

Applying consensus and replication securely with FLAQR

Priyanka Mondal

University of California, Santa Cruz
pmondal@ucsc.edu

Maximilian Algehed

Chalmers University of Technology
algehed@chalmers.se

Owen Arden

University of California, Santa Cruz
owen@soe.ucsc.edu

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 by 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.

I. 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 [1] and PBFT [2], and blockchains such as Bitcoin [3] replicate state and computation at independent nodes and use consensus protocols to ensure the integrity and availability of operations on system state. In these protocols, there is neither centralization of function nor centralization of trust: all honest nodes 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 [4] and DStar [5]). 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 [6] 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 Figure 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 Figure 1a, 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 Figure 1b 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 Figure 1c, 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




Fig. 1: Integrity-Availability Trade-off

of the protocol or its output.

Figure 1c 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 1d 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 as follows:

- An extension of the static fragment of the Flow Limited Authorization Model (FLAM) [6] with availability policies and algebraic operators representing the effective authority of consensus and replication protocols (§III-§V).
- A formalization of the FLAQR language (§IV) and accompanying results:
 - A liveness theorem for majority-quorum FLAQR protocols (§VII-A) 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 (§VII-B).

Non-goals: The design of FLAQR is motivated by application-agnostic consensus protocols such as Paxos [1] and PBFT [2], 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 nodes 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

```

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

among the primary security concerns of high-level application developers.

II. 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.

A. 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 Figure 1c. If a majority of nodes agree on an account balance, then the system can tolerate the remaining minority of nodes failing or returning corrupted results.

Consider a quorum system with three nodes: alice, bob, and carol. To tolerate the failure of a single node, balance queries attempt to contact all three nodes 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 nodes (lines 2-4). The function returns the balance if two of the fetched values match, 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

```

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) and (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

to unilaterally choose the return value of the function, even when *alice* and *carol* agree on a different value.

B. Using best available services

Real world applications often consist of communication between entities with mutual distrust. The pseudocode in Figure 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

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.

III. SPECIFYING FLAQR POLICIES

FLAQR policies are specified using an extension of the FLAM [6], [7] principal algebra that includes availability

policies.¹ 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 *alice^c* refers to Alice’s confidentiality authority: what Alice may read. Principal *aliceⁱ* refers to Alice’s integrity authority: what Alice may write or influence.² Principal *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 *alice* \succcurlyeq *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 *alice* \wedge *bob*, denoting the combined authority of Alice and Bob. Put another way, *alice* \wedge *bob* is a principal both Alice and Bob *trust*. The disjunction of two principals’ authority is written using the connective \vee as *alice* \vee *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 *alice* \vee *bob*. Put another way *alice* \vee *bob* is a principal that trusts both Alice and Bob. Authority projections distribute over \wedge and \vee , so for example $(\text{alice} \wedge \text{bob})^i \equiv \text{alice}^i \wedge \text{bob}^i$.

The confidentiality, integrity, and availability authorities make up the totality of a principal’s authority, so writing *alice^c* \wedge *aliceⁱ* \wedge *alice^a* is equivalent to writing *alice*. For brevity, we sometimes write *alice^{ci}* as a shorthand for *alice^c* \wedge *aliceⁱ* when we wish to include all but one kind of authority. Due to space constraints, we refer the reader to FLAC [7] and our expanded technical report [8] for the complete formalization of the \succcurlyeq relation. In this article, we present only the extensions to this relation introduced by FLAQR.

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 Figure 1a. 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, (*aliceⁱ* \wedge *bobⁱ*). A principal with integrity authority (*aliceⁱ* \wedge *bobⁱ*) 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 Figure 1a 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 *aliceⁱ* \boxplus *bobⁱ*.

¹Specifically, we extend the static fragment of FLAM’s principal algebra defined by FLAC [7].

²Prior FLAM-based formalizations have used \rightarrow and \leftarrow for confidentiality and integrity, respectively.

Similarly, for replication protocols like that in Figure 1b, 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: $\text{alice}^i \vee \text{bob}^i$. The integrity of a value produced by either Alice or Bob is written as the *partial disjunction* $\text{alice}^i \boxplus \text{bob}^i$. 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 $\text{alice}^i \vee \text{bob}^i$, 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 $\{\mathcal{N}, \top, \perp\}$ over the operations $\wedge, \vee, \boxplus, \boxminus$, and authority projections c, i and a . We say Alice acts for Bob (or equivalently, Bob trusts Alice) and write $\text{alice} \succcurlyeq \text{bob}$ 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 \perp .

In addition to the trust relationships such as $p \wedge q \succcurlyeq p$ and $p \succcurlyeq p^i$ implied by the principal algebra, explicit delegations of trust such as $p \succcurlyeq q$ (for any p, q in \mathcal{P}) may be 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. [7] 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. The complete rule set is presented in Figure 23 in the appendices. In Figure 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 .

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 $\{\perp, x, y, \top\}$. Figure 5 shows the trust ordering of all possible distinct combinations of elements that can be formed on the set $\{\perp, x, y, \top\}$ with operations \wedge, \vee, \boxplus and \boxminus over them. The relationship between principals $\perp, x, y, x \wedge y, x \vee y$, and \top is the same as that in FLAM, but Figure 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 (see Figure 4), which lets us simplify $x \wedge (x \boxplus y) \wedge x \boxplus y$ 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 $\text{alice}^i \sqsubseteq \text{bob}^i$ if Bob trusts information influenced by Alice at least as much as information he influenced himself. Likewise, we write $\text{alice}^c \sqsubseteq \text{bob}^c$ if Alice trusts Bob to protect the confidentiality of her information, and $\text{alice}^a \sqsubseteq \text{bob}^a$ if Bob is trusted to keep Alice's data available. The flows-to relation behaves similarly to a sub-typing relation. Treating

information labeled $\text{alice}^{\text{cia}}$ (i.e. *alice*) as though it was labeled bob^{cia} (i.e. *bob*) is only safe (doesn't violate anyone's policies) if $\text{alice}^{\text{cia}} \sqsubseteq \text{bob}^{\text{cia}}$ (i.e. $\text{alice} \sqsubseteq \text{bob}$).

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.

$$\begin{aligned} p \sqsubseteq q &\Leftrightarrow q^c \succcurlyeq p^c \text{ and } p^i \succcurlyeq q^i \text{ and } p^a \succcurlyeq q^a \\ p \sqcup q &\triangleq (p^c \wedge q^c) \wedge (p^i \vee q^i) \wedge (p^a \vee q^a) \\ p \sqcap q &\triangleq (p^c \vee q^c) \wedge (p^i \wedge q^i) \wedge (p^a \wedge q^a) \end{aligned}$$

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^c is at least as confidential as p^c , q^i trusts information influenced by p^i , and q^a cannot cause failures that p^a cannot.

IV. FLAQR SYNTAX AND SEMANTICS

Figures 6 and 7 present the FLAQR syntax and selected evaluation rules. 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 the extended technical report [8].

FLAQR is based on FLAC [7], [9], a monadic calculus in the style of Abadi's Polymorphic DCC [10]. In addition to standard extensions to System F [11]–[13] 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 $\ell \text{ says } \tau$, 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 readers to the technical report [8] 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.

The monadic unit or return term $\eta_\ell e$ protects the value that e evaluates to at level ℓ (E-SEALED).³ Protected values, $(\bar{\eta}_\ell v)$ 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 Figure 7 are specific to FLAQR.

The primary novelty in the FLAQR calculus is the introduction of *compare* and *select* terms for expressing consensus

³Polymorphic DCC does not define a term similar to $(\bar{\eta}_\ell v)$ and thus does not have a 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.

$$\begin{array}{l}
\text{[PANDL]} \frac{\Pi \Vdash p_i \succcurlyeq p}{\Pi \Vdash p_1 \boxplus p_2 \succcurlyeq p} \quad k \in \{1, 2\} \\
\text{[PANDR]} \frac{\Pi \Vdash p \succcurlyeq p_1 \quad \Pi \Vdash p \succcurlyeq p_2}{\Pi \Vdash p \succcurlyeq p_1 \boxplus p_2} \\
\text{[ANDPAND]} \Pi \Vdash p \wedge q \succcurlyeq p \boxplus q \\
\text{[PANDPOR]} \Pi \Vdash p \boxplus q \succcurlyeq p \boxplus q \\
\text{[PROJPANDL]} \Pi \Vdash p^\pi \boxplus q^\pi \succcurlyeq (p \boxplus q)^\pi \\
\text{[PROJPANDR]} \Pi \Vdash (p \boxplus q)^\pi \succcurlyeq p^\pi \boxplus q^\pi \\
\text{[PROJPORL]} \Pi \Vdash p^\pi \boxplus q^\pi \succcurlyeq (p \boxplus q)^\pi \\
\text{[PROJPORR]} \Pi \Vdash (p \boxplus q)^\pi \succcurlyeq p^\pi \boxplus q^\pi \\
\text{[POROR]} \Pi \Vdash p \boxplus q \succcurlyeq p \vee q
\end{array}$$

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




Fig. 5: The FLAQR authority lattice for the principal set $\{\perp, x, y, \top\}$. 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, $\ell_1 \oplus \ell_2$ 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).

