Nomination of the paper:

Applying consensus and replication securely with FLAQR

By Priyanka Mondal, Maximilian Algehed, and Owen Arden

The paper introduces FLAQR, a new calculus that aims to facilitate the design of distributed systems based on quorum replication protocols, allowing the composition of protocols whose participants are allowed to have arbitrary trust relationships. It specializes the slogan that 'well-typed programs do not go wrong' to this setting, where going wrong means unrecoverable errors that violate program specifications. To achieve these protections, it extends the Flow-Limited Authorization Model (FLAM). This model features an algebra of principals with lattice structure that represents the acts-for relation on authority. The extension includes a new kind of principal, an availability authority, as well as two new authority operators, partial conjunction and partial disjunction, for consensus and replication protocols. These help represent the trade-offs between integrity and availability.

Availability attackers are studied under an active static attacker model where malicious principals are fixed prior to program execution. The power of an availability attacker is defined in the context of quorum systems, and compared to integrity and confidentiality attackers. The authority of an attacker is reduced to the authority of a particular principal in the so-called toleration set, essentially collections of upper bounds on attacker authorities that can be tolerated by the system. Using this, quorum types w.r.t. a quorum system are defined. Type soundness is phrased in terms of blame constraints describing sets of principals that may cause run-time failures, and a more general liveness theorem is provided for the special case majority quorum protocols. Finally, the paper presents noninterference theorems for confidentiality, integrity and availability.

FLAQR is motivated by application-agnostic consensus protocols such as Paxos, but focuses on higher level goals --availability, integrity, confidentiality-- than have been addressed in prior work on protocol verification. In fact, the FLAQR framework is not intended for verifying protocol implementations but rather it provides security abstractions to design and compose components with application-specific availability and integrity guarantees.

The technical novelties include a 'blame semantics' which associates a failure with a set of principals that may have caused it. The results include a blame soundness theorem that confirms this intuition. The other theorems are consensus liveness, confidentiality-integrity noninterference, and availability noninterference. Reviewers found the technical development to be correct and elegant, and the exposition to be clear and to the point.

In light of the significance of the results and the quality of presentation, the work was awarded Distinguished Paper at IEEE CSF 2022, and it is likely to have lasting impact.