The Next 700 BFT Protocols
(Invited Talk)

Rachid Guerraoui

EPFL LPD, Bat INR 310, Station 14, 1015 Lausanne, Switzerland

Byzantine fault-tolerant state machine replication (BFT) has reached a reasonable level of maturity as an appealing, software-based technique, to building robust distributed services with commodity hardware. The current tendency however is to implement a new BFT protocol from scratch for each new application and network environment. This is notoriously difficult. Modern BFT protocols require each more than 20.000 lines of sophisticated C code and proving their correctness involves an entire PhD. Maintainning and testing each new protocol seems just impossible.

This talk will present a candidate abstraction, named ABSTRACT (Abortable State Machine Replication), to remedy this situation. A BFT protocol is viewed as a, possibly dynamic, composition of instances of ABSTRACT, each instance developed and analyzed independently. A new effective BFT protocol can be developped by adding less than 10% of code to an existing one. Correctness proofs become at human reach and even model checking techniques can be envisaged. To illustrate the ABSTRACT approach, we describe a new BFT protocol we name Aliph: the first of a hopefully long series of effective yet modular BFT protocols. The Aliph protocol has a peak throughput that outperforms those of all BFT protocols we know of by 300% and a best case latency that is less than 30% of that of state of the art BFT protocols.

This is joint work with Dr V. Quema (CNRS) and Dr M. Vukolic (IBM).

T.P. Baker, A. Bui, and S. Tixeuil (Eds.): OPODIS 2008, LNCS 5401, p. 1, 2008.