

A dynamically adaptive, unstructured multicast overlay

Matei Ripeanu¹

(work in cooperation with Ian Foster¹, Adriana Iamnitchi², Anne Rogers¹)

¹ Computer Science Department, The University of Chicago,
1100 E. 58th Street, Chicago, IL 60637, US

{matei, foster, amr}@cs.uchicago.edu

² Computer Science Department, Duke University,
D306 Levine Science Research Center Durham, NC 27708-0129, US
anda@cs.uchicago.edu

Summary

Multicast is a broadly useful communications primitive in distributed systems which provides the ability to deliver data to every member of a set of nodes.

Recently, considerable effort has been spent designing multicast solutions based on structured overlays. Yet unstructured solutions may be less expensive to create and maintain, and map more naturally on inherently heterogeneous sets of end nodes linked by networks that display scale-free and small-world characteristics.

We argue that it is possible to design an unstructured multi-source multicast overlay that is less complex than, but as efficient as, current state-of-the-art solutions, based on structured overlays or on full routing protocols at the overlay level. In brief, our solution works as follows: it builds a relatively rich base overlay and uses heuristics to adapt its topology to the state of the underlying network. On top of the base overlay, UMM selects efficient distribution trees for each source by first flooding the base overlay and then, using the implicit information contained in duplicate messages received at nodes, by selecting and filtering out unnecessary overlay tunnels.

Our solution is efficient, simple, and adaptive:

- *Efficient.* We use traditional measures of stretch and stress to compare with IP layer multicast and, indirectly, with other proposed solutions. Experience gained with a deployment on over 60 PlanetLab nodes and on an emulated distributed testbed indicates that UMM performs well compared to alternative solutions.
- *Simple.* When two mechanisms offer similar efficiency for comparable costs, the discriminating factor is complexity. UMM preserves the simplicity of flooding based solutions and the flexibility of unstructured overlays. It decouples message routing and overlay construction so that the two can be optimized separately.
- *Adaptive.* UMM adapts to the underlying network topology, and recovers from node and network failures with minimal disruption to the multicast service offered.