

Used two sources for more background:

- <http://muratbuffalo.blogspot.com/2011/02/chain-replication-for-supporting-high.html> Murat Demirbas is a principal applied scientist at AWS on leave as a computer science and engineering professor at SUNY Buffalo
- <http://dsrg.pdos.csail.mit.edu/2013/08/08/chain-replication/> Distributed Systems Reading Group at MIT CSAIL

1 Summary

At the time this paper was published, *chain replication* was a novel approach to coordinating fail-stop servers that simultaneously supported high throughput, availability, and strong consistency. §2 is a discussion on the generic storage service interface of the time. §3 introduces the novel query and update operations using chain replication, and §4 compares chain replication with the primary/backup approach. Finally, §5 offers analysis w.r.t throughput and availability using experiments related to chain replication on a simulated network, §6 compares it with other scalable, storage systems, and the paper concludes in §7.

2 Strengths of the paper

Obviously, the primary strength of this paper is novelty related to the introduction of a new replication protocol, *chain replication* that, unlike the classic primary/backup approach, achieves high throughput and availability while providing strong consistency guarantees. I appreciate the journey the authors take to explain their approach without overwhelming the reader at any given point. For example, figure 2 is sufficient to gain the overview of the chain structure and helps in visualizing how read and write operations are implemented within. Also, rather than cherry-picking scenarios to showcase solely positive results of using chain replication, the authors point out relatively probably scenarios wherein, due to a failed server, performance may be severely impacted.

3 Major weakness of the paper

It is hard to imagine the approach scale past the homogeneous LAN clusters setting which the prototype was primarily intended for. While it is acknowledged that in order to have this be deployed in a heterogeneous wide-area setting the uniform random placement of volume replicas would be infeasible, I wish we got to see some work done on a prototype with replica placement curated in accordance with access patterns, network proximity, and observed host reliability. Essentially, employing the necessary and various techniques to control the load imbalances plaguing today's highly scalable and high throughput systems would make it even more informative, especially for the industry audience.

4 Future work opportunities

Given the primary intent of this paper was to introduce the novel *chain replication* method, I think Professor van Renesse and Professor Schneider did an excellent job in doing so. Further work could potentially expand upon this by experimenting with techniques to balance loads such that even if only a handful of items in a large database were the only ones requested for look-ups, the approach would still work well.