Probabilistically-Structured Overlay Networks*

José Orlando Pereira
Computer Science and Technology Center (CCTC) / U. Minho
Campus de Gualtar, 4710-057 Braga PORTUGAL
email: jop@di.uminho.pt, phone: +351 253604477, fax: +351 253604471

MAP-i Thesis Proposal

1 Context

Application level multicast protocols are designed to support the efficient dissemination of information in the Internet through the coordination of participants without explicit support from the network infrastructure. This is a topic that has deserved significant attention in recent years since some applications require additional guarantees, such as bimodal reliability, and IP multicast-based protocols are hard to deploy pervasively.

The dependency on IP multicast is avoided by creating overlay networks to efficiently route and disseminate data. Two alternative approaches have been proposed. The first consists in building a structured overlay that explicitly attempts to optimize according to some efficiency criteria, for instance, to build a minimum cost spanning tree. The structured overlay is then used to deterministically perform the routing of multicast messages. The second approach consists in structuring the overlay as a simple random graph and then use an epidemic algorithm to disseminate data.

Although the deterministic solution seems preferable (in fact, it has been used in a number of protocols) it has been shown that structured overlay networks are brittle in the face of faults, load, and churn (i.e. frequent joins and leaves). In contrast, the simplicity and the redundancy provided by the gossiping approach result in high resilience and scalability.

The P-SON project addresses the problem of combining both approaches. The goal is to approximate the efficiency of structured overlay networks while leveraging the resilience of gossiping to cope with unstable periods or components. The challenge is to efficiently combine these approaches. While it would be possible to operate the two types of overlays in parallel, this would be extremely inefficient: in fact, it is likely that resources would be consumed both in the maintenance of the structured overlay as well as in the redundant retransmissions within the gossip protocol.

Instead, the project aims at lightweight protocols that can impose the right amount of structure to an otherwise random overlay. The idea is to start by gossiping preferably to a sub-set of the random overlay that tentatively approximates the desired structure. We aim at protocols where structure emerges probabilistically only as much as can be obtained with very low resource usage, for instance, by locally diagnosing the system and judiciously adjusting membership and gossip parameters. By varying the adaptation policies, the overlay can be optimized according to different reliability and performance criteria.

More information about the project is available at *http://pson.lsd.di.uminho.pt* and *http://neem.sf.net*.

*In collaboration with the P-SON research project (POSC/EIA/60941/2004).

2 Objectives

The goal of this project is to advance the state of the art in epidemic multicast protocols by:

- Proposing and evaluating adaptation mechanisms and policies for overlay maintenance and message dissemination.
- Proposing a service decomposition of epidemic multicast protocols that supports heterogeneous deployments in the Internet.