Bitcoin 2.0 for Replicated Data

Pascal Urso
pascal.urso@loria.fr

17 février 2017

Nombre d'étudiants souhaités : 1-3

Description du sujet

Context

Modern web-scale applications store and manipulate a massive amount of data. To tolerate failures and to lower user's perceived latency, this data is replicated on different data-centers distributed across the planet, often "cloud-hosted". The CAP theorem implies that to allow high availability, these data is replicated "optimistically", i.e. any replica can be modified independently. The challenge is then to merge concurrent modifications and to ensure consistency of the replicas.

Objective

Your objective will be to combine two solutions for such an distributed data consistency issue. On one hand, the Operational Transformations is a framework for managing replicated data. It has been applied successfully in different application contexts, but it is particularly well suited when the network overlay can provide a total order on communication. For instance, Google uses it to empower Googles Docs, but requires a heavy infrastructure based on GPS and atomic clocks to synchronizes communication.

On the other hand, the BitCoin protocol ensures such an ordered communication - the blockchain - in a pure peer-to-peer context. Several propositions where made to use such a protocol to build a more general distributed computing platform - aka BitCoin 2.0. But such projects have to face several non-trivial problems such a deferred broadcasting or undo-redo management, that are already solved by Operational Transformations.

Approach (techniques and tools to use)

Several open-sourced implementations exist for blockchain protocols and operation transformations.

Schedule and sharing out of the work

The main task will be to build, deploy and evaluate a "naive" implementation of the solution using a basic version of the blockchain protocol. Future improvements that could be proposed

by the students are related to the inefficiency of the standard blockchain protocol in such a cloud computing context.

The schedule will be divided in a state of the art study (until end of may), and then a agile development of the prototype until the end of the project.

References

— Ahmed-Nacer M., Urso P., Balegas V., Preguiça N. Merging OT and CRDT Algorithms. Principles and Practice of Eventual Consistency (PaPEC), Apr 2014, Amsterdam, Netherlands. 2014

— Corbett, J. C., Dean, J., Epstein, M., Fikes, A., Frost, C., Furman, J. J., ... & Woodford, D. Spanner : Google's globally distributed database. ACM Transactions on Computer Systems (TOCS), 31(3), 8. 2013.

— Vidot, N., Cart, M., Ferrié, J., & Suleiman, M. Copies convergence in a distributed real-time collaborative environment. In Proceedings of the 2000 ACM conference on Computer supported cooperative work (pp. 171-180). ACM. 2000.

— Vogels, W. Eventually consistent. Communications of the ACM, 52(1), 40-44. 2009.

Lieu

Sessions de travail régulières au laboratoire I3S (Sophia Antipolis)