Chapter X
New Internet Protocols for Multimedia Transmission

Michael Welzl
University of Innsbruck, Austria

ABSTRACT

This chapter will introduce three new IETF transport layer protocols in support of multimedia data transmission and discuss their usage. First, the stream control transmission protocol (SCTP) will be described; this protocol was originally designed for telephony signaling across the Internet, but it is in fact broadly applicable. Second, UDP-Lite (an even simpler UDP) will be explained; this is an example of a small protocol change that opened a large can of worms. The chapter concludes with an overview of the datagram congestion control protocol (DCCP), a newly devised IETF protocol for the transmission of unreliable (typically real-time multimedia) data streams.

INTRODUCTION

For decades, two transport layer protocols of the TCP/IP suite were almost exclusively used: TCP and UDP. The services that these protocols provide are entirely different, and easy to grasp: while the latter simply makes the "best effort" service of the Internet accessible to applications, TCP reliably transfers a stream of bytes across the network. UDP only has port numbers that make it possible to distinguish between several communicating entities which share the same IP address and a checksum that ensures data integrity, but TCP encompasses a large number of additional functions:

- **Stream-based in-order delivery:** Packets are ordered according to sequence numbers, and only consecutive bytes are delivered
- **Reliability:** Missing packets are detected and retransmitted

- **Flow control:** The receiver is protected against overload with a sliding window scheme
- **Congestion control:** The network is protected against overload by appropriately limiting the window of the sender
- **Connection handling:** Since TCP is a connection oriented protocol, it must have the ability to explicitly set up and tear down connections
- **Full-duplex communication:** An acknowledgment (ACK) can also carry user data; this is usually referred to as "piggy-backing"

The importance of these mechanisms varies. A protocol could, for instance, easily do without the full-duplex communication capability; on the other hand, some form of end-to-end congestion control has been identified as an indispensable element of any protocol that is to be used on the Internet (Floyd & Fall, 1999). This does however not mean that there is only one way to carry out congestion control: TCP uses an "additive increase, multiplicative decrease" strategy which essentially probes for the available bandwidth by linearly increasing the rate until a limit is hit (causing a packet to be dropped or a congestion signaling bit to be set), whereupon the rate is reduced by half. There are proposals for congestion control that is fair towards TCP ("TCP-friendly") yet more suitable for multimedia applications because the rate fluctuations are less severe. One notable example is "TCP-friendly rate control (TFRC)" (Floyd, Handley, Padhye, & Widmer, 2000).

TCP does not provide the flexibility that today's applications need: it is neither possible to disable any of its aforementioned functions (in particular reliability, which adds delay but is typically not needed by real-time multimedia applications), nor can a user change the way they work (e.g., influence how congestion control is carried out). UDP, on the other hand, allows for more flexibility, but its feature set is so small that any additional protocol function must be implemented directly within the application that uses it. Sometimes, this is unacceptable — realizing TCP-friendly congestion control, for instance, is difficult, and may not be worth the effort from the perspective of a single application designer. Indeed, even the popular streaming media applications "RealPlayer" and "Windows Media Player" do not appear to properly adapt their rate in response to congestion (Hessler & Welzl, 2005).

In this chapter, we will take a look at three novel IETF protocols that change this situation somewhat: the "stream control transmission protocol (SCTP)," "UDP-Lite," and the "datagram congestion control protocol (DCCP)." While SCTP could also be regarded as some sort of a "TCP++," these three protocols share one notable property: they can emulate the behavior of TCP (or UDP, in the case of UDP-Lite), but with *less* features. The ability to effectively *disable* TCP features is therefore a feature in itself; this gives new meaning to the saying "*less is more*." Historically, SCTP is by far the oldest of these protocols; its main specification (Stewart et al., 2000) was published in 2000, and it is now going through the difficult post-standardization phase of achieving large-scale Internet deployment. Notably, the IETF recommends this protocol for authentication, authorization, and accounting (AAA) in any future IP service networks, and SCTP has been required by the 3rd Generation Partnership Project (3GPP) (Stewart & Xie, 2002, p. 17). UDP-Lite was recently published as a "Proposed Standard" — the same status as SCTP — by the IETF (Larzon, Degermark, Pink, Jonsson, & Fairhurst, 2004), and DCCP has not even reached this status yet; at the time of writing, its specification (Kohler, Handley, & Floyd, 2005) was still an Internet-draft, which is

Related Content