Introduction

Today, the fault-tolerance techniques applied in practice are almost solely dedicated to handling crash-stop failures, for example, by employing replication, that is, by using multiple redundant instances of a single service. Apart from that, only specific techniques are used to selectively address the most common or most severe non-crash faults, for example, by using checksums to detect bit flips. In consequence, a wide spectrum of threats remains largely unaddressed, including software bugs, spurious hardware errors, and intrusions. Handling such arbitrary faults in a generic fashion requires Byzantine fault tolerance (BFT).

To date, Byzantine fault-tolerant systems have mainly been of academic interest. Despite all progress made in this field, industry is reluctant to actually utilize the research results available. One reason for this is the high resource demand, current BFT systems exhibit.

To tackle this problem, we present *CheapBFT*, a resource-efficient system for providing Byzantine fault-tolerant services. CheapBFT employs a so-called hybrid fault model in which a trusted subsystem is supposed to be subject to crash faults only whereas the rest of the system can behave in a Byzantine manner, that is, arbitrarily. For that purpose, CheapBFT relies on a novel FPGA-based hardware module, named *CASH*, that provides high performance but comprises only a small trusted computing base, which denotes the size of the system part that is trusted to not behave arbitrarily.

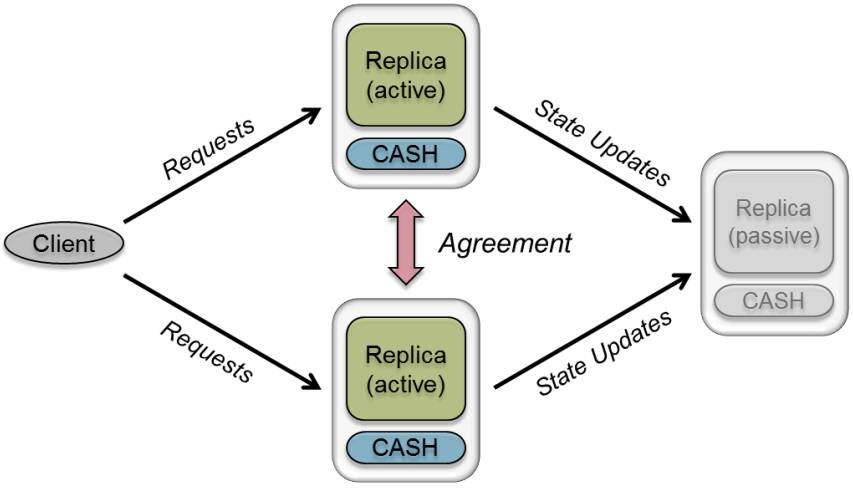

Figure 1: Overview of CheapBFT

From 3f + 1 Active Replicas to f + 1

Traditional BFT systems require 3f + 1 replicas which actively execute an instance of the service each to tolerate f arbitrary faults. By utilizing a hybrid fault model with a subsystem that cannot be compromised by an attack, the number of replicas can be lowered to 2f + 1.

Although they reduce the provisioning costs for BFT, such state-of-the-art systems have a major disadvantage: they either require a software-based trusted subsystem with a large trusted computing base or a hardware-based variant that imposes significant performance penalties.

Therefore, we developed CASH, an FPGA-based trusted hardware subsystem (see Figure 2) which achieves both the goal of a small trusted computing base and high performance rates.

