{F}$ clause. Figure 16 presents inference rules for the $\vDash$ relation.

The rules C-IN, C-OR and C-AND are defined for a primitive principal p^{π} in $\mathcal{N} \cup \{\bot, \top\}$, where $\pi \in \{\texttt{i}, \texttt{a}\}$. Whereas rules C-CONJ, C-DISJ, C-PARAND and C-PAROR are defined for compound principals such as $p \wedge q$ and $p \boxminus q$. The blame semantics rules (particularly, the $\mathcal{L}$ and $\texttt{NORM}$ functions) ensure all statementss added to the blame set only refer to primitive principals. This rule set differs from the rule set presented in the originally published one [19], which didn't correctly handle compound principals such as $p \wedge q$. For example, if $\mathcal{C} \vDash p$ and $\mathcal{C} \vDash q$, then with the old ruleset from [19], we can not prove $\mathcal{C} \vDash p \wedge q$, because the blame semantics did not add the compound principal $p \wedge q$ to $\mathcal{F}$. Instead the blame semantics add $p \in \mathcal{F}$ and $q \in \mathcal{F}$ to the blame set as two different statements. But with our corrected ruleset we can prove $\mathcal{C} \vDash p \wedge q$, given $\mathcal{C} \vDash p$ and $\mathcal{C} \vDash q$ (using C-CONJ and C-IN).

Let us see another example. Since ℓ_1 is included in all satisfying choices of $\mathcal{F}$ below, we can say $\mathcal{C} \vDash \ell_1 \in \mathcal{F}$(using C-CONJ,C-IN, and possibly C-IN).

$$\begin{aligned} \mathcal{C} = (\ell_1 \in \mathcal{F} \texttt{ AND } \ell_2 \in \mathcal{F}) \quad & \texttt{OR} \quad (\ell_1 \in \mathcal{F} \texttt{ AND } \ell_3 \in \mathcal{F}) \\ \texttt{OR} \quad (\ell_1 \in \mathcal{F} \texttt{ AND } \ell_4 \in \mathcal{F}) \quad & \texttt{OR} \quad (\ell_1 \in \mathcal{F} \texttt{ AND } \ell_5 \in \mathcal{F}) \end{aligned}$$

The $\mathcal{L}$ function (full definition in Fig. 45) is used by rule C-COMPAREFAIL to update $\mathcal{C}$. For an expression:

$$\texttt{compare } (\overline{\eta}_{\ell_1}\ v_1) \texttt{ and } (\overline{\eta}_{\ell_2}\ v_2)$$

with $v_1 \neq v_2$, $\mathcal{L}(v_1, v_2, \mathcal{C}, \ell_1, \ell_2)$ updates the formulas in $\mathcal{C}$ to reflect that either ℓ_1 or ℓ_2 is faulty. If ℓ_1 or ℓ_2 already *must* be faulty, specifically if $\mathcal{C} \vDash \ell_1 \in \mathcal{F}$ or $\mathcal{C} \vDash \ell_2 \in \mathcal{F}$, then the function does not update any formulas. This approach avoids blaming honest principals when the other principal is already known to be faulty.

If neither ℓ_1 nor ℓ_2 are known to be faulty. then function $\mathcal{L}$ is called recursively on inner layers (i.e., nested $(\overline{\eta}\ )$ expressions) of v_1 and v_2 until a subexpression protected by a known-faulty principal is found. If no such layer is present, then the principal protecting the innermost layer is added to $\mathcal{C}$ (or the outer principals if there are no inner layers). Only this principal has seen the unprotected value and thus could have knowingly protected the wrong value. Observe that for well-typed `compare` expressions, only the outer layer of `compared` terms may differ in protection level, so there is less ambiguity when blaming an inner principal.

Updated constraints are kept in disjunctive normal form. Specifically, for compared terms $(\overline{\eta}_{\ell_1}\ v_1)$ and $(\overline{\eta}_{\ell_2}\ v_2)$, with $v_1 \neq v_2$, with initial constraint: $\mathcal{C}_{\text{init}} ::= (p \in \mathcal{F} \texttt{ AND } q \in \mathcal{F}) \texttt{ OR } (r \in \mathcal{F})$, then

AUTHOR COPY

$\mathcal{L}(v_1, v_2, \mathcal{C}_{\text{init}}, \ell_1, \ell_2)$ returns

$$\begin{aligned}\mathcal{C}' = {} & (p \in \mathcal{F} \texttt{ AND } q \in \mathcal{F} \texttt{ AND } \ell_1 \in \mathcal{F}) \\ & \texttt{OR} \quad (p \in \mathcal{F} \texttt{ AND } q \in \mathcal{F} \texttt{ AND } \ell_2 \in \mathcal{F}) \\ & \texttt{OR} \quad (r \in \mathcal{F} \texttt{ AND } \ell_1 \in \mathcal{F}) \quad \texttt{OR} \quad (r \in \mathcal{F} \texttt{ AND } \ell_2 \in \mathcal{F})\end{aligned}$$

We can now state the soundness theorem for our blame semantics, and apply it to prove a liveness result. Theorem 1 states that for any well-typed FLAQR program with a failing execution, and the faulty sets $\mathcal{F}_\text{i}$ implied by $\mathcal{C}'$ (the final constraint computed by the blame semantics), it must be the case that the program's type τ reflects the ability of the (possibly colluding) principals in $\mathcal{F}_\text{i}$ to fail the program.

Theorem 1 (Sound blame)**.** *Given,*

(1) $\Pi; \Gamma; pc; c \vdash \langle\langle e, c\rangle \;\&\; \texttt{empty}\rangle^{\mathcal{C}_{\text{init}}} : \tau$
(2) $\langle\langle e, c\rangle \;\&\; \texttt{empty}\rangle^{\mathcal{C}_{\text{init}}} \longrightarrow^* \langle\langle \texttt{fail}^\tau, c\rangle \;\&\; \texttt{empty}\rangle^{\mathcal{C}'}$

where e is a source-level expression,[7] *then for each possible faulty set $\mathcal{F}_i$ implied by $\mathcal{C}'$, there is a principal $b_i = \bigwedge_{p\in\mathcal{F}_i} p$ such that $\Pi \Vdash {b_i}^{\texttt{ia}} \succcurlyeq \tau$.*

Proof. Either e takes single step or multiple steps to produce the $\texttt{fail}^\tau$ term as the end result. For both the cases we prove it by induction over structure of e. See Appendix B for full proof. □

While Theorem 1 characterizes the relationship between a program's type and the possible faulty sets for a failing execution, it does not explicitly tell us anything about the fault-tolerance of a particular program. Since the type of a FLAQR program specifies its availability policy (in addition to its confidentiality and integrity), different FLAQR types will be tolerant of different failures. Below, we prove a liveness result for a common case, majority quorum protocols.

Definition 5 (Majority quorum system)**.** An m/n majority quorum system is a quorum system that always requires at least m of its hosts to reach consensus, where $m > n - m$.

Theorem 2 (Majority Liveness)**.** *If e is a source-level expression and:*

(1) $\Pi; \Gamma; pc; c \vdash \langle\langle e, c\rangle \;\&\; \texttt{empty}\rangle^{\mathcal{C}_{\text{init}}} : \tau$
(2) $\Pi \Vdash \mathcal{Q}$ *guards* τ
(3) $\mathcal{Q}$ *is a m/n majority quorum system*
(4) $\langle\langle e, c\rangle \;\&\; \texttt{empty}\rangle^{\mathcal{C}_{\text{init}}} \longrightarrow^* \langle\langle \texttt{fail}^\tau, c\rangle \;\&\; \texttt{empty}\rangle^{\mathcal{C}'}$

then for every possible faulty set $\mathcal{F}'$ implied by $\mathcal{C}'$, $|\mathcal{F}'| > (n - m)$.

