

Araneola: A Scalable Reliable Multicast System for Dynamic Environments

Roie Melamed

CS Department, Technion
mroi@cs.technion.ac.il

Idit Keidar

EE Department, Technion
idish@ee.technion.ac.il

Abstract

We present Araneola*, a scalable reliable application-level multicast system for highly dynamic wide-area environments. Araneola supports multi-point to multi-point reliable communication in a fully distributed manner while incurring constant load on each node. For a tunable parameter $k \geq 3$, Araneola constructs and dynamically maintains an overlay structure in which each node's degree is either k or $k + 1$, and roughly 90% of the nodes have degree k . Empirical evaluation shows that Araneola's overlay structure achieves three important mathematical properties of k -regular random graphs (i.e., random graphs in which each node has exactly k neighbors) with N nodes: (i) its diameter grows logarithmically with N ; (ii) it is generally k -connected; and (iii) it remains highly connected following random removal of linear-size subsets of edges or nodes. The overlay is constructed at a very low cost: each join, leave, or failure is handled locally, and entails the sending of only about $3k$ messages in total.

Given this overlay, Araneola disseminates multicast messages by gossiping over the overlay's links. We show that compared to a standard gossip-based multicast protocol, Araneola achieves substantial improvements in load, reliability, and latency. Finally, we present an extension to Araneola in which the basic overlay is enhanced with additional links chosen according to geographic proximity and available bandwidth. We show that this approach reduces the number of physical hops messages traverse without hurting the overlay's robustness.

* Araneola means "little spider" in Latin.