Multicast Routing in Datagram Internetworks and Extended LANs

Reviewer: Sayandeep Sen

March 5, 2008

1 Summary

The goal of the paper is to extend the scope of (then) popular multicast LAN applications over internetworks of LANS. The authors stipulate the following as the desirable characteristics of such extension protocols: 1. Group Addressing, 2. High probability delivery, 3. Low delay. The authors then present a set of modifications to unicast routing algorithms (single-spanning tree, distance-vector, and link-state routing) to make them multicast capable over an internetwork. The proposed modifications can be summarized as follows:

a). For networks which have a spanning tree topology to begin with, the authors propose *Single Spanning Tree multicast* protocol. This protocol works by letting the interconnecting bridges learning the group memberships of their connected nodes and then broadcasting this information periodically to all other bridges as membership reports,

b) For networks which have a graph like connectivity (opposed to spanning tree connectivity) the authors provide a series of modifications for Distance Vector Routing algorithms. The authors start by modifying Reverse Path Flooding (RPF) which necessitates a router to forward broadcast packet (on all but the incoming link). This algorithm suffers from broadcast flooding (all routers forward all packets). To remedy this authors propose Reverse Path Broadcasting (RPB), in which a router forwards a packet from a source S if and only if it arrived via the shortest path back to S. The conditional flooding reduces the amount of broadcast flood and ensures (along with packet TTLs) that the flooding stops a some point. The then proposed Truncated Reverse Path Broadcasting (TRPB) to prune the broadcast trees to only include those non-leaf LANS which have membership of a given group. Finally, Reverse Path Multicasting (RPM) was proposed, which uses non-membership reports from routers to prune the shortest path broadcast trees on-demand. These non-membership reports are generated when a router receives a multicast packet for group G and it is a leaf node (doesn't forward traffic to other routers) and none of its child links contain members of the group.

c) The authors then show that the Link state routing algorithms can be easily enhanced to accommodate Multicast. The algorithm needs extra information regarding the group memberships of the child nodes of each router and then it can form paths(spanning trees) for each source to a set of destinations on demand.

The authors finally mention that the inherent hierarchical structure of the Internet would be beneficial for scaling of their proposed multicast routing algorithms.

A major critic of this paper is the absence of any validation of the proposed schemes (understandable given the nascent stage of the networking field itself). While the proposed algorithms are easy to understand (and most probably have correct) behavior, their actual implementation on routers and bridges would had be non trivial (lots of engineering issues). Also, them memory, bandwidth and time overheads of the proposed solutions would also be high, While the above mentioned issues were most probably not the focus of the work, the authors should had thoroughly evaluated their algorithms on the degree of achieving the desired behavior properties (low delay, high probability delivery etc.).