Proof. From (2), we know τ is a valid quorum type for $\mathcal{Q}$ so $\forall \ell \in \mathcal{A}_{[\![\mathcal{Q}]\!]}.\Pi \nVdash \ell \succcurlyeq \tau$. Since $\mathcal{A}_{[\![\mathcal{Q}]\!]}$ is a superset of $[\![\mathcal{Q}]\!]$, we also have $\forall t \in [\![\mathcal{Q}]\!].\Pi \nVdash t \succcurlyeq \tau$. Furthermore, from Definition 4, for each possible faulty set $\mathcal{F}_i$ implied by $\mathcal{C}_{\text{init}}$, we know there is a principal $t_\text{i} \in [\![\mathcal{Q}]\!]$ such that $t_\text{i} = {b_\text{i}}^{\texttt{ia}}$, where $b_\text{i} = \bigwedge_{p\in\mathcal{F}_\text{i}} p$. Therefore, for each such b_i, we know $\Pi \nVdash {b_\text{i}}^{\texttt{ia}} \succcurlyeq \tau$.

Since $\mathcal{Q}$ is an m/n majority quorum system, every quorum is of size m and every faulty set in $\mathcal{C}_{\text{init}}$ is of size $(n - m)$. For contradiction, assume there exists a faulty set $\mathcal{F}'$ satisfying $\mathcal{C}'$ that has size

[7] In other words, e does not contain any $\texttt{fail}$ terms.

AUTHOR COPY

$(n-m)$. Then by the definition of $\mathcal{L}$, all possible faulty sets implied by $\mathcal{C}'$ also have size $(n-m)$ since $\mathcal{L}$ monotonically increases the size of all possible faulty sets or none of them. Furthermore, each possible faulty set implied by $\mathcal{C}_{\text{init}}$ is a subset (or equal to) a possible faulty set implied by $\mathcal{C}'$, so $|\mathcal{F}'| = (n-m)$ implies $\mathcal{C}_{\text{init}} = \mathcal{C}'$.

From Theorem 1 we know for every possible faulty set $\mathcal{F}'_{\text{i}}$ implied by $\mathcal{C}'$, it must be the case that $\Pi \Vdash {b'_{\text{i}}}^{\texttt{ia}} \geqslant \tau$, where $\bigwedge_{p \in \mathcal{F}'_{\text{i}}} p$. However, since $\mathcal{C}_{\text{init}} = \mathcal{C}'$, we have a contradiction since (2) implies $\Pi \nVdash {b'_{\text{i}}}^{\texttt{ia}} \geqslant \tau$. Thus there cannot exist a possible faulty set of size (at least) $(n-m)$ implied by $\mathcal{C}'$, and all possible faulty sets must have size greater than $(n-m)$. $\square$

7.2. Noninterference

We prove noninterference by extending the FLAQR syntax with bracketed expressions in the style of Pottier and Simonet [22]. Figure 43 shows selected bracketed evaluation rules and Fig. 42a and 42b show the typing rules for bracketed terms. The soundness and completeness of the bracketed semantics are proved the Appendix A (Lemmata 16–21).

Noninterference often is expressed with a distinct attacker label. We use H to denote the attacker. This means the attacker can read data with label ℓ if $\Pi \Vdash \ell^{\texttt{c}} \sqsubseteq H^{\texttt{c}}$ and can forge or influence it if $\Pi \Vdash H^{\texttt{i}} \sqsubseteq \ell^{\texttt{i}}$ and can make it unavailable if $\Pi \Vdash H^{\texttt{a}} \sqsubseteq \ell^{\texttt{a}}$.

An issue in typing brackets is how to deal with `fail` terms. Our confidentiality and integrity results are *failure-insensitive* in the sense that they only apply to terminating executions. This is similar to how termination-insensitive noninterference is typically characterized for potentially non-terminating programs.

Traditionally, bracketed typing rules require that bracketed terms have a restrictive type, ensuring that only values derived from secret (or untrusted) inputs are bracketed. In FLAQR, there are several scenarios where a bracketed value may not have a restrictive type. For example, when a `run` expression is evaluated within a bracket, it pushes an element onto the configuration stack, but only in one of the executions. Another example is when a bracketed value occurs in a `compare` expression, but the result is no longer influenceable by the attacker H. For these scenarios, several of the typing rules in Fig. 42a permit bracketed values to have less restrictive types. Because of these rules, subject reduction does not directly imply noninterference as it does in most bracketed approaches, but the additional proof obligations are relatively easy to discharge.

AUTHOR COPY

	Can have less restrictive type	
Term	$\pi = \texttt{i}$	$\pi = \texttt{a}$
$(v \mid v')$	No	Yes
$(v \mid \texttt{fail}^{\tau})$	Yes	No
$(v \mid v)$	Yes	Yes
$(\texttt{fail}^{\tau} \mid \texttt{fail}^{\tau})$	Yes	Yes

The table above summarizes how bracketed terms are typed depending on whether we are concerned with integrity or availability. For integrity, unequal bracketed values must have a restrictive type (i.e., one that protects H), but equal bracketed values may have a less restrictive type. For availability, only bracketed terms where one side contains a value and the other a failure must have a restrictive type.

7.2.1. Confidentiality and integrity noninterference

To prove confidentiality (integrity) noninterference we need to show that given two different secret (untrusted) inputs to an expression e the evaluated public (trusted) outputs are equivalent. Equivalence is defined in terms of an observation function $\mathcal{O}$ adapted from FLAC [2] in Appendix A, Fig. 44.

Theorem 3 (c-i Noninterference)**.** *If* $\Pi; \Gamma, x : \ell' \texttt{ says } \tau' \vdash \langle e, c\rangle \ \& \ \texttt{empty} : \ell \texttt{ says } \tau$ *where*

(1) $\Pi; \Gamma; pc; c \vdash v_i : \ell' \texttt{ says } \tau', i \in \{1, 2\}$
(2) $\langle e[x \mapsto (v_1 \mid v_2)], c\rangle \ \& \ \texttt{empty} \longrightarrow^* \langle v, c\rangle \ \& \ \texttt{empty}$
(3) $\Pi \Vdash H^{\pi} \sqsubseteq \ell'$ *and* $\Pi \nVdash H^{\pi} \sqsubseteq \ell, \pi \in \{\texttt{c}, \texttt{i}\}$.

then, $\mathcal{O}(\lfloor v\rfloor_1, \Pi, \ell, \pi) = \mathcal{O}(\lfloor v\rfloor_2, \Pi, \ell, \pi)$

Proof. From subject reduction we can prove that $\lfloor v\rfloor_1$ and $\lfloor v\rfloor_2$ have same type. By induction over the structure of projected values, $\lfloor v\rfloor_i$, we can show $\mathcal{O}(\lfloor v\rfloor_1, \Pi, \ell, \pi) = \mathcal{O}(\lfloor v\rfloor_2, \Pi, \ell, \pi)$ Please refer to the Appendix A for full proof. □

7.2.2. Availability noninterference

Similar to [30] our end-to-end availability guarantee is also expressed as noninterference property. Specifically, if one run of a well-typed FLAQR program running on a quorum system terminates successfully (does not fail), then all other runs of the program also terminate.

This approach treats "buggy" programs where every execution returns `fail` regardless of the choice of inputs as noninterfering. This behavior is desirable because here we are concerned with proving the absence of failures that attackers can *control*. For structured quorum systems with a liveness result such as Theorem 2 for m/n majority quorums, we can further constrain when failures may occur. For example, Theorem 2 proves failures can only occur when more than $(n - m)$ principals are faulty. In contrast, Theorem 4 applies to arbitrary quorum systems provided they guard the program's type, but cannot distinguish programs where all executions fail.