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

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

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

As discussed in Section III, 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

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

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

Fig. 6: FLAQR Syntax. Shaded terms are new to FLAQR. Underlined terms are used during evaluation and not available at the source level. integrity authority⁴ using a partial conjunction in Definition 1. In contrast, the decrease in availability is represented by a (full) $\ell_1^a \vee \ell_2^a$ 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 *fail*^τ 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 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 8 presents selected failure propagation rules (remaining

⁴Strictly speaking, $x \boxplus y$ is not an increase in integrity over x (or y); $x \boxplus y$ and x are incomparable.

[E-SEALED]	$\eta_\ell v \longrightarrow (\bar{\eta}_\ell v)$
[E-BINDM]	$\text{bind } x = (\bar{\eta}_\ell v) \text{ in } e \longrightarrow e[x \mapsto v]$
[E-COMPARE]	$\text{compare } (\bar{\eta}_{\ell_1} v) \text{ and } (\bar{\eta}_{\ell_2} v) \longrightarrow (\bar{\eta}_{\ell_1 \oplus \ell_2} v)$
[E-COMPAREFAIL]	$\frac{v_1 \neq v_2 \quad \tau = (\ell_1 \oplus \ell_2) \text{ says } \tau'}{\text{compare } (\bar{\eta}_{\ell_1} v_1) \text{ and } (\bar{\eta}_{\ell_2} v_2) \longrightarrow \text{fail}^\tau}$
[E-COMPAREFAILL]	$\frac{f_2 = \begin{cases} (\bar{\eta}_{\ell_2} v) \\ \text{fail}^{\ell_2} \text{ says } \tau \end{cases} \quad \tau_1 = \ell_1 \text{ says } \tau \quad \tau' = (\ell_1 \oplus \ell_2) \text{ says } \tau}{\text{compare } (\text{fail}^{\tau_1}) \text{ and } f_2 \longrightarrow \text{fail}^{\tau'}}$
[E-COMPAREFAILR]	$\frac{\tau_2 = \ell_2 \text{ says } \tau \quad \tau' = (\ell_1 \oplus \ell_2) \text{ says } \tau}{\text{compare } (\bar{\eta}_{\ell_1} v) \text{ and } \text{fail}^{\tau_2} \longrightarrow \text{fail}^{\tau'}}$
[E-SELECT]	$\frac{f_2 = \begin{cases} (\bar{\eta}_{\ell_2} v_2) \\ \text{fail}^{\ell_2} \text{ says } \tau \end{cases}}{\text{select } (\bar{\eta}_{\ell_1} v_1) \text{ or } f_2 \longrightarrow (\bar{\eta}_{\ell_1 \ominus \ell_2} v_1)}$
[E-SELECTR]	$\text{select } (\text{fail}^{\ell_1} \text{ says } \tau) \text{ or } (\bar{\eta}_{\ell_2} v) \longrightarrow (\bar{\eta}_{\ell_1 \ominus \ell_2} v)$
[E-SELECTFAIL]	$\frac{\forall i \in \{1, 2\} \quad \tau_i = \ell_i \text{ says } \tau \quad \tau' = (\ell_1 \ominus \ell_2) \text{ says } \tau}{\text{select } (\text{fail}^{\tau_1}) \text{ or } (\text{fail}^{\tau_2}) \longrightarrow \text{fail}^{\tau'}}$
[E-RETSTEP]	$\frac{e \longrightarrow e'}{\text{ret } e@c \longrightarrow \text{ret } e@c} \quad \text{[E-STEP]} \quad \frac{e \longrightarrow e'}{E[e] \longrightarrow E[e']}$
$E ::= [\cdot] \mid E e \mid v E \mid \eta_\ell E \mid \text{bind } x = E \text{ in } e$ $\quad \mid \text{ret } E@p \mid \text{select } E \text{ or } e \mid \text{select } f \text{ or } E$ $\quad \mid \text{compare } E \text{ and } e \mid \text{compare } f \text{ and } E$	

Fig. 7: FLAQR local semantics and evaluation context

[E-APPFAIL]	$\lambda(x:\tau)[pc]. e \text{ fail}^\tau \longrightarrow e[x \mapsto \text{fail}^\tau]$
[E-SEALEDFAIL]	$\eta_\ell \text{ fail}^\tau \longrightarrow \text{fail}^\ell \text{ says } \tau$
[E-INJFAIL]	$\text{inj}_i^{(\tau_1 + \tau_2)} \text{ fail}^{\tau_i} \longrightarrow \text{fail}^{(\tau_1 + \tau_2)}$
[E-PROJFAIL]	$\text{proj}_i \text{ fail}^{(\tau_1 \times \tau_2)} \longrightarrow \text{fail}^{\tau_i}$
[E-PAIRFAILL]	$\langle \text{fail}^{\tau_1}, f_2 \rangle^{(\tau_1 \times \tau_2)} \longrightarrow \text{fail}^{(\tau_1 \times \tau_2)}$

Fig. 8: fail propagation rules.

failure propagation rules are presented in Figure 19). 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-SELECTR returns its right subexpression when the left 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).

$$\ell_1 \ominus \ell_2 \triangleq (\ell_1^c \wedge \ell_2^c) \wedge (\ell_1^i \boxplus \ell_2^i) \wedge (\ell_1^a \wedge \ell_2^a)$$

We define the select action on types similarly to compare:

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

The end result of a *select* statement, **select** $(\bar{\eta}_{\ell_1} v) \text{ or } (\bar{\eta}_{\ell_2} v)$, will have integrity of either ℓ_1 or ℓ_2 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.

A. Global semantics

We capture the distributed nature of quorum replication by embedding the local semantic rules within a global distributed semantics defined in Figure 9. This semantics uses a *configuration stack* $s = \langle e, c \rangle \& t$ (Figure 10) 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 Figure 7 to identify the reducible subterms across stack elements.

The core operation for distributed computation is $\text{run}^\tau e@c$ 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 $(\bar{\eta}_{pc^{ia}} v)$. Rule E-RETFAIL serves the same purpose for **fail** terms, but is necessary since **fail** terms are not considered values (see Figure 6). The label pc^{ia} reflects both the integrity and availability context of the

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

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

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

Fig. 9: Global semantics

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

Fig. 10: Global configuration stack

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 V.

V. FLAQR TYPING RULES

As we have a local and global semantics, we have two corresponding forms of typing judgements: local typing judgements for expressions and global typing judgments for the stack. Local typing judgments 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 Figures 4 and 23. 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 11 presents a selection of local typing rules. Each typing rule includes an acts-for premise of the form $\Pi \Vdash c \succ pc$. This enforces the invariant that each host principal c has control of the program it executes locally. Thus for any judgment $\Pi; \Gamma; pc; c \vdash e : \tau$, the 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 $\Pi \Vdash c \succ pc$. The clearance of a type τ , written $C(\tau)$, is an upper bound on the pc annotations of the function types in τ . By checking that $\Pi \Vdash c \succ C(\tau_1 \xrightarrow{pc'} \tau_2)$ 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 $C(\cdot)$ is presented in Appendix A, Figure 18. 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 $\ell \text{ says } \tau$ 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 by policy ℓ if they are used by e . The SEALED rule types protected values $(\bar{\eta}_\ell v)$. 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 $\text{bind } x = e' \text{ in } e$. 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 $pc'^{ia} \text{ says } \tau'$, 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'^{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 well-typed at c and that c' is authorized to receive the return value based on the clearance of τ .

The COMPARE rule gives type $(\ell_1 \oplus \ell_2) \text{ says } \tau$ to the expression $\text{compare } e_1 \text{ and } e_2$ where e_1 and e_2 have types $\ell_1 \text{ says } \tau$ and $\ell_2 \text{ says } \tau$ respectively. Additionally, it requires that c , the host executing the compare, is authorized to fully examine the results of evaluating e_1 and e_2 so that they may be checked for equality.⁵ This requirement is captured by the premise $\Pi \Vdash c \triangleright \ell_i \text{ says } \tau$, pronounced “ c reads $\ell_i \text{ says } \tau$ ”. The inference rules for the reads judgment are found in Figure 21 in Appendix A.

