A Classification of Reliable Multicast Protocols

J. William Atwood, Concordia University

Abstract

The range of user requirements on multicast protocols is so wide that no single protocol will ever satisfy them. The set of multicast protocols can be classified using the user requirements, and the architectures, mechanisms, communications patterns, and policies used to satisfy these requirements. We provide such a classification, and illustrate it with several example protocols chosen to cover the range of features described.

For point-to-point (unicast) communication, the Transmission Control Protocol (TCP) [1] has dominated the Internet data communications landscape for many years.

TCP provides absolutely reliable service unless the underlying network fails. It assumes the user requires that all data be delivered, in sequence.

To provide multicast services (point-to-multipoint and multipoint-to-multipoint), a variety of protocols is available. However, when considering reliable multicast equivalents to TCP, the situation is vastly different from the unicast situation, as the number of possible failure modes is larger, and the definition of "reliable" can take on many shades of meaning. Each protocol proposed in the literature is intended to meet the needs of a particular set of applications. Each has a slightly different definition of reliability, and each operates in a slightly or significantly different environment. Given the wide variety of requirements and environments, there will never be a "one size fits all" multicast protocol design [2]. However, this does not mean it is impossible to develop a family of protocols (or a single protocol with a variety of selectable features) that satisfies a wide range of requirements.

As a step toward identifying the features of such a family, we have developed a taxonomy of reliable multicast protocols. Previous taxonomies include Diot *et al.* [3], which is a survey of multicast protocols in terms of functionality and mechanisms for reliable multicast transmission, and Obraczka [4], which reviews several existing multicast transport mechanisms and classifies them according to their distinct features. The Reliable Multicast Transport Working Group (RMTWG) within the Internet Engineering Task force (IETF) has a mandate to standardize reliable multicast transport. It has developed a framework for standardization of reliable multicast transport for the specific subset of reliable multicast protocols that relate to bulk data transfer (point-to-multipoint), and is now populating that framework. The reader is referred to the RMTWG Website [5] for current information on the progress and status of the RFCs and Internet drafts.

To help identify the components that are common across all classes of reliable multicast protocols (point-to-multipont and multipoint-to-multipoint), we have developed a taxonomy based on requirements, architectures, mechanisms, communications patterns, and policies. This brings out the orthogonality between the various requirements of reliable multicast transport protocols and the ways they are used to satisfy those requirements.

Requirements

Traditional protocol design has started with a statement of the requirements, and followed with the exploration of possible designs that can meet these requirements. In this section we identify the range of requirements that characterize and differentiate reliable multicast protocols.

To make the issues more concrete, capsule summaries of the following multicast protocols are given later in the article: UDP, XTP (two variants), RMP, PGM, SRM, LGMP, RMTP (two variants), and LPC. (LPC is a representative hybrid forward error eorrecting protocol, which is included in the study to illustrate an upper bound on receiver scalability.) For each requirement stated below, we note the example protocols that have this requirement. Table 1 gives the correlation between user requirements and the example protocols.

Multicast applications vary in their requirements along a number of dimensions:

- Number of senders
- Group organization and receiver scalability
- Data reliability
- Congestion control
- Group management
- Ordering

These parameters all interact with each other: for example, the data reliability requirements that can be met are dependent on the receiver scalability required.

Number of Senders

The first dimension of our classification is based on the number of senders: single-sender (point-to-multipoint) and multiple-sender (multipoint-to-multipoint) applications.

Point-to-multipoint or 1-to-N multicast applications require data delivery from a single source to multiple receivers, and usually run without human interaction. A few example applications for this set include software distribution, data distribution and replication, and mailing list delivery.

Multipoint-to-multipoint or M-to-N applications require data delivery from multiple sources to multiple receivers.

0890-8044/04/$20.00 © 2004 IEEE