Theorem 4 (Availability Noninterference)**.** *If* $\Pi; \Gamma, x : \ell \texttt{ says } \tau' \vdash \langle e, c\rangle \ \& \ \texttt{empty} : \ell_{\mathcal{Q}} \texttt{ says } \tau$ *where*

(1) $\Pi; \Gamma; pc; c \vdash f_i : \ell \texttt{ says } \tau', i \in \{1, 2\}$
(2) $\langle e[x \mapsto (f_1 \mid f_2)], c\rangle \ \& \ \texttt{empty} \longrightarrow^* \langle f, c\rangle \ \& \ \texttt{empty}$
(3) $\Pi \Vdash H \succcurlyeq \ell \texttt{ says } \tau'$ *and* $H^{\texttt{ia}} \in \mathcal{A}_{[\![\mathcal{Q}]\!]}$ *and* $\Pi \Vdash \mathcal{Q}$ *guards* $(\ell_{\mathcal{Q}} \texttt{ says } \tau)$

then $\lfloor f\rfloor_1 \neq \texttt{fail}^{\ell_{\mathcal{Q}} \texttt{ says } \tau} \iff \lfloor f\rfloor_2 \neq \texttt{fail}^{\ell_{\mathcal{Q}} \texttt{ says } \tau}$

Proof. From subject reduction (see Lemma 25 in the Appendix) we know, $\Pi; \Gamma; pc; c \vdash \lfloor f\rfloor_i : \ell_{\mathcal{Q}} \texttt{ says } \tau$. Because $\Pi \Vdash \mathcal{Q}$ *guards* $(\ell_{\mathcal{Q}} \texttt{ says } \tau)$ and $H^{\texttt{ia}} \in \mathcal{A}_{[\![\mathcal{Q}]\!]}$ we can write $\Pi \nVdash H^{\texttt{ia}} \succcurlyeq \ell_{\mathcal{Q}} \texttt{ says } \tau$ from rule Q-GUARD. This ensures if $\lfloor f\rfloor_1 \neq \texttt{fail}^{\ell_{\mathcal{Q}} \texttt{ says } \tau}$, then $\lfloor f\rfloor_2 \neq \texttt{fail}^{\ell_{\mathcal{Q}} \texttt{ says } \tau}$, and vice-versa. □

8. Examples revisited

We are now ready to implement the examples from Section 2 with FLAQR semantics. To make these implementations intuitive we assume that our language supports integer (*int*) types, a mathematical

AUTHOR COPY

```
1 (λ(x:τ_a)[pc].λ(y:τ_b)[pc].λ(z:τ_c)[pc].
2   (select
3     (compare x and y)
4    or
5     (select
6       (compare y and z)
7      or
8       (compare x and z))))
9  (run^{τ_a} e_a@a) (run^{τ_b} e_b@b) (run^{τ_c} e_c@c)
```

Fig. 18. FLAQR implementation of majority quorum example.

operator > (greater than), and ternary operator :?. Beacuse `int` is a base type $\mathtt{C}(\mathtt{int})$ returns $\bot$. The examples also read from the local state of the participating principals. Which is fine because there are standard ways to encode memory (reads/writes) into lambda-calculus.

8.1. Tolerating failure and corruption

In this FLAQR implementation (Fig. 18) of 2/3 majority quorum example of Section 2.1, we refer principals representing `alice`, `bob` and `carol` as a, b and c respectively. The program is executed at host c' with program counter pc. Which means condition $\Pi \Vdash c' \succcurlyeq pc$ holds. The program body consists of a function of type $\tau_f = (\tau_a \xrightarrow{pc} \tau_b \xrightarrow{pc} \tau_c \xrightarrow{pc} (((a^{\mathtt{ia}} \oplus b^{\mathtt{ia}}) \ominus (b^{\mathtt{ia}} \oplus c^{\mathtt{ia}}) \ominus (a^{\mathtt{ia}} \oplus c^{\mathtt{ia}})) \ \mathtt{says}\ \tau))$ and the three arguments to the function are `run` statements. Here τ is $(a \wedge b \wedge c)^{\mathtt{c}}$ `says int`. Which means $\mathtt{C}(\tau_f) = pc$. The function body can be evaluated at c', as condition $\Pi \Vdash c' \succcurlyeq pc$ is true.

Here e_a, e_b and e_c are the expressions that read the balances for account *acct* from the local states of a, b and c respectively. The program counter at a, b, and c are a, b and c respectively. The data returned from a has type τ_a, which is basically $a^{\mathtt{ia}}\ \mathtt{says}\ \tau$. Similarly τ_b is $b^{\mathtt{ia}}\ \mathtt{says}\ \tau$ and τ_c is $c^{\mathtt{ia}}\ \mathtt{says}\ \tau$. Because each `run` returns a balance, the base type of τ is an *int* type, and it is protected with confidentiality label $(a \wedge b \wedge c)^{\mathtt{c}}$, meaning anyone who can read all the three labels (a, b and c), can read the returned balances.

In order to typecheck the `run` statements the conditions $\Pi \Vdash pc \sqsubseteq a$, $\Pi \Vdash pc \sqsubseteq b$, and $\Pi \Vdash pc \sqsubseteq c$ need to hold. The condition $\Pi \Vdash c' \succcurlyeq \mathtt{C}(\tau_a)$ is trivially true as $\mathtt{C}(\tau_a) = \bot$. Similarly $\mathtt{C}(\tau_b) = \bot$ and $\mathtt{C}(\tau_c) = \bot$ as well.

The host executing the code need to be able to read the return values from the three hosts. This means conditions $\Pi \Vdash c' \rhd a^{\mathtt{ia}}\ \mathtt{says}\ \tau$ $\Pi \Vdash c' \rhd b^{\mathtt{ia}}\ \mathtt{says}\ \tau$ and $\Pi \Vdash c' \rhd c^{\mathtt{ia}}\ \mathtt{says}\ \tau$ need to hold in order to typecheck the `compare` statements. The type of the whole program is $(((a^{\mathtt{ia}} \oplus b^{\mathtt{ia}}) \ominus (b^{\mathtt{ia}} \oplus c^{\mathtt{ia}}) \ominus (a^{\mathtt{ia}} \oplus c^{\mathtt{ia}}))\ \mathtt{says}\ \tau)$, which is a valid quorum type for $\mathcal{Q} = \{q_1 := \{a, b\}; q_2 := \{b, c\}; q_3 := \{a, c\}\}$.

Based on the security properties defined in Section 7 this program offers the confidentiality, integrity and availability guaranteed by quorum system $\mathcal{Q}$. Therefore, the result cannot be learned or influenced by unauthorized principals, and will be available as long as two hosts out of a, b, and c are non-faulty.

The toleration set here is $[\![\mathcal{Q}]\!] = \{a^{\mathtt{ia}}, b^{\mathtt{ia}}, c^{\mathtt{ia}}\}$. So, the program is not safe against an attacker with label $l_a = a^{\mathtt{ia}} \wedge b^{\mathtt{ia}}$ (or, $a^{\mathtt{i}} \wedge b^{\mathtt{a}}$), for example. This is because $\nexists t \in [\![\mathcal{Q}]\!].\Pi \Vdash t \succcurlyeq l_a$. Since $\Pi \Vdash l_a \succcurlyeq a^{\mathtt{ia}}$, principal l_a can fail two `compare` statements on lines 3 and 8. And, because $\Pi \Vdash l_a \succcurlyeq b^{\mathtt{ia}}$, l_a can also fail another two `compare` statements (one overlapping `compare` statment) on lines 3 and 6. Thus the whole program evaluates to `fail`. This FLAQR code also helps prevent incorrect comparisons. For instance, replacing z with y on line 8 will not typecheck.

AUTHOR COPY

```
1 (λ(arg₁:τ_b)[pc].(λ(arg₂:τ_b′)[pc].
2   (select
3     (bind x = arg₁ in (bind y = arg₂ in
4       (bind x′ = x in (bind y′ = y in
5         (η̄_d (η̄_(bᶜ∧b′ᶜ) (x′ > y′ ? x′ : y′)))))))
6    or
7   (select (arg₁) or (arg₂))))))(run^τ_b′ e′@b′))(run^τ_b e@b)
```

Fig. 19. FLAQR implementation of available largest balance example.