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

The typing judgment for the global configuration is presented in Figure 12 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 $t \neq \text{empty}$) and EMP (when $t = \text{empty}$) to typecheck the tail t .

⁵Assuming 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.

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

$$\begin{array}{c}
\text{[UNIT]} \frac{\Pi \Vdash c \not\approx pc}{\Pi; \Gamma; pc; c \vdash () : \mathbf{unit}} \quad \text{[FAIL]} \frac{\Pi \Vdash c \not\approx pc}{\Pi; \Gamma; pc; c \vdash \mathbf{fail}^\tau : \tau} \quad \text{[EXPECT]} \frac{\Pi \Vdash c \not\approx pc}{\Pi; \Gamma; pc; c \vdash \mathbf{expect}^\tau : \tau} \\
\text{[LAM]} \frac{\Pi; \Gamma, x: \tau_1; pc'; u \vdash e : \tau_2 \quad \Pi \Vdash c \not\approx pc \quad u = C(\tau_1 \xrightarrow{pc'} \tau_2) \quad \Pi \Vdash c \not\approx u}{\Pi; \Gamma; pc; c \vdash \lambda(x: \tau_1)[pc']. e : \tau_1 \xrightarrow{pc'} \tau_2} \quad \text{[APP]} \frac{\Pi; \Gamma; pc; c \vdash e_1 : \tau' \xrightarrow{pc'} \tau \quad \Pi; \Gamma; pc; c \vdash e_2 : \tau' \quad \Pi \Vdash pc \sqsubseteq pc' \quad \Pi \Vdash c \not\approx pc}{\Pi; \Gamma; pc; c \vdash e_1 e_2 : \tau} \\
\text{[UNITM]} \frac{\Pi; \Gamma; pc; c \vdash e : \tau \quad \Pi \Vdash pc \sqsubseteq \ell}{\Pi; \Gamma; pc; c \vdash \eta_e e : \ell \mathbf{says} \tau} \quad \text{[SEALED]} \frac{\Pi; \Gamma; pc; c \vdash v : \tau \quad \Pi \Vdash c \not\approx pc}{\Pi; \Gamma; pc; c \vdash (\bar{\eta}_\ell v) : \ell \mathbf{says} \tau} \\
\text{[BINDM]} \frac{\Pi; \Gamma; pc; c \vdash e' : \ell \mathbf{says} \tau' \quad \Pi \Vdash \ell \sqcup pc \sqsubseteq \tau \quad \Pi; \Gamma, x: \tau'; \ell \sqcup pc; c \vdash e : \tau \quad \Pi \Vdash c \not\approx pc}{\Pi; \Gamma; pc; c \vdash \mathbf{bind} x = e' \mathbf{in} e : \tau} \quad \text{[RUN]} \frac{\Pi; \Gamma; pc'; c' \vdash e : \tau' \quad \Pi \Vdash pc \sqsubseteq pc' \quad \Pi \Vdash c \not\approx pc \quad \Pi \Vdash c \not\approx C(\tau') \quad \tau = pc'^{ia} \mathbf{says} \tau'}{\Pi; \Gamma; pc; c \vdash \mathbf{run}^\tau e@{c'} : \tau} \\
\text{[RET]} \frac{\Pi; \Gamma; pc; c \vdash e : \tau \quad \Pi \Vdash c' \not\approx C(\tau)}{\Pi; \Gamma; pc; c \vdash \mathbf{ret} e@{c'} : pc'^{ia} \mathbf{says} \tau} \quad \text{[COMPARE]} \frac{\forall i \in \{1, 2\}. \Pi; \Gamma; pc; c \vdash e_i : \ell_i \mathbf{says} \tau \quad \Pi \Vdash c \triangleright \ell_i \mathbf{says} \tau \quad \Pi \Vdash c \not\approx pc}{\Pi; \Gamma; pc; c \vdash \mathbf{compare} e_1 \mathbf{and} e_2 : (\ell_1 \oplus \ell_2) \mathbf{says} \tau} \\
\text{[SELECT]} \frac{\forall i \in \{1, 2\}. \Pi; \Gamma; pc; c \vdash e_i : \ell_i \mathbf{says} \tau \quad \Pi \Vdash c \not\approx pc}{\Pi; \Gamma; pc; c \vdash \mathbf{select} e_1 \mathbf{or} e_2 : (\ell_1 \ominus \ell_2) \mathbf{says} \tau}
\end{array}$$

Fig. 11: Typing rules for expressions

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

$$\text{[HEAD]} \frac{\Pi; \Gamma; pc'; c \vdash e : \tau' \quad \Pi; \Gamma; pc \vdash t : [\tau']\tau \quad \Pi \Vdash pc \sqsubseteq pc' \quad \Pi \Vdash c \not\approx pc}{\Pi; \Gamma; pc \vdash \langle e, c \rangle \& s : \tau}$$

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

$$\text{[TAIL]} \frac{\Pi; \Gamma; pc'; c \vdash E[\mathbf{expect}^{\tau'}] : \hat{\tau} \quad \Pi; \Gamma; pc \vdash t : [\hat{\tau}]\tau \quad \Pi \Vdash pc \sqsubseteq pc' \quad \Pi \Vdash c \not\approx pc}{\Pi; \Gamma; pc \vdash \langle E[\mathbf{expect}^{\tau'}], c \rangle :: t : [\tau']\tau}$$

$$\text{[EMP]} \Pi; \Gamma; pc \vdash \mathbf{empty} : [\tau]\tau$$

Fig. 12: Typing rules for configuration stack

$\langle E[\mathbf{expect}^{\tau'}], c \rangle :: t$ is well-typed with type $[\tau']\tau$, if expression $E[\mathbf{expect}^{\tau'}]$ 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 $[\tau]\tau$.

VI. 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 ℓ^{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* $\llbracket \mathcal{Q} \rrbracket$ of a quorum system \mathcal{Q} . The toleration

$$\begin{array}{l}
\text{[A-PAIR]} \frac{\Pi \Vdash \ell \triangleright \tau_i \quad i \in \{1, 2\}}{\Pi \Vdash \ell \triangleright (\tau_1 \times \tau_2)} \quad \text{[A-SUM]} \frac{\Pi \Vdash \ell \triangleright \tau_i \quad i \in \{1, 2\}}{\Pi \Vdash \ell \triangleright (\tau_1 + \tau_2)} \\
\text{[A-FUN]} \frac{\Pi \Vdash \ell \triangleright \tau_2}{\Pi \Vdash \ell \triangleright \tau_1 \xrightarrow{pc'} \tau_2} \quad \text{[A-TYPE]} \frac{\Pi \Vdash \ell \triangleright \tau}{\Pi \Vdash \ell \triangleright \text{says } \tau} \\
\text{[A-AVAIL]} \frac{\Pi \Vdash \ell^a \triangleright \ell'^a}{\Pi \Vdash \ell \triangleright \ell' \text{ says } \tau} \quad \text{[A-INTEGCOM]} \frac{\Pi \Vdash \ell^i \triangleright \ell_j^i, j \in \{1, 2\}}{\Pi \Vdash \ell \triangleright (\ell_1 \oplus \ell_2) \text{ says } \tau}
\end{array}$$

Fig. 13: fails judgments.

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.

- 1) The toleration set for quorum $\mathcal{Q}_1 = \{q_1 := \{a, b\}; q_2 := \{b, c\}; q_3 := \{a, c\}\}$ is $\llbracket \mathcal{Q}_1 \rrbracket = \{a^{ia}, b^{ia}, c^{ia}\}$,
- 2) For heterogeneous quorum system $\mathcal{Q}_2 = \{q_1 := \{p, q\}; q_2 := \{r\}\}$ the toleration set is $\llbracket \mathcal{Q}_2 \rrbracket = \{p^{ia} \wedge q^{ia}, r^{ia}\}$
- 3) For $\mathcal{Q}_3 = \{q := \{\text{alice}\}\}$ the toleration set is $\llbracket \mathcal{Q}_3 \rrbracket = \{\}$, 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

$$\mathcal{A}_{\llbracket \mathcal{Q} \rrbracket} = \{\ell \mid \exists \ell' \in \llbracket \mathcal{Q} \rrbracket. \Pi \Vdash \ell' \triangleright \ell\}.$$

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

The fails relation (\triangleright) 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 13 defines the fails judgment, written $\Pi \Vdash \ell \triangleright \tau$, which describes when a principal ℓ can fail an expression of type τ in delegation context Π .

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

In contrast, if $\Pi \Vdash \ell_a^a \triangleright \ell'^a$, and $\Pi \not\Vdash \ell_a^a \triangleright \ell'^a$, an availability attacker may cause $\eta_{\ell'}(e)$ to evaluate to `fail` ^{ℓ'} says τ , which steps to `fail` ^{ℓ} says (ℓ' says τ) 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 $\Pi \Vdash \ell_a \triangleright \ell \text{ says } (\ell' \text{ says } \tau)$.

