

Extreme Web Caching for Faster Web Browsing

Ali Raza^{*}, Yasir Zaki^{*}, Thomas Pötsch⁺, Jay Chen^{*}, Lakshmi Subramanian[#]

^{*}New York University Abu Dhabi, ⁺University of Bremen, [#]New York University

{ali.raza|yasir.zaki|jchen}@nyu.edu, thp@comnets.uni-bremen.de, lakshmi@nyu.edu

Problem

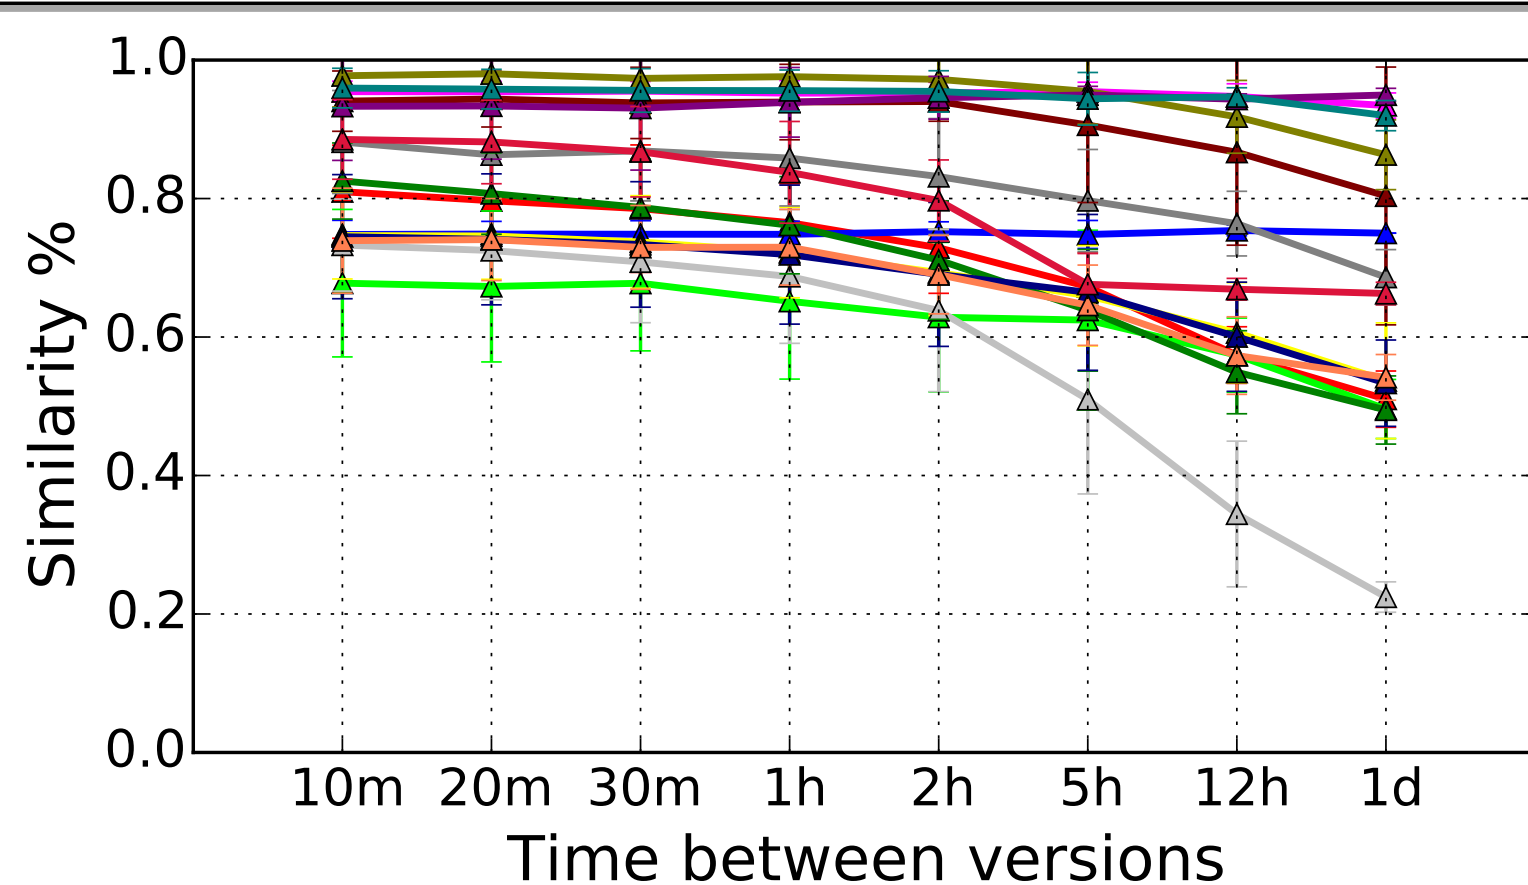

- Modern web pages are very complex
- Severe impact in developing regions [1]
 - Scarce bandwidth
 - High latencies

Idea


- Bringing content closer to client/user
- Better utilization of caching related Headers
- Predicting change rate of objects

Motivation

- 70% of objects do not change even after one day
- Changes in text-based objects are not significant [2]
- Publishers violate cache-control



Proposed Architecture



- Edge Cache
 - Close to the client
 - Regular caching proxy
 - Collect statistics of client requests
 - Update Cloud Controller with statistics
- Cloud Controller
 - Control over multiple Edge Caches
 - Periodically request objects of set of web pages
 - Hold record of versions
 - Predict change rate of objects
 - Estimate and modify Max-age option
 - Push cached objects to Edge Cache

Benefits

- Extreme Cache would lead to:
 - Better cache utilization
 - Faster page load times

References

- [1] Y. Zaki, J. Chen, T. Pötsch, T. Ahmad, and L. Subramanian, Dissecting web latency in Ghana, In Proceedings of the 2014 ACM Conference on Internet Measurement Conference, IMC'14, Vancouver, Canada, 2014.
- [2] X. S. Wang, A. Krishnamurthy, and D. Wetherall, How much can we micro-cache web pages?, In Proceedings of the 2014 ACM Conference on Internet Measurement Conference, IMC'14, Vancouver, Canada, 2014.