8.2. *Using best available services*

The code in Fig. 19 is the FLAQR implementation of Fig. 3. The program runs at a host c with program counter pc. The expressions e and e' read account balances from principals b and b', representing the banks. The values returned from b and b' have types $\tau_b = (b^{\mathtt{ia}}\ \mathtt{says}\ (b^{\mathtt{c}} \wedge b'^{\mathtt{c}})\ \mathtt{says}\ \mathtt{int})$ and $\tau_{b'} = (b'^{\mathtt{ia}}\ \mathtt{says}\ (b^{\mathtt{c}} \wedge b'^{\mathtt{c}})\ \mathtt{says}\ \mathtt{int})$ respectively.

The type of the whole program is $((d \ominus b^{\mathtt{ia}} \ominus b'^{\mathtt{ia}})\ \mathtt{says}\ (b^{\mathtt{c}} \wedge b'^{\mathtt{c}})\ \mathtt{says}\ int)$. Here $d = pc \sqcup b \sqcup b'$. In order to typecheck the `run` statements, the conditions $\Pi \Vdash pc \sqsubseteq b$ and $\Pi \Vdash pc \sqsubseteq b'$ need to hold. The program counter at b is b and b' is b'. The bind statements (lines 3–4) typecheck because conditions $\Pi \Vdash pc \sqcup b^{\mathtt{ia}} \sqsubseteq d$, $\Pi \Vdash pc \sqcup b^{\mathtt{ia}} \sqcup b'^{\mathtt{ia}} \sqsubseteq d$, $\Pi \Vdash pc \sqcup b^{\mathtt{ia}} \sqcup b'^{\mathtt{ia}} \sqcup b^{\mathtt{c}} \sqsubseteq d$, and $\Pi \Vdash pc \sqcup b^{\mathtt{ia}} \sqcup b'^{\mathtt{ia}} \sqcup b^{\mathtt{c}} \sqcup b'^{\mathtt{c}} \sqsubseteq d$ hold, because of our choice of d.

AUTHOR COPY

9. Secret sharing with **FLAQR**

Secret sharing is a cryptographic mechanism used in several distributed systems protocols such as oblivious transfer, multiparty communication, and Byzantine agreement. In this section we extend FLAQR with two new language constructs to support secret sharing. We call this extended version of our programming model FLAQR$^+$. Secret sharing is a process of splitting a secret into n shares and distributing the shares among n hosts (or principals in our setting). When an adequate number of hosts, say t, combine their respective shares, the secret is reconstructed [10,24]. Sometimes this process is referred as *(t,n)-threshold* secret sharing scheme [24,26], i.e. a quorum of t hosts (where $1 < t \leqslant n$) out of those n hosts need to combine their shares to retrieve the initial secret value. With $(t-1)$ or fewer shares, adversaries learn nothing about the secret.

Secret sharing is most commonly described in terms of a mathematical polynomial, say $p(x)$, of degree $(t-1)$, $p(0)$ being the secret value [24]. The polynomial's values at n different co-ordinates are distributed as the secret shares, and by knowing t of these values one can reconstruct[8] the polynomial and find the secret value $p(0)$.

For simplicity, we extend FLAQR to support *(2,2)-threshold* secret sharing, but later sketch how support for *(t,n)-threshold* secret sharing where $2 < n$ and $t < n$ would be straightforward. We model secret shares abstractly using a value sealed by new kinds of principals $k.L$ and $k.R$. We call k a *key principal* because, unlike the principals in $\mathcal{P}$, a new, unique key principal k is created each time secret shares are created. In contrast, the principals in $\mathcal{P}$ are statically known. The principals $k.L$ and $k.R$ represent the two associated shares of the key principal k. We will be refering to *(2,2)-threshold* secret sharing simply as *(2,2)* secret sharing in the following sections.

[8]Typically using Lagrange interpolation [10].

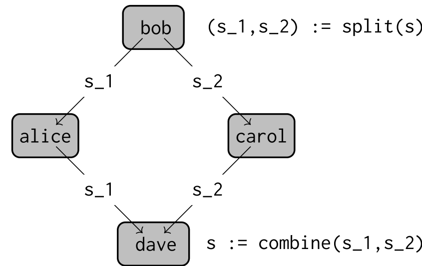

Fig. 20. Overview of (2, 2) secret sharing: `bob` shares his secret shares with `alice` and `carol`. Later, `alice` and `carol` forward their respective shares to `dave`. Finally, `dave` reproduces the initial secret `s` with the shares he received from `alice` and `carol`.

```
1  splitCombinePassword():
2    (s_1,s_2) := split(s) @ bob;
3    send s_1 to alice;
4    send s_2 to carol;
5    con := func(); // func() returns a bool
6    if(con)
7      fetch s_1 from alice;
8      fetch s_2 from carol;
9      s' := combine(s_1,s_2) @ dave;
10   else return;
```

Fig. 21. Creating two secret shares from a secret and then reconstructing the secret from the two secret shares using functions of a (2, 2) secret sharing protocol.

AUTHOR COPY

9.1. Motivating example of secret sharing : Password splitting

Figure 20 presents a simple *(2,2)* secret sharing example and the corresponding pseudocode is shown in Fig. 21. A (secret) password `s` belongs to `bob` who wants keep a backup of it. Bob creates two secret-shares of `s` as `s_1` and `s_2` and sends them to `alice` and `carol` respectively. That is, anyone who wants to get access to `s` has to either get it directly from `bob`, or needs to get both `s_1` and `s_2` from `alice` and `carol` and reproduce it. Later, `dave` fetches the secret shares from `alice` and `carol` and combines them to produce the password `s`.

An advantage of secret sharing is that it permits the secure transmission of secrets without requiring key distribution or public-key infrastructure (PKI). Instead, any party who possesses t shares may recover the secret. This can also be a liability, however, since an adversary needs to only obtain the shares to access the secret, too. Thus, considering the flow of shares between principals is central to the security of a secret sharing scheme. In the example in Fig. 21, `alice` and `carol` are unable to access the secret because they only possess a single share, but a coding error could transmit both shares to `alice` or `carol`, obviating the cryptographic protection. For this reason, embedding secret sharing in a language like FLAQR makes sense because the type system ensures the code only permits authorized flows.

9.2. (2, 2) *secret sharing in FLAQR*$^+$

Our abstractions for secret sharing in FLAQR$^+$make use of the $\overline{\eta}_\ell$ term to represent sealed secret shares. However, aspects of secret sharing schemes differ from the use of $\overline{\eta}_\ell$ in prior FLAC-based languages [2,8,14]. Here, in addition to the sealed values generated by the η_ℓ term, sealed values may also be created when splitting a secret into shares. In the previous approaches, a value sealed by η_ℓ serves

$$[\text{E-SPLIT}] \quad \frac{k \text{ is fresh}}{\texttt{split}_\ell\ v \longrightarrow \langle(\overline{\eta}_{k.L^c \wedge \ell}\ v), (\overline{\eta}_{k.R^c \wedge \ell}\ v)\rangle}$$

$$[\text{E-COMBINE}] \quad \texttt{combine}\ x = \langle(\overline{\eta}_{k.L^c \wedge \ell}\ v), (\overline{\eta}_{k.R^c \wedge \ell}\ v)\rangle @ pc\ \texttt{in}\ e \longrightarrow e[x \mapsto v]$$

Fig. 22. FLAQR$^+$ semantics for secret sharing (splitting secrets and combining shares).

as a reasonable model for signed and encrypted values. Specifically, the confidentiality component ℓ^c behaves like a public-key encrypted value: anyone can encrypt values at ℓ^c, but only authorized parties (which possess the associated private key) can distinguish the values protected at ℓ^c. Since η_ℓ can only be applied in contexts where $pc \succcurlyeq \ell^i$ (see UNITM), the integrity component behaves like a digitally signed value: only authorized principals can cause a value to be signed with ℓ^i integrity, but anyone can use high-integrity values.[9] Obviously then, enforcing these policies cryptographically would require public-key infrastructure.

