AG Nodes With Source Specific Multicast

Maintaining "Shared Spaces" With Finer-Grained Bandwidth Control

Jeff Eschbach, Motorola Labs
Access Grid Node Retreat, January 31, 2001

IGMPv3 (Internet Group Management Protocol Version 3) is a protocol recently proposed to enhance the operation of multicast. IGMPv3 employs SSM (Source Specific Multicast), which enables hosts to selectively receive traffic from individual sources of a multicast group. This new technology reduces the distribution of unnecessary multicast traffic, which serves to alleviate network congestion.

Currently a meeting involving several AG nodes can overwhelm even high-bandwidth links. Using IGMPv2, outgoing traffic from each AG node is distributed to all other hosts participating in the meeting. With only a few Mbps of traffic from each AG node, tens of participants could consume the bandwidth of a DS-3 (45 Mbps) or even an OC-3 (155Mbps) link. Non-AG node users at a location may also vie for the same bandwidth and exacerbate the problem. To counteract this, the use of SSM may be incorporated into future versions of AG nodes to provide finer-grained bandwidth control. By accepting audio/video streams only from selected sites, a node can avoid incoming traffic from locations of less interest (e.g.: those not being displayed locally), thereby reducing network congestion.

However, this model jeopardizes the notion of "shared spaces" for AG nodes. A driving vision for AG node research is to enable collaboration between geographically separate locations that is comparable to sharing the same physical space. Within a given virtual room, the space is "shared" by transmitting audio/video feeds between all participating nodes, and it is expected that each participant can communicate with all others in the virtual space. Yet a simple deployment of SSM capabilities in AG nodes would violate that principle. If a given node does not specifically select all transmitting sources, it could be unaware of some participants in the virtual room. In this case the virtual space is not truly "shared". Therefore a need exists to control bandwidth via SSM while not violating the shared spaces model for AG nodes.

A potential solution uses two multicast groups per audio/video transmission in a virtual room, one to enable shared spaces and another to provide specific feeds via SSM. The first multicast group distributes traffic between all nodes in the space, similar to the operation of today's system. However, nodes only supply low-bandwidth feeds to these groups, enough to provide a minimal shared spaces experience. Some level of audio would be distributed in this fashion, possibly coupled with low-resolution video to aid operators in distinguishing locations. In contrast, the second multicast group employs SSM technology to distribute individual high-bandwidth feeds, but only to requesting nodes. Such transmissions would include high-resolution video streams and perhaps additional traffic for enhanced audio quality. This overall method maintains shared spaces through the low-bandwidth transmissions, yet provides finer-grained control over high-bandwidth traffic with SSM.

A variety of specific implementations are possible. Audio could simply be distributed between all nodes with standard multicast, while all video feeds are distributed via SSM. Or lower-

quality (requiring less traffic) audio and/or video could use standard multicast groups, while corresponding higher-quality feeds are distributed through SSM. In this case, traffic for a lower-quality feed is dropped if the corresponding SSM feed is selected and used. Adaptive audio/video feeds with layering techniques could also be employed, where a base layer stream can be combined with an enhancement layer stream to provide a higher bit rate and thus obtain better quality. Base layer streams would reach all nodes through low-bandwidth multicast groups, thereby maintaining shared spaces with lower-quality transmissions. The enhancement layer streams would be available through SSM, and combined with the base layer if selected by an AG node. Certainly other implementations are conceivable as well, and future analysis is required to validate this overall method and to understand the tradeoffs between possible approaches.