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:

$$\text{[Q-GUARD]} \frac{\forall \ell \in \mathcal{A}_{\llbracket \mathcal{Q} \rrbracket}. \Pi \not\Vdash \ell \triangleright \tau}{\Pi \Vdash \mathcal{Q} \text{ guards } \tau}$$

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 $\Pi \Vdash \mathcal{Q} \text{ guards } \tau$ is satisfied.

$$\begin{array}{l}
\mathcal{C} ::= \mathcal{F} = \emptyset \mid \mathcal{B} \\
\mathcal{B} ::= \ell \in \mathcal{F} \mid \mathcal{B}_1 \text{ OR } \mathcal{B}_2 \mid \mathcal{B}_1 \text{ AND } \mathcal{B}_2
\end{array}$$

Fig. 14: Blame constraint syntax

$$\begin{array}{l}
\text{[C-IN]} \frac{\Pi \Vdash \ell' \triangleright \ell}{\ell' \in \mathcal{F} \Vdash \ell \in \mathcal{F}} \quad \text{[C-OR]} \frac{\mathcal{C}_1 \Vdash \ell \in \mathcal{F} \quad \mathcal{C}_2 \Vdash \ell \in \mathcal{F}}{\mathcal{C}_1 \text{ OR } \mathcal{C}_2 \Vdash \ell \in \mathcal{F}} \\
\text{[C-ANDL]} \frac{\exists i \in \{1, 2\}. \mathcal{C}_i \Vdash \ell \in \mathcal{F}}{\mathcal{C}_1 \text{ AND } \mathcal{C}_2 \Vdash \ell \in \mathcal{F}}
\end{array}$$

Fig. 15: Blame membership

Example 2. If $\mathcal{Q} = \{q_1 := \{a, b\}; q_2 := \{b, c\}; q_3 := \{a, c\}\}$ and $\ell_{\mathcal{Q}} = (a \oplus b) \oplus (b \oplus c) \oplus (a \oplus c)$ then $\ell_{\mathcal{Q}}$ says (a says τ) is not a valid quorum type because $\Pi \not\Vdash \mathcal{Q}$ guards ($\ell_{\mathcal{Q}}$ says (a says τ)) as $\Pi \Vdash a^{ia} \triangleright \ell_{\mathcal{Q}}$ says (a says τ) and $a^{ia} \in \mathcal{A}_{\llbracket \mathcal{Q} \rrbracket}$. But it is a valid quorum type for heterogeneous quorum system $\mathcal{Q}' = \{q_1 := \{a, b\}; q_2 := \{a, c\}\}$ as $a^{ia} \notin \mathcal{A}_{\llbracket \mathcal{Q}' \rrbracket}$.

VII. SECURITY PROPERTIES

To evaluate the formal properties of FLAQR, we prove that FLAQR preserves noninterference for confidentiality, integrity, and availability (section VII-B). 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.

A. 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 nodes \mathcal{F} . Figure 14 presents the syntax of *blame constraints*, which are boolean formulas representing a lower bound on the contents of \mathcal{F} . Atomic constraints $\ell \in \mathcal{F}$ denote that label ℓ is in faulty set \mathcal{F} . This initial blame constraint (\mathcal{C}_{init}) is represented using the toleration set of the implied quorum system.

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

$$\begin{array}{l}
\mathcal{C}_{init} \triangleq (p_1^1 \in \mathcal{F} \text{ AND } \dots \text{ AND } p_{m_1}^1 \in \mathcal{F}) \text{ OR } \dots \\
\text{OR } (p_1^k \in \mathcal{F} \text{ AND } \dots \text{ AND } p_{m_k}^k \in \mathcal{F})
\end{array}$$

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 $t_i \in \llbracket \mathcal{Q} \rrbracket$ and every possible faulty set $\mathcal{F}_1, \dots, \mathcal{F}_k$ in \mathcal{C}_{init} where \mathcal{F}_i is the set implied by the i^{th} disjunction in \mathcal{C}_{init} such that $t_i = b_i^{ia}$, where $b_i = \bigwedge_{p \in \mathcal{F}_i} p$.

Evaluation rule C-COMPAREFAIL, in Figure 16, 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 3.

- 1) Quorum system $\mathcal{Q}_1 = \{q_1 = \{a, b\}; q_2 = \{b, c\}; q_3 = \{a, c\}\}$ has toleration set $\llbracket \mathcal{Q}_1 \rrbracket = \{a^{ia}, b^{ia}, c^{ia}\}$ and three possible faulty sets in \mathcal{C}_{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 $\llbracket \mathcal{Q}_2 \rrbracket = \{p^{ia} \wedge q^{ia}, r^{ia}\}$ and two possible faulty sets in \mathcal{C}_{init} : $\mathcal{F} = \{p, q\}$ or $\mathcal{F} = \{r\}$.

While \mathcal{C}_{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}_{init} .

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 $\mathcal{C}_{init} ::= (p \in \mathcal{F} \text{ AND } q \in \mathcal{F}) \text{ OR } (r \in \mathcal{F})$

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 $\mathcal{C} \models \ell \in \mathcal{F}$ (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 15 presents inference rules for the \models relation.

For example, since ℓ_1 is included in all satisfying choices of \mathcal{F} below, we can say $\mathcal{C} \models \ell_1 \in \mathcal{F}$.

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

The \mathcal{L} function (full definition in Figure 22) is used by rule C-COMPAREFAIL to update \mathcal{C} . For an expression:

$$\text{compare } (\bar{\eta}_{\ell_1} v_1) \text{ and } (\bar{\eta}_{\ell_2} v_2)$$

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 $\mathcal{C} \models \ell_1 \in \mathcal{F}$ or $\mathcal{C} \models \ell_2 \in \mathcal{F}$, 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 $(\bar{\eta})$ 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 $(\bar{\eta}_{\ell_1} v_1)$ and $(\bar{\eta}_{\ell_2} v_2)$, with $v_1 \neq v_2$, with initial constraint: $\mathcal{C}_{init} ::= (p \in \mathcal{F} \text{ AND } q \in \mathcal{F}) \text{ OR } (r \in \mathcal{F})$, then $\mathcal{L}(v_1, v_2, \mathcal{C}_{init}, \ell_1, \ell_2)$ returns

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

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}_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}_i to fail the program.

Theorem 1 (Sound blame). *Given,*

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

where e is a source-level expression,⁶

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 $\Pi \Vdash b_i^{ia} > \tau$.

Proof. Either e takes single step or multiple steps to produce the fail^τ term as the end result. For both the cases we prove it by induction over structure of e . See [8] for full proof. \square

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) $\Pi; \Gamma; pc; c \vdash \langle \langle e, c \rangle \& \text{empty} \rangle^{\mathcal{C}_{init}} : \tau$

⁶In other words, e does not contain any fail terms.

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

Fig. 16: E-COMPAREFAIL with Blame Semantics.

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

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

Proof. From given condition (2), we know τ is a valid quorum type for \mathcal{Q} so $\forall \ell \in \mathcal{A}_{[\mathcal{Q}]}. \Pi \not\ll \ell > \tau$. Since $\mathcal{A}_{[\mathcal{Q}]}$ is a superset of $[\mathcal{Q}]$, we also have $\forall t \in [\mathcal{Q}]. \Pi \not\ll t > \tau$. Furthermore, from Definition 4, for each possible faulty set \mathcal{F}_i implied by \mathcal{C}_{init} , we know there is a principal $t_i \in [\mathcal{Q}]$ such that $t_i = b_i^{ia}$, where $b_i = \bigwedge_{p \in \mathcal{F}_i} p$. Therefore, for each such b_i , we know $\Pi \not\ll b_i^{ia} > \tau$.

Since \mathcal{Q} is an m/n majority quorum system, every quorum is of size m and every faulty set in \mathcal{C}_{init} is of size $(n - m)$. For contradiction, assume there exists a faulty set \mathcal{F}' satisfying \mathcal{C}' that has size $(n - m)$. 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}_{init} is a subset (or equal to) a possible faulty set implied by \mathcal{C}' , so $|\mathcal{F}'| = (n - m)$ implies $\mathcal{C}_{init} = \mathcal{C}'$.