Secret sharing behaves differently from the above interpretations: rather than authorization being based on possession of a long-lived private key (a reasonable proxy for identity), secret sharing implicitly authorizes *any party* possessing t shares. Therefore using identity-based principals such as `alice` or `bob` is inappropriate since, even if a shared secret is intended for `alice`, anyone with t shares will be able to distinguish the secret. Even `alice` must have t shares to access the secret. An abstraction for secret sharing should capture this behavior, but doing so in FLAQR$^+$requires new concepts.

We extend the set of principals $\mathcal{P}$ with new primitive principals L and R representing the left and right shares of our $(2, 2)$ secret sharing scheme. The set of all principals $\mathcal{P}$ for FLAQR$^+$is thus the closure of the set $\mathcal{N} \cup \{\top, \bot, L, R\}$ over the same operations as FLAQR. In the following we are only interested in the confidentiality projections L^c and R^c, since secret sharing only concerns enforcing the confidentiality of the secret.

Another aspect of secret sharing that departs from prior uses of FLAC principals is that each time shares are created, they are protected by a different secret. Consequently, shares created from different invocations cannot be mixed, even when the underlying value is the same. For this reason, we define *key principals*, a new type of principal generated dynamically at runtime. For our purposes, each $k \in \mathcal{K}$, where $\mathcal{K}$ is the set of all key principals, is equipped with a left and right principal, $k.L^c$ and $k.R^c$. Importantly, since key principals are generated dynamically, they are not directly representable statically. The principals L^c and R^c are the static representation for the left and right principals of any key principal, but the shares of different key principals cannot be distinguished at the type level.

AUTHOR COPY

9.3. Semantics and types for secret sharing

Figure 22 presents the semantic rules added to FLAQR$^+$. Expression $\texttt{split}_\ell\ v$ produces two secret shares, sealed with principals $k.L^c \wedge \ell$ and $k.R^c \wedge \ell$ from the secret value v (rule E-SPLIT) using a fresh key principal k. The ℓ annotation specifies an additional policy to seal the secret with an addition to the key principal. Primarily ℓ is used for integrity and availability components since $k.L^c$ and $k.R^c$ are confidentiality projections.

Two shares are combined with expression

$$\texttt{combine}\ x = \langle(\overline{\eta}_{k.L^c \wedge \ell}\ v), (\overline{\eta}_{k.R^c \wedge \ell}\ v)\rangle @ pc\ \texttt{in}\ e$$

[9] There doesn't appear to be a natural cryptographic analog for the availability component.

$$[\textsc{Split}]\ \frac{\Pi;\Gamma;pc;c \vdash e : \tau \qquad \Pi \Vdash c \succcurlyeq pc \qquad \Pi \Vdash pc \sqsubseteq \ell \sqcup \triangle(pc^{\mathsf{i}})}{\Pi;\Gamma;pc;c \vdash \mathtt{split}_{\ell}\ e : (L^{\mathsf{c}} \wedge \ell\ \mathsf{says}\ \tau \times R^{\mathsf{c}} \wedge \ell\ \mathsf{says}\ \tau)}$$

$$[\textsc{Combine}]\ \frac{\Pi;\Gamma;pc;c \vdash e : (L^{\mathsf{c}} \wedge \ell\ \mathsf{says}\ \tau \times R^{\mathsf{c}} \wedge \ell\ \mathsf{says}\ \tau) \qquad \Pi \Vdash c \succcurlyeq pc \qquad \Pi;\Gamma,x{:}\tau;\ell \sqcup pc;c \vdash e' : \ell'\ \mathsf{says}\ \tau \qquad \Pi \Vdash \ell \sqcup pc \sqsubseteq \ell'\ \mathsf{says}\ \tau}{\Pi;\Gamma;pc;c \vdash \mathtt{combine}\ x = e@pc\ \mathtt{in}\ e' : \ell'\ \mathsf{says}\ \tau}$$

$$[\textsc{SealedK}]\ \frac{\Pi;\Gamma;pc;c \vdash v : \tau \qquad \Pi \Vdash c \succcurlyeq pc \qquad K \in \{L^{\mathsf{c}}, R^{\mathsf{c}}\}}{\Pi;\Gamma;pc;c \vdash (\overline{\eta}_{k.K \wedge \ell}\ v) : K \wedge \ell\ \mathsf{says}\ \tau}$$

Fig. 23. FLAQR^{+} typing rules for secret sharing.

Rule E-COMBINE evaluates these terms to $e[x \mapsto v]$ revealing the secret v and substituting it for x in the body e. Notice that the key principal is the same on both sides of the pair. As discussed below in Section 9.4, mismatched key principals result in failure. The additional *pc* annotation on combine terms is used by the extended blame semantics, discussed in Section 9.4.

For simplicity, our extension only supports *(2,2)-threshold* secret sharing, but we believe extending this framework to support *(t,n)-threshold* secret sharing for $2 < n$ and $t < n$ would be straightforward. For example, given some t and n, we could redefine E-SPLIT to generate a tuple containing n shares sealed by principals $k.S_1^c, k.S_2^c, \ldots, k.S_n^c$. E-COMBINE would be replaced by $\binom{n}{t}$ rules: one for each valid t-sized subset of shares.

Figure 23 presents the FLAQR^{+} typing rules for `split` and `combine`. The last premise in the SPLIT rule involves the *view* of the *pc*'s integrity, $\triangle(pc^{\mathtt{i}})$. The view of a principal was introduced by Ceccetti et al. [8] to specify an upper bound on the confidentiality that may be *robustly declassified* [27] based on the integrity of the context performing the declassification and the data itself. These restrictions ensure an attacker cannot influence what (or whether) information is declassified. Below, we extend the definition of *view* with the principals' availability projection counterpart as well.

Definition 6 (*view* of a principal)**.** Let $\ell = p^{\mathtt{c}} \wedge q^{\mathtt{i}} \wedge r^{\mathtt{a}}$ be a FLAM [3] label (principal) expressed in normal form. The view of ℓ, written as $\triangle(\ell)$, is defined as $\triangle(p^{\mathtt{c}} \wedge q^{\mathtt{i}} \wedge r^{\mathtt{a}}) \triangleq q^{\mathtt{c}}$.

The premise $\Pi \Vdash pc \sqsubseteq \ell \sqcup \triangle(pc^{\mathtt{i}})$ in SPLIT serves two purposes. First, it ensures the confidentiality of control flow and the unsealed values in the context, represented by *pc*, are no more restrictive than the upper bound on declassification $\triangle(pc^{\mathtt{i}})$ (or the confidentiality of ℓ if no declassification takes place). Second, it ensures the label ℓ protects the availability and integrity of the context; only confidentiality may be downgraded by `split` terms.

When shares are combined to reveal the secret, the rule COMBINE ensures the combined pair contains a left and right share (although not which key principal they are associated with), and that the body of the `combine` term protects the result with a principal at least as restrictive as the upper bound of ℓ and the context *pc* the `combine` occurs in.

