Abstract

Proofs of Writing for Efficient Robust Storage

Dan Dobre

In this talk I will present PoWerStore, an efficient and robust distributed storage protocol recently developed at NEC Labs Europe.

PoWerStore's robustness comprises tolerating asynchrony, maximum number of Byzantine storage servers, any number of Byzantine readers and crash-faulty writers, and guaranteeing unconditional progress (i.e. wait-freedom) and strong consistency (i.e. linearizability) of read/write operations. PoWerStore's efficiency stems from combining lightweight cryptography, erasure coding and metadata write-backs, where readers write-back only metadata to achieve linearizability.

Central to PoWerStore are "Proofs of Writing" (PoW), a technique inspired from commitment schemes. PoW rely on a 2-round write procedure, in which the first round writes the actual data and the second round only serves at "proving" the occurrence of the first round. PoW enable efficient implementations of linearizable storage through metadata write-backs and low latency reads (comparable to reading from crash tolerant storage). We implemented PoWerStore and showed its superior performance when compared to existing robust storage protocols, including crash tolerant ones.