From Theorem 1 we know for every possible faulty set \mathcal{F}'_i implied by \mathcal{C}' , it must be the case that $\Pi \not\ll b_i'^{ia} > \tau$, where $\bigwedge_{p \in \mathcal{F}'_i} p$. However, since $\mathcal{C}_{init} = \mathcal{C}'$, we have a contradiction since condition (2) implies $\Pi \not\ll b_i'^{ia} > \tau$. 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

B. Noninterference

We prove noninterference by extending the FLAQR syntax with bracketed expressions in the style of Pottier and Simonet [14]. Figure 25 shows selected bracketed evaluation rules and Figure 24a and 24b show the typing rules for bracketed terms. The soundness and completeness of the bracketed semantics are proved in [8].

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 $\Pi \Vdash \ell^c \sqsubseteq H^c$ and can forge or influence it if $\Pi \Vdash H^i \sqsubseteq \ell^i$ and can make it unavailable if $\Pi \Vdash H^a \sqsubseteq \ell^a$.

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 Figure 24a 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.

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

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 [7] (see Figure 26 in Appendix A).

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

- 1) $\Pi; \Gamma; pc; c \vdash v_i : \ell'$ says $\tau', i \in \{1, 2\}$
- 2) $\langle e[x \mapsto (v_1 \mid v_2)], c \rangle \& \text{empty} \longrightarrow^* \langle v, c \rangle \& \text{empty}$
- 3) $\Pi \Vdash H^\pi \sqsubseteq \ell'$ and $\Pi \not\ll H^\pi \sqsubseteq \ell, \pi \in \{c, i\}$.

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 $\mathcal{O}(\lfloor v \rfloor_1, \Pi, \ell, \pi) = \mathcal{O}(\lfloor v \rfloor_2, \Pi, \ell, \pi)$. See our technical report [8] for full proof. \square

2) *Availability Noninterference:* Similar to [15] 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*

$\Pi; \Gamma, x : \ell$ *says* $\tau' \vdash \langle e, c \rangle$ & *empty* : ℓ_Q *says* τ *where*

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

then $[f]_1 \neq \text{fail}^{\ell_Q}$ *says* $\tau \iff [f]_2 \neq \text{fail}^{\ell_Q}$ *says* τ

Proof. From subject reduction (see [8]) we know, $\Pi; \Gamma; pc; c \vdash [f]_i : \ell_Q$ *says* τ . Because $\Pi \Vdash Q$ *guards* (ℓ_Q *says* τ) and $H^{ia} \in \mathcal{A}_{[Q]}$ we can write $\Pi \not\ll H^{ia} \triangleright \ell_Q$ *says* τ from rule Q-GUARD. This ensures if $[f]_1 \neq \text{fail}^{\ell_Q}$ *says* τ , then $[f]_2 \neq \text{fail}^{\ell_Q}$ *says* τ , and vice-versa. \square

VIII. EXAMPLES REVISITED

We are now ready to implement the examples from section II with FLAQR semantics. To make these implementations intuitive we assume that our language supports integer (*int*) types, a mathematical operator $>$ (greater than), and ternary operator $?:$. Because *int* is a base type $C(\text{int})$ returns \perp . 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.

A. Tolerating failure and corruption

In this FLAQR implementation (Figure 17a) of 2/3 majority quorum example of section II-A, 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 $\Pi \Vdash c' \triangleright pc$ holds. The program body consists of a function of type $\tau_f = (\tau_a \xrightarrow{pc} \tau_b \xrightarrow{pc} \tau_c \xrightarrow{pc} ((a^{ia} \oplus b^{ia}) \oplus (b^{ia} \oplus c^{ia}) \oplus (a^{ia} \oplus c^{ia})) \text{ says } \tau)$ and the three arguments to the function are *run* statements. Here τ is $(a \wedge b \wedge c)^c$ *says* *int*. Which means $C(\tau_f) = pc$. The function body can be evaluated at c' , as condition $\Pi \Vdash c' \triangleright pc$ 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^{ia} *says* τ . Similarly τ_b is b^{ia} *says* τ and τ_c is c^{ia} *says* τ . Because each *run* returns a balance, the base type of τ is an *int* type, and it is protected with confidentiality label $(a \wedge b \wedge c)^c$, 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 $\Pi \Vdash c' \triangleright C(\tau_a)$ is trivially true as $C(\tau_a) = \perp$. Similarly $C(\tau_b) = \perp$ and $C(\tau_c) = \perp$ as well.

```

1  ( $\lambda(x:\tau_a)[pc]. \lambda(y:\tau_b)[pc]. \lambda(z:\tau_c)[pc].$ 
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 $\tau_a$   $e_a@_a$ ) (run $\tau_b$   $e_b@_b$ ) (run $\tau_c$   $e_c@_c$ )
```

(a) FLAQR implementation of majority quorum example

```

1  ( $\lambda(arg_1:\tau_b)[pc]. (\lambda(arg_2:\tau_{b'})[pc].$ 
2  (select
3  (bind  $x = arg_1$  in (bind  $y = arg_2$  in
4  (bind  $x' = x$  in (bind  $y' = y$  in
5  ( $\bar{\eta}_a (\bar{\eta}_{(b^c \wedge b'^c)} (x' > y' ? x' : y'))$ )))
6  or
7  (select ( $arg_1$ ) or ( $arg_2$ )))))(run $\tau_{b'}$   $e'_b@_b'$ )(run $\tau_b$   $e_b@_b$ )
```

(b) FLAQR implementation of available largest balance example

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

Based on the security properties defined in section VII this program offers the confidentiality, integrity and availability guaranteed by quorum system 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 $\ll[Q] = \{a^{ia}, b^{ia}, c^{ia}\}$. So, the program is not safe against an attacker with label $l_a = a^{ia} \wedge b^{ia}$ (or, $a^i \wedge b^a$), for example. This is because $\nexists t \in \ll[Q]. \Pi \Vdash t \triangleright l_a$. Since $\Pi \Vdash l_a \triangleright a^{ia}$, principal l_a can fail two *compare* statements on lines 3 and 8. And, because $\Pi \Vdash l_a \triangleright b^{ia}$, 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.

B. Using best available services

The code in Figure 17b is the FLAQR implementation of Figure 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 $\tau_b = (b^{ia} \text{ says } (b^c \wedge b'^c) \text{ says } \text{int})$ and $\tau_{b'} = (b'^{ia} \text{ says } (b^c \wedge b'^c) \text{ says } \text{int})$ respectively.

The type of the whole program is $((d \oplus b^{ia} \oplus b'^{ia}) \text{ says } (b^c \wedge b'^c) \text{ says } \text{int})$. 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 $\Pi \Vdash pc \sqcup b^{ia} \sqsubseteq d$, $\Pi \Vdash pc \sqcup b'^{ia} \sqcup b^{ia} \sqsubseteq d$, $\Pi \Vdash pc \sqcup b^{ia} \sqcup b'^{ia} \sqcup b^c \sqsubseteq d$, and $\Pi \Vdash pc \sqcup b^{ia} \sqcup b'^{ia} \sqcup b^c \sqcup b'^c \sqsubseteq d$ hold, because of our choice of d .

IX. RELATED WORK

FLAM [6] [16] 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 [9] [7] embeds its types with FLAM information flow policies. FLAC supports dynamic delegation of authority, but this feature is omitted in FLAQR.

A limited number of previous approaches [15], [17] combine availability with more common confidentiality and integrity policies in distributed systems. Zheng and Myers [17] extend the Decentralized Label Model [18] 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 [15] 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 [19] 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+ [20]. 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 [2], [21] 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 contrast, routinely compose consensus and replication primitives to enforce multiple policies while also providing end-to-end system guarantees.

Our blame semantics presented in Section VII-A has some resemblance to the idea of blame used to detect contract violations [22] and applied to gradual typing [23]. 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.

X. CONCLUSION

In this work, we extend Flow Limited Authorization Model [6] 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.