In some sense, $\mathtt{split}_{\ell}$ and `combine` function as an alternative to η_ℓ and `bind`. The difference is that `split` seals values using a key principal in addition to a label ℓ, and permits secrets more restrictive than ℓ to be sealed. Combine is similar to a `bind` that declassifies its bound value, dropping the key principals $k.L$ and $k.R$ from the protection requirements on the body of the `combine`. Note that it is not possible for a type-safe program to `bind` a secret share. Since a premise of bind would require the

AUTHOR COPY

$$[\text{E-SPLITFAIL}] \quad \frac{k \text{ is fresh}}{\texttt{split}_\ell\ \texttt{fail}^\tau \longrightarrow \texttt{fail}^{(L^c \wedge \ell \texttt{ says } \tau \times R^c \wedge \ell \texttt{ says } \tau)}}$$

$$[\text{E-COMBINEFAIL}] \quad \frac{k_1 \neq k_2}{\texttt{combine}\ x = \langle(\overline{\eta}_{k_1.L^c \wedge \ell}\ v), (\overline{\eta}_{k_2.R^c \wedge \ell}\ v)\rangle @pc\ \texttt{in}\ e \longrightarrow e[x \mapsto \texttt{fail}^\tau]}$$

$$[\text{E-COMBINEFAILL}] \quad \texttt{combine}\ x = \langle \texttt{fail}^{L^c \wedge \ell \texttt{ says } \tau}, f\rangle @pc\ \texttt{in}\ e \longrightarrow e[x \mapsto \texttt{fail}^\tau]$$

$$[\text{E-COMBINEFAILR}] \quad \texttt{combine}\ x = \langle v, \texttt{fail}^{R^c \wedge \ell \texttt{ says } \tau}\rangle @pc\ \texttt{in}\ e \longrightarrow e[x \mapsto \texttt{fail}^\tau]$$

Fig. 24. `fail` propagation rules in FLAQR$^+$.

$$[\text{C-COMBINEFAIL}] \quad \frac{k_1 \neq k_2 \quad \mathcal{C}' := \texttt{NORM}(pc, \mathcal{C})}{\langle\langle \texttt{combine}\ x = \langle(\overline{\eta}_{k_1.L^c \wedge \ell}\ v_1), (\overline{\eta}_{k_2.R^c \wedge \ell}\ v_2)\rangle @pc\ \texttt{in}\ (\overline{\eta}_\ell\ x), c\rangle \,\&\, s\rangle^{\mathcal{C}} \Longrightarrow \langle\langle \texttt{fail}^{\ell \texttt{ says } \tau}, c\rangle \,\&\, s\rangle^{\mathcal{C}'}}$$

Fig. 25. E-COMBINEFAIL with blame semantics.

body accessing the unsealed share on host c to typecheck at $pc \sqcup (L^c \wedge \ell)$, this would violate the invariant that c must act for the pc of the programs it executes.[10]

We need one more typing rule, SEALEDK, to preserve the types of the values sealed with labels $k.L^c$ and $k.R^c$, for any freshly generated key $k \in \mathcal{K}$. The existing SEALED rule is not enough as it does not handle the values protected with these new *key* principals. A consequence of this rule is that well-typed FLAQR$^+$ programs can produce mismatching shares with two different keys (say k_1 and k_2) during run-time. We handle mismatching shares with our extended blame semantics, discussed in Section 9.4.

AUTHOR COPY

9.4. Extending the blame semantics

As with other FLAQR terms, `fail` values propagate through `split` and `combine`. Figure 24 presents fail propagation rules for `split` and `combine` statements. These rules are straightforward propagation rules except for E-COMBINEFAIL (see Figure 25), which evaluates to `fail` if the key principals sealing the shares are mismatched.

The introduction of E-COMBINEFAIL rule creates an additional source of failure besides `compare` terms with mismatched values. Rule C-COMBINEFAIL extends FLAQR's blame semantics to account for this. Unlike the case for `compare`, we cannot blame the failure on the principal that sealed the mismatched values given to `combine`. The failure in this case is due to pairing together shares generated by different `split` evaluations. Rather than blaming the creators of the sealed value, we instead want to blame the principals that influenced this pairing. This influence is represented by the label of the pc. Hence, when k_1 and k_2 do not match, C-COMBINEFAIL adds pc to the blame set. The function $\mathcal{L}$ used in C-COMPAREFAIL is unnecessary here because it is unnecessary to examine any subterms of the combined pair – only the outer key principals contribute to a `combine` failure.

The NORM function[11] in the premise of C-COMBINEFAIL is used to add the new (potentially malicious) principal pc in the exisiting blame set $\mathcal{C}$ in normalized form. For example, if $pc = (a \wedge b) \vee c$ and if $\mathcal{C} := (\ell_1 \in \mathcal{F})\ \texttt{OR}\ (\ell_2 \in \mathcal{F})$, then calling $\texttt{NORM}(pc, \mathcal{C})$ will return a blame set

$$\mathcal{C}' := (a \in \mathcal{F}\ \text{AND}\ b \in \mathcal{F}\ \text{AND}\ \ell_1 \in \mathcal{F}) \quad \texttt{OR} \quad (a \in \mathcal{F}\ \text{AND}\ b \in \mathcal{F}\ \text{AND}\ \ell_2 \in \mathcal{F})$$

[10] Recall this invariant is enforced by the $\Pi \Vdash c \succcurlyeq pc$ premise included in all typing rules.

[11] This normalization function was not present in the original FLAQR publication [19], which is an error. Normalization of the statements added to the blame set is required to ensure compound principals are correctly handled.

$$[\text{B-SPLIT}] \quad \texttt{split}_\ell\,(v_1 \mid v_2) \longrightarrow \langle(\overline{\eta}_{k.L^c\wedge\ell}\,(v_1 \mid v_2)), (\overline{\eta}_{k.R^c\wedge\ell}\,(v_1 \mid v_2))\rangle$$

$$[\text{B-COMBINE}]\ \texttt{combine}\ x = (v_1 \mid v_2)@pc\ \texttt{in}\ e \longrightarrow (\texttt{combine}\ x = v_1@pc\ \texttt{in}\ e \mid \texttt{combine}\ x = v_2@pc\ \texttt{in}\ e)$$

Fig. 26. Bracketed semantics for FLAQR$^+$ terms.

$$\texttt{OR} \quad (c \in \mathcal{F}\ \text{AND}\ \ell_1 \in \mathcal{F}) \quad \texttt{OR} \quad (c \in \mathcal{F}\ \text{AND}\ \ell_1 \in \mathcal{F})$$

In case of C-COMPAREFAIL (Fig. 17), the NORM function was called from within the $\mathcal{L}$ function (see Fig. 45). Figure 46 contains the complete definition of the NORM function.

9.5. Security properties

By design, `split` and `combine` are interfering with respect to confidentiality: they can cause secret values to be declassified. However, we would like to ensure that integrity and availability noninterference are unaffected.

In order to prove integrity and availability noninterference for FLAQR$^+$ programs we extend the bracketed semantics (Fig. 26) and observation function (Fig. 27) with rules for `split` and `combine`, and add the corresponding cases for `split` and `combine` terms to the proofs for the lemmas and theorems of FLAQR$^+$. The noninterference theorem statements for FLAQR$^+$ are identical to Theorems 3 and 4, though Theorem 3 only holds for $\pi = \texttt{i}$ in FLAQR$^+$. Since the new static principals L and R are only used in confidentiality projections, rules such as Q-GUARD and the *fails* are unaffected by the new terms for secret sharing, the proofs of these theorems is largely unchanged from those for FLAQR. However, ensuring the new terms did not break an essential lemma such as subject reduction (Lemma 24) required careful design of the new rules for evaluation, failure propagation, bracketed semantics, and typing.

Although we protect the robustness (in theory) of what values may be declassified via `split` and `combine`, secret sharing is inherently non-robust since the party possessing the shares decides whether to reveal the secret. To formalize the protections that `split` and `combine` do offer, a weaker form of robust declassification would be required that permits secure uses of `split` and prohibits insecure ones (such as those violating the premise $\Pi \Vdash pc \sqsubseteq \ell \sqcup \Delta(pc^{\texttt{i}})$). Such a definition is not immediately clear[12] to us, and we leave further investigation to future work. Since `split` and `combine` permit non-robust declassification, they could potentially permit malleability attacks [8]. Since secret shares cannot be unsealed via `bind`, the possibility for such attacks is limited, but we leave formalization of the strength of these limitations to future work.

