Preliminary Measurement of the RMTP/RTIP*

Hui Zhang and Tom Fisher
`hzhang, fisher@tenet.Berkeley.EDU`

The Tenet Group
Computer Science Division
Department of Electrical Engineering and Computer Sciences
University of California
Berkeley, California 94720
and
International Computer Science Institute
1947 Center Street
Berkeley, California 94704

Abstract. The Real-time Message Transport Protocol (RMTP) and the Real-Time Internetwork Protocol (RTIP) are the transport and network layer data delivery protocols in the Tenet Protocol Suite. We implemented the protocols in Ultrix on the DECstation 5000 workstations and in HP/UX on the HP 9000/7000 workstations. A preliminary measurement study has been conducted to evaluate the performance of the prototype implementation. Some of the results are: the throughput obtained by using RMTP/RTIP is comparable to that obtained by using raw IP; the rate control mechanism in RMTP/RTIP effectively enforces the traffic specification of communication clients; the scheduling mechanism in RTIP protects the real-time channel so that the performance of a real-time channel is not affected by the presence of IP traffic or other real-time channels in the network.

1 Introduction

There is an increasing demand to support real-time applications such as video conferencing, scientific visualization and medical imaging in an internetworking environment. These applications have stringent performance requirements in terms of delay, delay jitter, bandwidth and loss rate [Fer90]. The best-effort

---

* This research was supported by the National Science Foundation and the Defense Advanced Research Projects Agency (DARPA) under Cooperative Agreement NCR-8919038 with the Corporation for National Research Initiatives, by AT&T Bell Laboratories, Digital Equipment Corporation, Hitachi, Ltd., Hitachi America, Ltd., Pacific Bell, the University of California under a MICRO grant, and the International Computer Science Institute. The views and conclusions contained in this document are those of the authors, and should not be interpreted as representing official policies, either expressed or implied, of the U.S. Government or any of the sponsoring organizations.

service provided by the Internet Protocol [Pos81] is not adequate to support these applications.

The Tenet Group of University of California at Berkeley and International Computer Science Institute has proposed a new type of service called *guaranteed performance service*. By a guaranteed performance service, we mean that, before the communication starts, the client specifies its traffic characteristics and desired performance requirements; when the network accepts the client's request, the network guarantees that the specified performance requirements will be met provided that the client obeys the restrictions implied in its traffic description.

We believe that a reservation-oriented network architecture is needed to provide such a guaranteed performance service. The current Internet Protocol, which is connectionless and reservationless, cannot be used to support the service.

The Tenet Group has also designed and implemented a new protocol suite that provides guaranteed performance service [Fer92, FBZ92]. The protocol suite consists of five protocols: three data delivery protocols and two control protocols. The three data delivery protocols are: the Real-Time Internet Protocol (RTIP), the Real-time Message Transport Protocol (RMTP) [VZ91, ZVF92] and the Continuous Media Transport Protocol (CMTP) [WM91, MGWF92]. RTIP is the network layer protocol, while RMTP and CMTP are two transport layer protocols that provide message-oriented and stream-oriented transport services, respectively, on top of RTIP. The two control protocols are: the Real-Time Channel Administration Protocol (RCAP) [BM91], and the Real-Time Control Message Protocol (RTCMP). RCAP is responsible for establishment, tear-down and modification of the real-time channels, while RTCMP is responsible for control and management during data transfers.

In this paper, we will focus on the two data delivery protocols, RMTP and RTIP. Both protocols have been implemented on DECstation 5000 and HP9000 workstations. A preliminary measurement study has been conducted to evaluate the performance of the prototype implementation. Some of the results are: the throughput obtained by using RMTP/RTIP is comparable to that obtained by using raw IP; the rate control mechanism in RMTP/RTIP effectively enforces the traffic specification of communication clients; the scheduling mechanism in RTIP protects real-time connections so that the performance of a real-time connection is not affected by the presence of IP traffic or other real-time connections in the network.

The paper is organized as follows: Section 2 brief reviews the real-time channel scheme, on which the Tenet Protocol Suite are based; Section 3 describes the services, functions, software structure and programming interface for RTIP and RMTP; Section 4, which is the core of this paper, presents the simulation results; Section 5 briefly describes related work; Section 6 gives the conclusion and provides future work.

2 Background

The Tenet protocol suite implements the real-time channel scheme, which is a new communication abstraction proposed to support guaranteed performance service in a general packet-switched internetwork environment [FV90].

A real-time channel is a simplex unicast end-to-end connection with performance guarantees and traffic restrictions. Once established, it guarantees that the performance bounds requested by the communication client are satisfied so long as the client does not violate the traffic restrictions. The performance parameters a client can request are:

- delay bound D
- delay violation probability bound Z
- buffer overflow probability bound W
- delay jitter bound J

The traffic parameters a client needs to specify are:

- minimum packet inter-arrival time x_min
- average packet inter-arrival time x_ave
- averaging interval I
- maximum packet size s_max

A channel with both Z and W being 1 is called a *deterministic* channel; a channel with either Z or W being less than 1 is called a *statistical* channel [Fer90].

The following paradigm is proposed in [FV90] to provide guaranteed services to clients in a packet-switching network: before communication starts, the client specifies its traffic characteristics and performance requirements to the network; the client's traffic and performance parameters are translated into local parameters, and a set of connection admission control conditions are tested at each switch or gateway; the new channel is accepted only if its admission would not cause the performance guarantees made to other channels to be violated; during data transfers, each switch or gateway will service packets from different channels according to a service discipline; by ensuring that the local performance requirements are met at each switch or gateway, the end-to-end performance requirements can be satisfied.

Notice that there are two levels of control in this paradigm: at the connection level, the admission control policy allocates and reserves resources for each channel; at the packet level, the service discipline allocates resources according to the reservations made. In the Tenet protocol suite, RCAP is responsible