XI. ACKNOWLEDGEMENTS

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

REFERENCES

- [1] L. Lamport, "The Part-time Parliament," *ACM Trans. on Computer Systems*, vol. 16, no. 2, pp. 133–169, May 1998.
- [2] M. Castro and B. Liskov, "Practical Byzantine fault tolerance and proactive recovery," *ACM Trans. on Computer Systems*, vol. 20, 2002.
- [3] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," *Consulted*, vol. 1, no. 2012, p. 28, 2008.
- [4] J. Liu, O. Arden, M. D. George, and A. C. Myers, "Fabric: Building open distributed systems securely by construction," *J. Computer Security*, vol. 25, no. 4–5, pp. 319–321, May 2017.
- [5] 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.
- [6] O. Arden, J. Liu, and A. C. Myers, "Flow-limited authorization," in *28th IEEE Symp. on Computer Security Foundations (CSF)*, Jul. 2015, pp. 569–583.
- [7] O. Arden, A. Gollamudi, E. Cecchetti, S. Chong, and A. C. Myers, "A calculus for flow-limited authorization: Technical report," 2021.
- [8] P. Mondal, M. Algehed, and O. Arden, "Applying consensus and replication securely with flaqr," Tech. Rep., 2022.
- [9] O. Arden and A. C. Myers, "A calculus for flow-limited authorization," in *29th IEEE Symp. on Computer Security Foundations (CSF)*, Jun. 2016, pp. 135–147.
- [10] M. Abadi, "Access control in a core calculus of dependency," in *11th ACM SIGPLAN Int'l Conf. on Functional Programming*. New York, NY, USA: ACM, 2006, pp. 263–273.
- [11] J.-Y. Girard, "Une extension de l'interprétation de gödel a l'analyse, et son application a l'élimination des coupures dans l'analyse et la théorie des types," in *Studies in Logic and the Foundations of Mathematics*. Elsevier, 1971, vol. 63, pp. 63–92.
- [12] —, "Interprétation fonctionnelle et élimination des coupures dans l'arithmétique d'ordre supérieur," *Ph. D. Thesis, L'université Paris VII*, 1972.
- [13] J. C. Reynolds, "Towards a theory of type structure," in *Programming Symposium*. Springer, 1974, pp. 408–425.
- [14] F. Pottier and V. Simonet, "Information flow inference for ML," in *29th ACM Symp. on Principles of Programming Languages (POPL)*, 2002, pp. 319–330.
- [15] 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)*, Aug. 2014.
- [16] O. Arden, J. Liu, and A. C. Myers, "Flow-limited authorization: Technical report," Cornell University Computing and Information Science, Tech. Rep. 1813–40138, May 2015.

- [17] L. Zheng and A. C. Myers, "End-to-end availability policies and noninterference," in *18th IEEE Computer Security Foundations Workshop (CSFW)*, Jun. 2005, pp. 272–286.
- [18] A. C. Myers and B. Liskov, "Protecting privacy using the decentralized label model," *ACM Transactions on Software Engineering and Methodology*, vol. 9, no. 4, pp. 410–442, Oct. 2000.
- [19] S. Hunt and D. Sands, "A quantale of information," in *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*, 2021.
- [20] L. Lamport, "The pluscal algorithm language," in *Theoretical Aspects of Computing - ICTAC 2009*, M. Leucker and C. Morgan, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 36–60.
- [21] 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.
- [22] R. B. Findler and M. Felleisen, "Contracts for higher-order functions," *SIGPLAN Not.*, vol. 37, no. 9, p. 48–59, sep 2002.
- [23] P. Wadler and R. B. Findler, "Well-Typed Programs Can't Be Blamed," in *Programming Languages and Systems*, G. Castagna, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 1–16.

APPENDIX

$$\begin{aligned}
C(\tau_1 \xrightarrow{pc} \tau_2) &= C(\tau_1) \sqcup pc \sqcup C(\tau_2) \\
C(\forall X[pc]. \tau) &= pc \sqcup C(\tau) \\
C(\ell \text{ says } \tau) &= C(\tau) \\
C((\tau_1 + \tau_2)) &= C(\tau_1) \sqcup C(\tau_2) \\
C((\tau_1 \times \tau_2)) &= C(\tau_1) \sqcup C(\tau_2) \\
C(\text{unit}) &= \perp
\end{aligned}$$

Fig. 18: Clearance function

$$\begin{aligned}
[\text{E-APPFAIL}] \quad & \lambda(x:\tau)[pc]. \text{fail}^{\tau \xrightarrow{pc} \tau'} e \longrightarrow \text{fail}^{\tau'} \\
[\text{E-TAPPFAIL}] \quad & \text{fail}^{\forall X[pc]. \tau} \tau' \longrightarrow \text{fail}^{\tau[X \rightarrow \tau']} \\
[\text{E-CASEFAIL}] \quad & \text{case}^{\tau} \text{fail}^{\tau'} \text{ of } \text{inj}_1(x).e_1 \mid \text{inj}_2(x).e_2 \longrightarrow \text{fail}^{\tau} \\
[\text{E-PAIRFAIL}] \quad & \langle \text{fail}^{\tau_1}, f_2 \rangle^{(\tau_1 \times \tau_2)} \longrightarrow \text{fail}^{(\tau_1 \times \tau_2)} \\
[\text{E-PAIRFAILR}] \quad & \langle f_1, \text{fail}^{\tau_2} \rangle^{(\tau_1 \times \tau_2)} \longrightarrow \text{fail}^{(\tau_1 \times \tau_2)}
\end{aligned}$$

Fig. 19: Remaining cases for propagation of fail terms.

$$\begin{aligned}
\mathcal{C}(\text{unit}) &= \text{unit} \\
\mathcal{C}((\tau_1 + \tau_2)) &= (\mathcal{C}(\tau_1) + \mathcal{C}(\tau_2)) \\
\mathcal{C}((\tau_1 \times \tau_2)) &= (\mathcal{C}(\tau_1) \times \mathcal{C}(\tau_2)) \\
\mathcal{C}(\tau_1 \xrightarrow{pc} \tau_2) &= \mathcal{C}(\tau_1) \xrightarrow{pc} \mathcal{C}(\tau_2) \\
\mathcal{C}(\Lambda X[pc]. \tau) &= \Lambda X[pc]. \mathcal{C}(\tau) \\
\mathcal{C}((\ell_1 \boxplus \ell_2) \text{ says } \tau) &= (\ell_1 \vee \ell_2) \text{ says } \mathcal{C}(\tau) \\
\mathcal{C}((\ell_1 \boxminus \ell_2) \text{ says } \tau) &= (\ell_1 \wedge \ell_2) \text{ says } \mathcal{C}(\tau) \\
(\text{otherwise}) \mathcal{C}(\ell \text{ says } \tau) &= \ell \text{ says } \mathcal{C}(\tau)
\end{aligned}$$

Fig. 20: \mathcal{C} function on types.

$$\begin{aligned}
[\text{R-UNIT}] \quad & \Pi \Vdash p \triangleright () \\
[\text{R-TFUN}] \quad & \frac{\Pi \Vdash p \triangleright \tau}{\Pi \Vdash p \triangleright \forall X[pc]. \tau} \\
[\text{R-SUM}] \quad & \frac{\Pi \Vdash p \triangleright \tau_1 \quad \Pi \Vdash p \triangleright \tau_2}{\Pi \Vdash p \triangleright (\tau_1 + \tau_2)} \\
[\text{R-PROD}] \quad & \frac{\Pi \Vdash p \triangleright \tau_1 \quad \Pi \Vdash p \triangleright \tau_2}{\Pi \Vdash p \triangleright (\tau_1 \times \tau_2)} \\
[\text{R-LBL}] \quad & \frac{\Pi \Vdash p^c \succ \ell^c \quad \Pi \Vdash p \triangleright \tau}{\Pi \Vdash p \triangleright \ell \text{ says } \tau} \\
[\text{R-FUN}] \quad & \frac{\Pi \Vdash p \triangleright \tau_1 \quad \Pi \Vdash p \triangleright \tau_2}{\Pi \Vdash p \triangleright \tau_1 \xrightarrow{pc} \tau_2}
\end{aligned}$$

Fig. 21: Reads judgments.