For `compare` statements, failures are generated because of two mismatching values. In contrast, the contents of the secret shares are irrelevant in `combine` statements. Instead, `combine` failures happen due to mismatching keys of the secret shares. Hence, we should only blame the control flow of the program for putting the two mismatching secret shares together. Since the program counter tracks the control flow of the program, we blame the program counter annotation *pc* in the C-COMBINEFAIL rule by adding it to the blame set when a `fail` term is returned while combining two shares. Our Theorem 1 (Sound blame) still holds, even though `combine` statement adds a new source of failure. The new cases hold because the premise $\Pi \Vdash pc \sqsubseteq \ell'\ \texttt{says}\ \tau$ in COMBINE allows us to show $\Pi \Vdash pc \geqslant \ell'\ \texttt{says}\ \tau$ (using P-LBL and A-AVAIL). Since, Theorem 1 holds, Theorem 2 (Majority liveness) holds as well, as it depends on Theorem 1.

[12] One reason such a definition is challenging in FLAC-based languages is that the *pc* label not only protects control flow, but also any values that have been unsealed using `bind` (which raises the *pc* label for its body).

AUTHOR COPY

$$\mathcal{O}(\texttt{split}_l\ e, \Pi, \ell, \pi) = \texttt{split}_l\ \mathcal{O}(e, \Pi, \ell, \pi)$$
$$\mathcal{O}(\texttt{combine}\ x = \langle e_1, e_2\rangle @pc\ \texttt{in}\ e, \Pi, \ell, \pi) =$$
$$\texttt{combine}\ x = \langle \mathcal{O}(e_1, \Pi, \ell, \pi), \mathcal{O}(e_2, \Pi, \ell, \pi)\rangle @pc\ \texttt{in}\ \mathcal{O}(e, \Pi, \ell, \pi)$$

Fig. 27. Observation function for intermediate FLAQR$^+$ terms (extended from FLAC [2]).

```
1 λ(arg:b^ia says (L^c ∧ b^ia says int × R^c ∧ b^ia says int))[pc].
2   (bind s = arg in
3     (bind s1 = (run^τa (proj1 s)@a) in
4       (bind s2 = (run^τc (proj2 s)@c) in
5         (combine sec = ⟨s1, s2⟩@pc in (η̄_ℓ sec)))))
6 (run^τb (split_b^ia v)@b)
```

where

$$\tau_a = a^{\text{ia}}\ \texttt{says}\ (L^{\text{c}} \wedge b\ \texttt{says int})$$
$$\tau_b = b^{\text{ia}}\ \texttt{says}\ (L^{\text{c}} \wedge b\ \texttt{says int} \times R^{\text{c}} \wedge b\ \texttt{says int})$$
$$\tau_c = c^{\text{ia}}\ \texttt{says}\ (R^{\text{c}} \wedge b\ \texttt{says int})$$

Fig. 28. A simple example of secret sharing in FLAQR$^+$.

9.6. *Password splitting example with FLAQR$^+$*

Figure 28 presents the FLAQR$^+$implementation of the example discussed in Section 9.1. The program executes at host c' with program counter pc, such that $\Pi \Vdash c' \succcurlyeq pc$. The host a has program counter a, host b has program counter b and host c has program counter c. The program consists of a function body (lines 1–5) and an argument to it (line 6). The function body is of type $\tau_b \xrightarrow{pc} \ell'$ says int, where τ_b = $b^{\texttt{ia}}$ says $(L^{\texttt{c}} \wedge b$ says int $\times$ $R^{\texttt{c}} \wedge b$ says int$)$, and takes the value of running a split statement at host b (i.e. $\texttt{run}^{\tau_b}\ (\texttt{split}_b\ v)@b$), which splits b's secret v. The argument type is τ_b. This means the pair of the secret shares created at and returned by b is tainted with b's integrity and availability. In order to typecheck the run statement pc needs to flow to b, i.e. the condition $\Pi \Vdash pc \sqsubseteq b$ needs to hold. The condition $\Pi \Vdash c \succcurlyeq \texttt{C}((L^{\texttt{c}} \wedge b\ \texttt{says int} \times R^{\texttt{c}} \wedge b\ \texttt{says int}))$ satisfies trivially, as $\texttt{C}((L^{\texttt{c}} \wedge b\ \texttt{says int} \times R^{\texttt{c}} \wedge b\ \texttt{says int})) = \bot$. The function body can be executed at c' as $\texttt{C}(\tau_b \xrightarrow{pc} \ell'\ \texttt{says int}) = pc$ and we mentioned earlier that the condition $\Pi \Vdash c' \succcurlyeq pc$ is true. The run statements on line 3 and 4 indicates that the left share is tainted by a's and the right share is tainted by c's integrity and availability. Which means a and c have seen and approved on the secret shares created by b. To make the run statements on lines 3 and 4 well-typed, the conditions $\Pi \Vdash pc \sqsubseteq a$ $\Pi \Vdash pc \sqsubseteq c$ should satisfy. We choose label ℓ such that $\Pi \Vdash pc \sqcup a^{\texttt{ia}} \sqcup b^{\texttt{ia}} \sqcup c^{\texttt{ia}} \sqsubseteq \ell$. The bind statements (lines 2–4) typecheck because the conditions $\Pi \Vdash pc \sqcup b^{\texttt{ia}} \sqsubseteq \ell$, $\Pi \Vdash pc \sqcup b^{\texttt{ia}} \sqcup a^{\texttt{ia}} \sqsubseteq \ell$ and $\Pi \Vdash pc \sqcup b^{\texttt{ia}} \sqcup a^{\texttt{ia}} \sqcup c^{\texttt{ia}} \sqsubseteq \ell$ hold due to our choice of ℓ.

10. Related work

FLAM [3,4] offers an algebra to integrate authorization logics and information flow control policies. FLAM also introduces a security condition, robust authorization, that is useful to ensure security when delegations and revocations change the meaning of confidentiality and integrity policies. In FLAQR we extend FLAM algebra with availability policies, and new binary operations to represent integrity and availability policies of the output of quorum based protocols. FLAC [2,5] embeds its types with FLAM

AUTHOR COPY

information flow policies. FLAC supports dynamic delegation of authority, but this feature is omitted in FLAQR.

A limited number of previous approaches [29,30] combine availability with more common confidentiality and integrity policies in distributed systems. Zheng and Myers [29] extend the Decentralized Label Model [20] with availability policies, but focus primarily tracking dependencies rather than applying mechanisms such as consensus and replication to improve availability and integrity. Zheng and Myers later introduce the language Qimp [30] with a type system explicitly parameterized on a quorum system for offloading computation while enforcing availability policies. Instead of treating quorums specially, FLAQR quorums emerge naturally using `compare` and `select` and enable application-specific integrity and availability policies that are secure by construction.

Hunt and Sands [15] present a novel generalisation of information flow lattices that captures disjunctive flows similar to the influence of replicas in FLAQR on a `select` result. Our partial-or operation was inspired by their treatment of disjunctive dependencies.

Models of distributed system protocols are often verified with model checking approaches such as TLA+ [17]. Model checking programs is typically undecidable, making it ill-suited to integrate directly into a programming model in the same manner as a (decidable) type system. To make verification tractable, TLA+ models are often simplified versions of the implementations they represent, potentially leading to discrepancies. FLAQR is designed as a core calculus for a distributed programming model, making direct verification of implementations more feasible.

BFT protocols [6,7] use consensus and replication to protect the integrity and availability of operations on a system's state. Each instance of a BFT protocol essentially enforces a single availability policy and a single integrity policy. While composing multiple instances is possible, doing so provides no end-to-end availability or integrity guarantees for the system as a whole. FLAQR programs, by constrast, routinely compose consensus and replication primitives to enforce multiple policies while also providing end-to-end system guarantees.