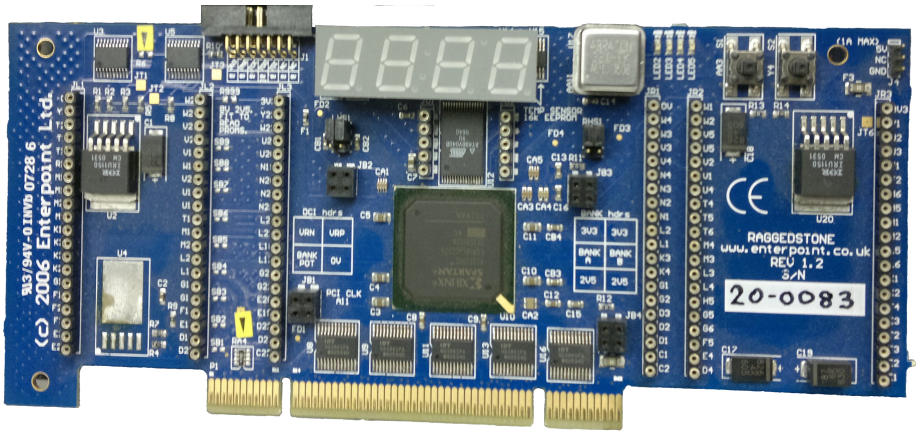

Figure 2: FPGA board used to implement the trusted subsystem CASH

CHEAPBFT: AN EFFICIENT BFT SYSTEM

Agreement: CheapTiny

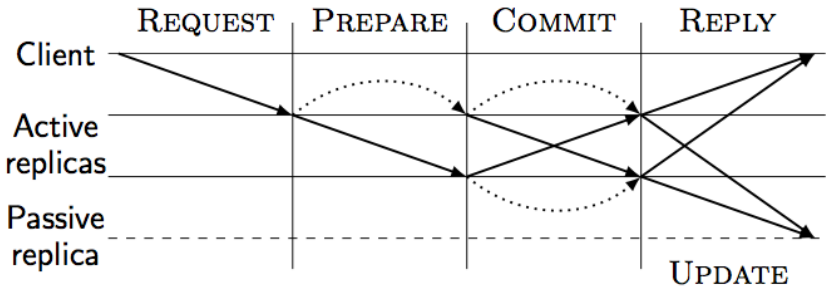

Agreement: MinBFT

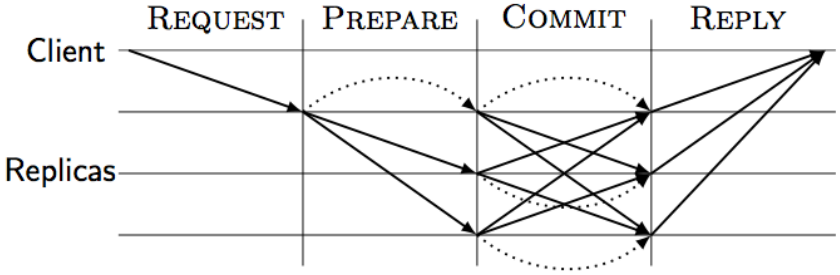

Figure 3: Message exchange by the agreement sub-protocols

In addition, CheapBFT advances the state of the art in resource-efficient BFT systems by running a composite agreement protocol, which requires only f + 1 actively participating replicas for agreeing on requests during normal-case operation. The agreement protocol of CheapBFT consists of three sub-protocols: the normal-case protocol CheapTiny, the transition protocol CheapChange, and the fall-back protocol MinBFT [1]. During normal-case operation, CheapTiny makes use of passive replication to save resources; it is the first BFT agreement protocol that requires only f + 1 active replicas. However, CheapTiny is not able to tolerate faults, so that in case of suspected or detected faulty behavior of replicas, CheapBFT runs CheapChange to bring all non-faulty replicas into a consistent state. Having completed CheapChange, the replicas temporarily execute the MinBFT protocol, which involves 2f + 1 active replicas (i.e., it can tolerate up to f faults), before eventually switching back to CheapTiny. To give an impression about the two agreement sub-protocols and their relative complexity, Figure 3 depicts simplified views of the message exchanges that take place when a client request is to be processed.

References

[1] Veronese, G. S.; Correia, M.; Bessani, A. N.; Lung, L. C. & Verissimo, P. Efficient Byzantine Fault-Tolerance. IEEE Transactions on Computers, Jan. 2013.

[2] Kapitza, R.; Behl, J.; Cachin, C.; Distler, T.; Kuhnle, S.; Mohammadi, S. V.; Schröder-Preikschat, W. & Stengel, K. CheapBFT: Resource-efficient Byzantine Fault Tolerance. Proc. of the 7th ACM European Conference on Computer Systems (EuroSys '12), ACM, 2012.

Further Information

Further information about CheapBFT can be found in [2] and under Deliverable „D2.1.2—Preliminary Description of Mechanisms and Components for Single Trusted Clouds".

Disclaimer

The TClouds project has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement number ICT-257243.

TClouds at a glance

Project number:
257243

TClouds mission:

- Develop an advanced cloud infrastructure that delivers computing and storage with a new level of security, privacy, and resilience.
- Change the perceptions of cloud computing by demonstrating the prototype infrastructure in socially significant application areas.

Project start:
01.10.2010

Project duration:
3 years

Total costs:
EUR 10.536.129

EC contribution:
EUR 7.500.000

Consortium:
14 partners from 7 different countries.

Project Coordinator:
Dr. Klaus-Michael Koch
coordination@tclouds-project.eu

Technical Leader:
Dr. Christian Cachin
cca@zurich.ibm.com

Project website:
www.tclouds-project.eu