$$\begin{aligned}
\mathcal{L}(x, y, \mathcal{C}, \ell_1, \ell_2) &= \text{match } (x, y) \text{ with} \\
& \mid ((\bar{\eta}_\ell v_1), (\bar{\eta}_\ell v_2)) = \\
& \quad \text{if } \mathcal{C} \Vdash \ell_1 \in \mathcal{F} \text{ then } \mathcal{C} \\
& \quad \text{else if } \mathcal{C} \Vdash \ell_2 \in \mathcal{F} \text{ then } \mathcal{C} \\
& \quad \text{else if } \mathcal{C} \Vdash \ell \in \mathcal{F} \text{ then } \mathcal{C} \\
& \quad \text{else } \mathcal{L}(v_1, v_2, \mathcal{C}, \ell, \ell) \\
& \mid (\eta_e e_1, \eta_e e_2) = \mathcal{L}(e_1, e_2, \mathcal{C}, \ell_1, \ell_2) \\
& \mid (\text{inj}_i^\tau e_1, \text{inj}_i^\tau e_2) = \mathcal{L}(e_1, e_2, \mathcal{C}, \ell_1, \ell_2) \\
& \mid \langle (e_{11}, e_{12})^\tau, \langle e_{21}, e_{22} \rangle^\tau \rangle = \\
& \quad \mathcal{L}(e_{11}, e_{21}, (\mathcal{L}(e_{12}, e_{22}, \mathcal{C}, \ell_1, \ell_2)), \ell_1, \ell_2) \\
& \mid (\text{run}^\tau e_1 @ p, \text{run}^\tau e_2 @ p) = \mathcal{L}(e_1, e_2, \mathcal{C}, \ell_1, \ell_2) \\
& \mid (\text{select}^\tau e_1 \text{ or } e_2, \text{select}^\tau e'_1 \text{ or } e'_2) = \\
& \quad \mathcal{L}(e_1, e'_1, (\mathcal{L}(e_2, e'_2, \mathcal{C}, \ell_1, \ell_2)), \ell_1, \ell_2) \\
& \mid (\text{compare}^\tau e_1 \text{ and } e_2, \text{compare}^\tau e'_1 \text{ and } e'_2) = \\
& \quad \mathcal{L}(e_1, e'_1, (\mathcal{L}(e_2, e'_2, \mathcal{C}, \ell_1, \ell_2)), \ell_1, \ell_2) \\
& \mid (\lambda(x:\tau)[pc]. e_1, \lambda(x:\tau)[pc]. e_2) = \mathcal{L}(e_1, e_2, \mathcal{C}, \ell_1, \ell_2) \\
& \mid (\Lambda X[pc]. e_1, \Lambda X[pc]. e_2) = \mathcal{L}(e_1, e_2, \mathcal{C}, \ell_1, \ell_2) \\
& \mid (\text{proj}_i e_1, \text{proj}_i e_2) = \mathcal{L}(e_1, e_2, \mathcal{C}, \ell_1, \ell_2) \\
& \mid (\text{bind } x_1 = e_1 \text{ in } e'_1, \text{bind } x_2 = e_2 \text{ in } e'_2) = \\
& \quad \mathcal{L}(e_1, e_2, \mathcal{L}(e'_1, e'_2, \mathcal{C}, \ell_1, \ell_2), \mathcal{C}, \ell_1, \ell_2) \\
& \mid (\text{case}^\tau e_1 \text{ of } \text{inj}_1^\tau(z).e_2 \mid \text{inj}_2^\tau(z).e_3, \\
& \quad \text{case}^\tau e'_1 \text{ of } \text{inj}_1^\tau(z).e'_2 \mid \text{inj}_2^\tau(z).e'_3) = \\
& \quad \text{last}(e_1, e'_1, \mathcal{L}(e_2, e'_2, \mathcal{L}(e_3, e'_3, \mathcal{C}, \ell_1, \ell_2), \ell_1, \ell_2), \ell_1, \ell_2), \ell_1, \ell_2) \\
& \mid (f_1, f_2) = \\
& \quad \text{if } f_1 = f_2 \text{ then } \mathcal{C} \\
& \quad \text{else if } \mathcal{C} \Vdash \ell_1 \in \mathcal{F} \text{ then } \mathcal{C} \\
& \quad \text{else if } \mathcal{C} \Vdash \ell_2 \in \mathcal{F} \text{ then } \mathcal{C} \\
& \quad \text{else } \text{DNF}(\ell_1 \in \mathcal{F} \text{ AND } \mathcal{C}) \text{ OR } \text{DNF}(\ell_2 \in \mathcal{F} \text{ AND } \mathcal{C}) \\
\text{DNF}(\ell \in \mathcal{F} \text{ AND } \mathcal{C}) &= \text{match } \mathcal{C} \text{ with} \\
& \mid \mathcal{F} = \emptyset \Rightarrow \ell \in \mathcal{F} \\
& \mid \ell' \in \mathcal{F} \Rightarrow \ell \in \mathcal{F} \text{ AND } \ell' \in \mathcal{F} \\
& \mid \mathcal{C}_1 \text{ OR } \mathcal{C}_2 \Rightarrow \text{DNF}(\ell \in \mathcal{F} \text{ AND } \mathcal{C}_1) \text{ OR } \text{DNF}(\ell \in \mathcal{F} \text{ AND } \mathcal{C}_2) \\
& \mid \mathcal{C}_1 \text{ AND } \mathcal{C}_2 \Rightarrow \mathcal{C}_1 \text{ AND } \mathcal{C}_2 \text{ AND } \ell \in \mathcal{F}
\end{aligned}$$

Fig. 22: Function to construct blame constraint \mathcal{C} .

$$\begin{array}{c}
\text{[PANDL]} \frac{\Pi \Vdash p_i \succcurlyeq p}{\Pi \Vdash p_1 \boxplus p_2 \succcurlyeq p} \quad \text{[PANDR]} \frac{\Pi \Vdash p \succcurlyeq p_1 \quad \Pi \Vdash p \succcurlyeq p_2}{\Pi \Vdash p \succcurlyeq p_1 \boxplus p_2} \\
\text{[ANDPAND]} \Pi \Vdash p \wedge q \succcurlyeq p \boxplus q \\
\text{[PANDPOR]} \Pi \Vdash p \boxplus q \succcurlyeq p \boxdot q \\
\text{[PROJPANDL]} \Pi \Vdash p^\pi \boxplus q^\pi \succcurlyeq (p \boxplus q)^\pi \\
\text{[PROJPANDR]} \Pi \Vdash (p \boxplus q)^\pi \succcurlyeq p^\pi \boxplus q^\pi \\
\text{[PROJPORL]} \Pi \Vdash p^\pi \boxdot q^\pi \succcurlyeq (p \boxdot q)^\pi \\
\text{[PROJPORR]} \Pi \Vdash (p \boxdot q)^\pi \succcurlyeq p^\pi \boxdot q^\pi \\
\text{[POROR]} \Pi \Vdash p \boxdot q \succcurlyeq p \vee q \\
\text{[ANDDISTPORR]} \Pi \Vdash p \wedge (q \boxdot r) \succcurlyeq (p \wedge q) \boxdot (p \wedge r) \\
\text{[PORDISTANDR]} \Pi \Vdash p \boxdot (q \wedge r) \succcurlyeq (p \boxdot q) \wedge (p \boxdot r) \\
\text{[ANDDISTPORL]} \Pi \Vdash (p \wedge q) \boxdot (p \wedge r) \succcurlyeq p \wedge (q \boxdot r) \\
\text{[PORDISTANDL]} \Pi \Vdash (p \boxdot q) \wedge (p \boxdot r) \succcurlyeq p \boxdot (q \wedge r) \\
\text{[ORDISTPORR]} \Pi \Vdash p \vee (q \boxdot r) \succcurlyeq (p \vee q) \boxdot (p \vee r) \\
\text{[ORDISTPORL]} \Pi \Vdash (p \vee q) \boxdot (p \vee r) \succcurlyeq p \vee (q \boxdot r) \\
\text{[PORDISTORR]} \Pi \Vdash p \boxdot (q \vee r) \succcurlyeq (p \boxdot q) \vee (p \boxdot r) \\
\text{[PORDISTORL]} \Pi \Vdash (p \boxdot q) \vee (p \boxdot r) \succcurlyeq p \boxdot (q \vee r) \\
\text{[ANDDISTPANDR]} \Pi \Vdash p \wedge (q \boxplus r) \succcurlyeq (p \wedge q) \boxplus (p \wedge r) \\
\text{[PANDDISTANDR]} \Pi \Vdash p \boxplus (q \wedge r) \succcurlyeq (p \boxplus q) \wedge (p \boxplus r) \\
\text{[ANDDISTPANDL]} \Pi \Vdash (p \wedge q) \boxplus (p \wedge r) \succcurlyeq p \wedge (q \boxplus r) \\
\text{[PANDDISTANDL]} \Pi \Vdash (p \boxplus q) \wedge (p \boxplus r) \succcurlyeq p \boxplus (q \wedge r) \\
\text{[ORDISTPANDR]} \Pi \Vdash p \vee (q \boxplus r) \succcurlyeq (p \vee q) \boxplus (p \vee r) \\
\text{[ORDISTPANDL]} \Pi \Vdash (p \vee q) \boxplus (p \vee r) \succcurlyeq p \vee (q \boxplus r) \\
\text{[PANDDISTORR]} \Pi \Vdash p \boxplus (q \vee r) \succcurlyeq (p \boxplus q) \vee (p \boxplus r) \\
\text{[PANDDISTORL]} \Pi \Vdash (p \boxplus q) \vee (p \boxplus r) \succcurlyeq p \boxplus (q \vee r)
\end{array}$$