Our blame semantics presented in Section 7.1 has some resemblance to the idea of blame used to detect contract violations [11] and applied to gradual typing [25]. In our system, blame is necessarily ambiguous since perfect fault detection is not possible. Hence, rather than identifying a single program point responsible for a contract or type violation, our semantics builds constraints that specify a set of principals that may be responsible for a given failure.

In [9] Clarkson and Schneider talks about integrity measures such as contamination and suppression. The idea of suppression can be equivalent to the idea of unavailability. Although in [9] suppression happens due to untrusted input making trusted output unavailable. In FLAQR, unavailibity is caused by both unavailable and untrusted (via `compare` statement) inputs.

AUTHOR COPY

11. Conclusion

In this work, we extend Flow Limited Authorization Model [3] with availability policies. We introduce a core calculus and type-system, FLAQR, for building decentralized applications that are secure by construction. We identify a trade-off relation between integrity and availability, and introduce two binary operations *partial-and* and *partial-or*, specifically to express integrities of quorum based replicated programs. We define $fails$ relation and judgments that help us reason about a principal's authority over availability of a type. We introduce blame semantics that associate failures with malicious hosts of a quorum system to ensure that quorums can not exceed a bounded number of failures without causing the whole system to fail. FLAQR ensures end-to-end information security with noninterference for

confidentiality, integrity and availability. Finally we present FLAQR$^+$, which is an extension of FLAQR with language constructs that support secret sharing between hosts with mutual distrust. We extend our failure propagation rules and blame semantics to assign blame to appropriate principals when a secret sharing round fails.

Supplementary data

The supplementary material is available at: http://dx.doi.org/10.3233/JCS-230048.

Acknowledgments

Funding for this work was provided in part by NSF CAREER CNS-1750060 and IARPA HECTOR CW3002436.

AUTHOR COPY

References

[1] M. Abadi, Access control in a core calculus of dependency, in: *11th ACM SIGPLAN Int'l Conf. on Functional Programming*, ACM, New York, NY, USA, 2006, pp. 263–273. doi:10.1145/1159803.1159839.
[2] O. Arden, A. Gollamudi, E. Cecchetti, S. Chong and A.C. Myers, A calculus for flow-limited authorization, Technical Report, 2021. doi:10.48550/ARXIV.2104.10379.
[3] O. Arden, J. Liu and A.C. Myers, Flow-limited authorization, in: *28th IEEE Computer Security Foundations Symp*, CSF, 2015, pp. 569–583.
[4] O. Arden, J. Liu and A.C. Myers, Flow-limited authorization: Technical Report, Technical Report, 1813–40138, Cornell University Computing and Information Science, 2015.
[5] O. Arden and A.C. Myers, A calculus for flow-limited authorization, in: *29th IEEE Computer Security Foundations Symp*, CSF, 2016, pp. 135–147.
[6] A. Bessani, J. Sousa and E.E. Alchieri, State machine replication for the masses with BFT-SMaRt, in: *Dependable Systems and Networks (DSN), 2014 44th Annual IEEE/IFIP International Conference on*, IEEE, 2014, pp. 355–362.
[7] M. Castro and B. Liskov, Practical Byzantine fault tolerance and proactive recovery, *ACM Trans. on Computer Systems* **20** (2002), 2002.
[8] E. Cecchetti, A.C. Myers and O. Arden, Nonmalleable information flow control, in: *24th ACM Conf. on Computer and Communications Security (CCS)*, 2017, pp. 1875–1891.
[9] M.R. Clarkson and F.B. Schneider, Quantification of integrity, in: *Proc. IEEE Computer Security Foundations Symposium*, 2010, pp. 28–43.
[10] E. Dawson and D. Donovan, The breadth of Shamir's secret-sharing scheme, *Computer & Security* (1994).
[11] R.B. Findler and M. Felleisen, Contracts for higher-order functions, *SIGPLAN Not.* **37**(9) (2002), 48–59. doi:10.1145/583852.581484.
[12] J.-Y. Girard, Une extension de L'interpretation de gödel a L'analyse, et son application a L'elimination des coupures dans L'analyse et la theorie des types, in: *Studies in Logic and the Foundations of Mathematics*, Vol. 63, Elsevier, 1971, pp. 63–92.
[13] J.-Y. Girard, Interpretation fonctionelle et elimination des coupure dans l'arithmetic d'ordre superieur, Ph. D. Thesis, L'universite Paris VII, 1972.
[14] A. Gollamudi, S. Chong and O. Arden, Information flow control for distributed trusted execution environments, in: *32nd IEEE Computer Security Foundations Symp*, CSF, 2019.
[15] S. Hunt and D. Sands, A quantale of information, in: *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*, 2021. doi:10.1109/CSF51468.2021.00031.
[16] L. Lamport, The Part-time Parliament, *ACM Trans. on Computer Systems* **16**(2) (1998), 133–169. doi:10.1145/279227.279229.
[17] L. Lamport, The PlusCal algorithm language, in: *Theoretical Aspects of Computing – ICTAC 2009*, M. Leucker and C. Morgan, eds, Springer, Berlin, Heidelberg, 2009, pp. 36–60. ISBN 978-3-642-03466-4. doi:10.1007/978-3-642-03466-4_2.

[18] J. Liu, O. Arden, M.D. George and A.C. Myers, Fabric: Building open distributed systems securely by construction, *J. Computer Security* **25**(4–5) (2017), 319–321. doi:10.3233/JCS-0559.
[19] P. Mondal, M. Algehed and O. Arden, Applying consensus and replication securely with FLAQR, in: *IEEE Computer Security Foundations Symp (CSF)*, 2022, pp. 163–178. doi:10.1109/CSF54842.2022.9919637.
[20] A.C. Myers and B. Liskov, Protecting privacy using the decentralized label model, *ACM Transactions on Software Engineering and Methodology* **9**(4) (2000), 410–442. doi:10.1145/363516.363526.
[21] S. Nakamoto, Bitcoin: A peer-to-peer electronic cash system, *Consulted* **1**(2012) (2008), 28.
[22] F. Pottier and V. Simonet, Information flow inference for ML, in: *29th ACM Symp. on Principles of Programming Languages (POPL)*, 2002, pp. 319–330.
[23] J.C. Reynolds, Towards a theory of type structure, in: *Programming Symposium*, Springer, 1974, pp. 408–425. doi:10.1007/3-540-06859-7_148.
[24] A. Shamir, How to share a secret, *Communications of the ACM* **22**(11) (1979), 612–613. doi:10.1145/359168.359176.
[25] P. Wadler and R.B. Findler, Well-typed programs can't be blamed, in: *Programming Languages and Systems*, G. Castagna, ed., Springer, Berlin, Heidelberg, 2009, pp. 1–16. ISBN 978-3-642-00590-9.
[26] C.-C. Yang, T.-Y. Chang and M.-S. Hwang, A (t,n) multi-secret sharing scheme, in: *Applied Mathematics and Computation*, 2004, pp. 483–490.
[27] S. Zdancewic and A.C. Myers, Robust declassification, in: *14th IEEE Computer Security Foundations Workshop (CSFW)*, 2001, pp. 15–23, ISSN 1063-6900. doi:10.1109/CSFW.2001.930133.
[28] N. Zeldovich, S. Boyd-Wickizer and D. Mazières, Securing distributed systems with information flow control, in: *5th USENIX Symp. on Networked Systems Design and Implementation (NSDI)*, 2008, pp. 293–308.
[29] L. Zheng and A.C. Myers, End-to-end availability policies and noninterference, in: *18th IEEE Computer Security Foundations Workshop (CSFW)*, 2005, pp. 272–286. doi:10.1109/CSFW.2005.16.
[30] L. Zheng and A.C. Myers, A language-based approach to secure quorum replication, in: *9th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security (PLAS)*, 2014.

AUTHOR COPY