

Nathan Matsuda
CS443
1/15/07

1. Paper title and its author(s).

Corona: A High Performance Publish-Subscribe System for the World Wide Web
Venugopalan Ramasubramanian, Ryan Peterson, and Emin G'un Sirer

2. Brief one-line summary.

The paper presents and evaluates an overlay-based publish-subscribe system which optimizes the tradeoff between update latency and bandwidth usage/server load.

3. A paragraph of the most important ideas: perhaps a combination of their motivations, observations, interesting parts of the design, or clever parts of their implementation.

Earlier publish-subscribe systems like RSS use clients that simply poll each subscribed source for updates at a fixed rate. This introduces high lag for users who set the polling frequency very low and introduces huge load on servers when users poll often to see updates faster. The corona implements a structured network of peer to peer nodes overlaid on standard instant messaging protocols.

The nodes then have a mathematical model of the optimization function for the specific structure of the network overlay. Each channel (news source, website, etc) has it's own managing node in the network, assigned based on proximity to the source (with regards to the overlay structure), so that the system is completely decentralized. No central server is needed, making the system infinitely scaleable, in theory. Instead of each subscriber polling the server independently, they share update information amongst each other.

The paper presents several different versions of the optimization function, each with it's own specific attributes. These include the average load on servers, the average latency time, the latency time with regard to update of channel, the popularity of a channel and so on. Several of these functions were tested on a real network with subscriptions to real RSS feeds and compared to the performance by standard RSS clients. The system showed update latency improvements up to an order of magnitude over the RSS system while maintaining the same load on servers.

4. A paragraph of the largest flaws; maybe an experiment was poorly designed or the main idea had a narrow scope or applicability. Being able to assess weaknesses as well as strengths is an important skill for this course and beyond.

One major limitation of these test setup was that it used IM for inter-node communication. The instant messaging service used sends all messages through a centralized server, nullifying the decentralized goals of the system. The authors list many, many examples for other publish-subscribe systems, network overlays, and micronews services, which requires extensive outside research to make sense of.

5. A last paragraph where you state the relevance of the ideas today, potential future research suggested by the article, etc.

Subscription services are widespread in use so improvements in both user

experienced lag and bandwidth usage have benefits on both ends of the publish-subscribe relationship. There is a lot of related work mentioned in the paper regarding overlay network structures of various types, so this paper contributes to that discussion by providing a simple and clearly useful real world application for such network overlays.