$$\begin{array}{c}
\text{[BRACKET]} \frac{\Pi \Vdash (H^\pi \sqcup pc) \sqsubseteq pc' \quad e_1 = v_1 \iff e_2 \neq v_2 \quad \Pi; \Gamma; pc'; c \vdash e_1 : \tau \quad \Pi; \Gamma; pc'; c \vdash e_2 : \tau}{\Pi \Vdash H^\pi \sqsubseteq \mathcal{C}(\tau) \quad \Pi \Vdash c \succcurlyeq pc} \\
\text{[BRACKET-VALUES]} \frac{\Pi; \Gamma; pc; c \vdash v_1 : \tau \quad \Pi; \Gamma; pc; c \vdash v_2 : \tau}{\Pi \Vdash H^\pi \sqsubseteq \mathcal{C}(\tau) \quad \Pi \Vdash c \succcurlyeq pc} \\
\text{[BULLR]} \frac{\Pi; \Gamma; pc; c \vdash e : \tau}{\Pi; \Gamma; pc; c \vdash (e \mid \bullet) : \tau} \quad \text{[BULLL]} \frac{\Pi; \Gamma; pc; c \vdash e : \tau}{\Pi; \Gamma; pc; c \vdash (\bullet \mid e) : \tau} \\
\text{[BRACKET-FAIL-L]} \frac{\Pi; \Gamma; pc; c \vdash e : \tau}{\Pi; \Gamma; pc; c \vdash (e \mid \text{fail}^\tau) : \tau} \\
\text{[BRACKET-FAIL-R]} \frac{\Pi; \Gamma; pc; c \vdash e : \tau}{\Pi; \Gamma; pc; c \vdash (\text{fail}^\tau \mid e) : \tau} \\
\text{[BRACKET-FAIL-A]} \frac{\Pi; \Gamma; pc; c \vdash e_i : \tau \quad e_i \neq \text{fail}^\tau \quad \pi = \mathbf{a}}{\Pi; \Gamma; pc; c \vdash (e_1 \mid e_2) : \tau} \\
\text{[BRACKET-SAME]} \frac{\Pi; \Gamma; pc; c \vdash v : \tau}{\Pi; \Gamma; pc; c \vdash (v \mid v) : \tau} \\
\text{(a) Typing rules for bracketed expressions.} \\
\text{[BRACKET-STACK]} \frac{\Pi; \Gamma; pc'; c \vdash e : \tau' \quad \Pi \Vdash pc \sqsubseteq pc' \quad \forall i \in \{1, 2\}. \Pi; \Gamma; pc \vdash s_i : [\tau']\tau}{\Pi; \Gamma; pc \vdash \langle e, c \rangle \& (s_1 \mid s_2) : \tau} \\
\text{[BRACKET-HEAD]} \frac{\Pi; \Gamma; pc'; c \vdash (e_1 \mid e_2) : \tau' \quad \Pi \Vdash pc \sqsubseteq pc' \quad \Pi; \Gamma; pc \vdash s : [\tau']\tau}{\Pi; \Gamma; pc \vdash \langle (e_1 \mid e_2), c \rangle \& s : \tau} \\
\text{(b) Typing rules for bracketed configuration stack.}
\end{array}$$

Fig. 23: FLAQR Partial conjunction and disjunction acts-for rules.

$$\begin{array}{l}
\text{[B-COMPARECOMMON]} \quad \frac{\lfloor \text{compare}^{\ell_1 \oplus \ell_2} \text{ says } \tau (f_{11} \mid f_{12}) \text{ and } (f_{21} \mid f_{22}) \rfloor_i \longrightarrow f_i \quad \forall i \in \{1, 2\}}{\text{compare}^{\ell_1 \oplus \ell_2} \text{ says } \tau (f_{11} \mid f_{12}) \text{ and } (f_{21} \mid f_{22}) \longrightarrow (f_1 \mid f_2)} \\
\text{[B-COMPARECOMMONRIGHT]} \quad \frac{\lfloor \text{compare}^{\ell_1 \oplus \ell_2} \text{ says } \tau (f_{11} \mid f_{12}) \text{ and } f \rfloor_i \longrightarrow f_i \quad \forall i \in \{1, 2\}}{\text{compare}^{\ell_1 \oplus \ell_2} \text{ says } \tau (f_{11} \mid f_{12}) \text{ and } f \longrightarrow (f_1 \mid f_2)} \\
\text{[B-SELECTCOMMON]} \quad \frac{\lfloor \text{select}^{\ell_1 \oplus \ell_2} \text{ says } \tau (f_{11} \mid f_{12}) \text{ or } (f_{21} \mid f_{22}) \rfloor_i \longrightarrow f_i \quad \forall i \in \{1, 2\}}{\text{select}^{\ell_1 \oplus \ell_2} \text{ says } \tau (f_{11} \mid f_{12}) \text{ or } (f_{21} \mid f_{22}) \longrightarrow (f_1 \mid f_2)} \\
\text{[B-SELECTCOMMONLEFT]} \quad \frac{\lfloor \text{select}^{\ell_1 \oplus \ell_2} \text{ says } \tau (f_{11} \mid f_{12}) \text{ or } f \rfloor_i \longrightarrow f_i \quad \forall i \in \{1, 2\}}{\text{select}^{\ell_1 \oplus \ell_2} \text{ says } \tau (f_{11} \mid f_{12}) \text{ or } f \longrightarrow (f_1 \mid f_2)} \\
\text{[B-FAIL1]} \quad \eta_\ell (v \mid \text{fail}^\tau) \longrightarrow ((\bar{\eta}_\ell v) \mid \text{fail}^\ell \text{ says } \tau) \quad \text{[B-FAIL2]} \quad \eta_\ell (\text{fail}^\tau \mid v) \longrightarrow (\text{fail}^\ell \text{ says } \tau \mid (\bar{\eta}_\ell v)) \\
\text{[B-FAIL]} \quad \eta_\ell (\text{fail}^\tau \mid \text{fail}^\tau) \longrightarrow \text{fail}^\tau \\
\text{[B-RUNLEFT]} \quad \langle (E[\text{run}^\tau e_1 @ c'] \mid e_2), c \rangle \& s \Longrightarrow \langle (\text{ret } e_1 @ c \mid \bullet), c' \rangle \& \langle (E[\text{expect}^\tau] \mid e_2), c \rangle :: s \\
\text{[B-RETRIGHT]} \quad \frac{f' = \begin{cases} (\bar{\eta}_\ell v) & \text{if } f = v \\ \text{fail}^\ell \text{ says } \tau & \text{if } f = \text{fail}^\tau \end{cases}}{\langle (\bullet \mid \text{ret } f @ c), c' \rangle \& \langle (e_2 \mid E[\text{expect}^\ell \text{ says } \tau]), c \rangle :: s \Longrightarrow \langle (e_2 \mid E[f']), c \rangle \& s} \\
\text{[B-RETV]} \quad \frac{f'_i = \begin{cases} (\bar{\eta}_\ell v) & \text{if } f_i = v \\ \text{fail}^\ell \text{ says } \tau & \text{if } f_i = \text{fail}^\tau \end{cases}}{\langle \text{ret } (f_1 \mid f_2) @ c, c' \rangle \& \langle E[\text{expect}^\ell \text{ says } \tau], c \rangle :: s \Longrightarrow \langle E[(f'_1 \mid f'_2)], c \rangle \& s}
\end{array}$$

Fig. 25: Selected bracketed Evaluation Rules.

$$\begin{array}{l}
\mathcal{O}(\text{fail}^\tau, \Pi, p, \pi) = \circ \\
\mathcal{O}(\text{select } e_1 \text{ or } e_2, \Pi, p, \pi) = \text{select } \mathcal{O}(e_1, \Pi, p, \pi) \text{ or } \mathcal{O}(e_2, \Pi, p, \pi) \\
\mathcal{O}(\text{compare } e_1 \text{ and } e_2, \Pi, p, \pi) = \text{compare } \mathcal{O}(e_1, \Pi, p, \pi) \text{ and } \mathcal{O}(e_2, \Pi, p, \pi) \\
\mathcal{O}(\langle e, c \rangle \& s, \Pi, \ell, \pi) = \begin{cases} \mathcal{O}(e, \Pi, \ell, \pi) & s = \text{empty} \\ \mathcal{O}(e, \Pi, \ell, \pi) \& \mathcal{O}(s, \Pi, \ell, \pi) \end{cases} \\
\mathcal{O}(\langle e, c \rangle :: s, \Pi, \ell, \pi) = \begin{cases} \mathcal{O}(e, \Pi, \ell, \pi) & s = \text{empty} \\ \mathcal{O}(e, \Pi, \ell, \pi) :: \mathcal{O}(s, \Pi, \ell, \pi) \end{cases} \\
\mathcal{O}(E[\text{run}^\tau e @ c], \Pi, \ell, \pi) = \mathcal{O}(E[e], \Pi, \ell, \pi) \\
\mathcal{O}(\text{ret } e @ c, \Pi, \ell, \pi) = \mathcal{O}(e, \Pi, \ell, \pi)
\end{array}$$

Fig. 26: Observation function for intermediate FLAQR terms (extended from FLAC [7]).