

A Locality-Aware Approach to Distributed Systems

Abstract

Today's high-bandwidth and real time applications place stringent, new demands on the Internet. For example, voice-over-Internet-Protocol, video-on-demand, content distribution, and real-time multi-player gaming require real-time communication and dissemination across a potentially wide area. In all these applications, placing services, creating groups, or selecting servers in a fashion that takes network latency into account can dramatically improve performance.

Network coordinates are a promising technique for providing locality-awareness for these applications. They produce scalable latency estimates with minimal overhead. Previous work has shown the feasibility of network coordinates, but only in limited contexts. In this thesis, I measure the performance of the largest existing network coordinate system, improve its accuracy and stability through several key techniques, develop a locality-aware routing substrate, and build locality-aware applications with network coordinates.

I construct accurate coordinate systems in live networks. I introduce three techniques to improve the accuracy and stability of live coordinate systems and study their performance within Azureus, a popular BitTorrent client with more than a million nodes. Released as the open source Pyxida library, these techniques minimize overhead, adapt to latency anomalies, and increase coordinate stability, improving Azureus's accuracy by 43% and its stability by four orders-of-magnitude. Studying this system has also generated long-term traces for other researchers to use.

I also examine locality-aware routing and resource selection. I develop a practical routing algorithm using network coordinates, creating a building block for higher-level abstractions such as multicast and remote service discovery. I measure two applications that use network coordinates to optimize resource selection decisions: overlay routing, where delay was cut by 33%, and swarm-based file exchange, where network usage is reduced by 12% and download times are improved by 11% in my experiments.

Network coordinates are not always the best tool to provide locality-awareness. I conclude with an examination of the difficulties in directly embedding network characteristics other than latency — in particular, bandwidth — and qualify what contexts are appropriate for achieving locality-awareness with network coordinates.