SUMMARY:
Haifeng Yu and Amin Vahdat.
Design and Evaluation of a Continuous Consistency Model for Replicated Services.
In Proc. of the Symp. on Operating System Design and Implementation (OSDI'00), pages 21-21, 2000.

DATE: 25 January 2010

The paper studies the tradeoff between consistency and availability/performance in replicated systems. Supporting strong consistency is desirable, but it leads to poor availability/performance. Therefore some systems only provide optimistic consistency to improve performance. Unfortunately, there is no guarantee on the consistency level of these systems. The objective of this paper is to design a model that provides consistency levels between strong and optimistic consistency models. The motivation is that in various applications we can tolerate some inconsistency to achieve better availability/performance. However, at the same time we want to have some guarantee/bound on the amount of inconsistency in the system. The authors propose TACT, a system that provides a continuous consistency model in which applications can dynamically specify their own consistency levels and the system efficiently satisfies the specified consistency requirements.

Three metrics are used to specify the consistency requirements of applications, namely *Numerical Error*, *Order Error*, and *Staleness*. Informally, numerical error of a replica reflects the updates unseen by that replica, order error reflects the difference between order of updates at the current replica from their real order, and staleness reflects the delay of update propagation. The paper argues that this specification framework is expressive enough to specify consistency requirements for a broad range of applications. To support this claim, they consider three representative applications and explain how different consistency requirements can be specified for these applications. The corresponding metrics are defined on *conits*: a conit is a physical or logical unit of consistency. A good feature of the model is that different replicas can have different consistency levels. This enables the system to assign each request to a replica with appropriate consistency bounds. The paper also describes some algorithms for enforcing the specified consistency requirements. It uses *pushes* and *pulls* to make sure the required metric bounds are not violated.

The paper also provides an experimental evaluation of the proposed model for the three applications. The results reflect the expectation that the performance improves as consistency is relaxed. The authors compare their framework with other relevant systems and argue that their model is more expressive than most alternatives.

SUMMARIZED BY: Somayyeh Zangooei