An Efficient Reservation Connection Control Protocol for Gigabit Networks

Emmanouel A. Varvarigos and Vishal Sharma[1]
Department of Elect. & Comp. Eng., University of California,
Santa Barbara, CA 93106-9560, USA

Abstract — **The *Efficient Reservation Virtual Circuit* (or ERVC) protocol is a novel connection control protocol designed for constant-rate delay-insensitive traffic in gigabit networks. In the ERVC protocol, session durations are recorded and capacity is reserved only for the duration of the session, starting at the time it is actually needed. The protocol also has the "reservation ahead" feature, which allows a node to calculate the time at which the requested capacity will be available and reserve it in advance, thus avoiding wasteful repetition of the call setup phase. In addition, the protocol is robust to link and node failures, and allows soft recovery from processor failures.**

I. Introduction

The ERVC protocol is one of the two candidate protocols that we are considering for implementation in the 40 Gbit/s ATM-based fiber-optic Thunder and Lightning network currently being developed at UCSB. In designing the connection and flow control algorithms for this network our objectives were to ensure lossless transmission, efficient utilization of capacity, minimum pre-transmission delay for delay-sensitive traffic, and packet arrival in correct order. To meet these objectives, we have proposed the ERVC protocol for constant-rate traffic, and the Ready-to-Go Virtual Circuit (or RGVC) protocol for best-effort traffic and traffic with little delay tolerance. The RGVC protocol, described in [1], uses back-pressure and requires buffering at intermediate nodes, whereas the ERVC protocol, described in [2], uses reservations and requires little buffering at intermediate nodes.

II. Why the ERVC protocol ?

In standard reservation schemes (abbreviated SRVC) the capacity required by a session at an intermediate node is reserved starting at the time the setup packet arrives at the node. This is inefficient since the capacity reserved will actually be used at least one round-trip delay after the arrival of the packet at the node. This is because the setup packet has to travel from the intermediate node to the destination, an acknowledgement has to be sent to the source, and the first data packet of the session has to travel to the intermediate node. Over long transmission distances, the round-trip propagation delay may be comparable to, or even larger than, the holding time of a session. In particular, if a typical session requests capacity r bits/sec, and transfers a total of M bits over a distance of L kilometers, then the maximum percentage of time that the capacity is efficiently used in a SRVC protocol is

$$e = \frac{\frac{M}{r}}{\frac{2Lc}{\eta} + \frac{M}{r}}, \qquad (1)$$

where $c/\eta = 5\,\mu$s/km is the propagation delay in the fiber. Typical values of these parameters for the Thunder and Lightning network are $r = 10$ Gbit/s, $M = 0.5$ Gbit, and $L = 3000$ km (coast-to-coast communication), which yields $e = 0.625$. In contrast, the efficiency factor e for the ERVC protocol can be as large as $e = 1$, independently of the parameters r, L, and M.

The "reservation ahead" feature of the ERVC protocol allows sessions to reserve capacity in advance for use at a later time. Thus, if capacity is available for a session starting at a time that is within the delay that the session can tolerate, the call is accepted on its first attempt. This feature, therefore, avoids unnecessarily prolonged call setup phases, reduces a session's susceptibility to blocking, and leads to efficient utilization of the available capacity.

III. Basic description of the protocol

In the ERVC protocol, each network node keeps track of the *utilization profile* of each outgoing link, which describes the residual capacity available on the link as a function of time. The utilizatio profile is stored as a linked-list of records, and is updated efficiently. Each intermediate node reserves the required capacity starting at the time at which this capacity will actually be used (which is at least one round-trip delay after the arrival of the setup packet at the node), and for time equal to the session duration. If the session duration is unknown, it is treated as infinite, and capacity is reserved for that session for an unspecified duration (as in standard reservation schemes). If the capacity is not available at the time requested, the setup packet may make a reservation starting at the first time the capacity becomes available, if the session can tolerate the delay. Since, capacity is blocked for other sessions only for the duration of the call and is available for the remaining time, this allows a considerably greater number of sessions to be served. It also avoids the wasteful repetition of the call setup process, because it enables a session to reserve the required capacity in its first attempt, possibly at a time later than the requested time. If adequate capacity is available at every intermediate node, the source eventually receives an acknowledgement from the destination and begins transmitting data. If the time at which adequate bandwidth first becomes available is exceeds the delay tolerance of the session, the call is blocked and is reattempted later, probably via a different path. The ERVC protocol requires a pre-transmission delay at least equal to the round-trip propagation delay between the source and the destination (as all reservation protocols do).

[1] Research supported by ARPA under Contract DABT63-93-C-0039

References

[1] E. A. Varvarigos and V. Sharma, "An efficient reservation connection control protocol for gigabit networks," submitted *IEEE/ACM Trans. on Networking*, January 1995.

[2] E. A. Varvarigos and V. Sharma, "A loss-free connection control protocol for the Thunder and Lightning network," submitted *Globecom'95*, March 1995.