A REAL-TIME TRANSPORT PROTOCOL*

Yechiam Yemini

CUCS-88-83

* This research has been supported in part by an NSF grant No. MCS-8110319 and by the Defense Advanced Research Project Agency of the Department of Defense.

ABSTRACT

A real-time protocol is usually concerned with the transportation of a real-time data stream over a packet switched network. Among the major issues distinguishing real-time protocols from ordinary transport protocols is the problem of trading delay for loss. That is, if some loss of packets may be acceptable, and usually inevitable, the objective of the protocol is to minimize the delay of packets, subject to constraints on the acceptable loss. This is unlike usual transport protocols, which are designed to guarantee no loss at the expense of increased delays.

This paper presents a model for the delay-loss tradeoffs in real-time transport protocols. It is demonstrated that, under very general assumptions, an optimal real-time protocol is a bang-bang protocol: there exists a threshold queue length such that as long as the packet queue length at the sender is less than the threshold the protocol should be an ordinary positive acknowledgement with retransmissions transport protocol. However, as soon as the threshold queue length is exceeded, a newly arriving packet causes the first packet in the queue to be discarded. Closed form expressions for the threshold buffer size are obtained and analyzed in terms of the given parameters of the system.

Table of Contents

1. INTRODUCTION

Problems of real-time communication over packet-switched computer communication networks arise in the context of packet voice communication [COHEN 78, BIALLY 80], video communication, and distributed sensors networks [DSN 78]. A real-time communication problem typically involves a source that needs to deliver information about an observed signal to a destination that needs this information to perform some task based upon the signal (e.g., estimate some of the signal parameters).

Among the chief concerns in the design of a real-time communication protocol is the problem of trading speed for reliability. This arises because insisting that any packet of the real-time stream be delivered to the receiver, and in the order of arrival, may result in costly delays and buffers overflowing. Moreover, in many real-time communication problems (e.g., voice or sensors communications) losing a few packets would be less harmful than delaying packets. Therefore, a more appropriate policy would allow some packets to be lost in order to decrease the expected delays.

A number of transmission schemes may be designed to facilitate the trading of reliability against speed. For instance, the receiver may discard packets arriving out of order, or the sender may discard packets when its queue grows. The objective of this paper is to study the fundamental tradeoff between reliability and time, for a class of real-time transport protocols.

2. THE PROTOCOL

Consider a transport station to which a stream of real-time packet traffic arrives. The sender end of the transport protocol transmits packets in a first come first serve manner. Packets may be lost in the medium (or equivalently, arrive out of order and be discarded at the receiver). Furthermore, the sender may discard packets as its queue grows, in order to to meet the real-time delay constraints. Packet loss results in a reduction in the quality of a system's response to the real-time task being processed; when packetized voice transmission is used the cost of packet loss is reduced voice quality, while in the case of a distributed sensors network it may be increased tracking errors.

The main question to be answered by this paper is what policy the sender should be using to determine which packets to discard, so as to meet externally specified real-time and quality of delivery constraints. We will first present a model for a real-time transport protocol in which the problems may be clearly posed, and then solve them to produce an optimal transport policy which, given some constraints on the expected rate at which packets may be discarded, minimizes the expected delay.

To fix our model let us assume that all packets are of a uniform size and that time is slotted to packet-size slots. Transmissions are restricted to occur only within respective time-slots*

Packets arrive at the sender from a Bernoulli source of rate λ and are stored in a buffer of an infinite length*. As soon as a packet reaches the head of the queue the sender transmits it, and then either retains a copy waiting for an acknowledgement, or discards the packet. Let p_n denote the probability that the packet is retained, where n is the number of packets in the queue, or discards it with probability $1-p_n$. A transmitted packet is delivered successfully with probability s, or lost with probability 1-s. A packet that is lost in the medium but has not been discarded by the sender returns to the head of the queue (as a result of a time-out condition or a negative acknowledgement) and gets retransmitted. This model is depicted in figure 2-1.

The assumption that the sender always discards the first packet in the queue is not essential; a simple variation would allow any packet to be dicarded. Also note that our model ignores a number of important parameters (for example, the propagation delay of acknowledgments). This reflects our decision to use a stripped-down model that will enable us to focus on the

---

*This discrete model of the transmission mechanism describes a Slotted-ALOHA broadcast network, a TDMA network, and some ring-network protocols. However, this assumption is not essential to our solution which, with some minor modifications, can also be applied to continuous-time transmission models.

*One of the surprising results of our analysis will eliminate this requirement later.

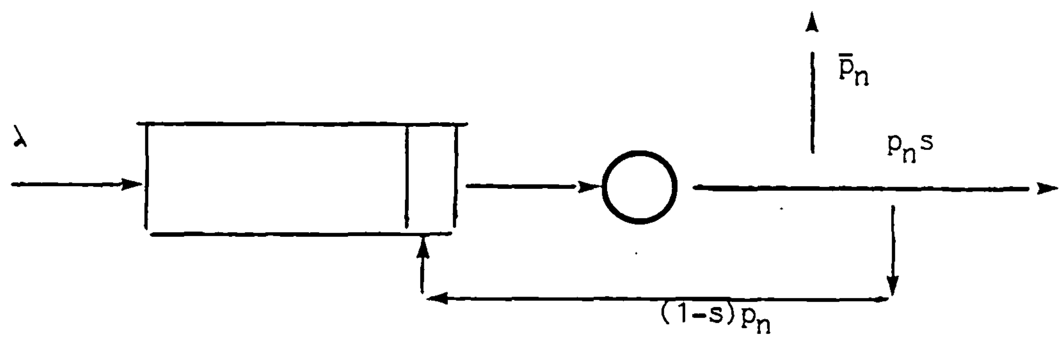

Figure 2-1: A queueing model for the sender protocol

significant tradeoffs of real-time communication, rather than a model where these tradeoffs are obscured by too many details.

We will focus on two performance measures: the expected delay of a packet and the expected loss. The problem to be solved can be stated as follows; given a limit to the expected loss, find a policy $\{p_n\}_0^\infty$ for the discarding of packets as a function of the queue length, which minimizes the expected delay. The major difficulties are caused by the dependence of the service process upon the queueing process, and the necessity for optimization over the set of possible probability sequences $\{p_n\}_0^\infty$.

3. OPTIMUM REAL-TIME PROTOCOLS

Let us proceed to define and solve the optimization problem stated above. Let q_n denote the number of packets queued at the beginning slot n. The process q_n exhibits a nearest neighbor random walk whose transition behavior is depicted in figure 3-1.

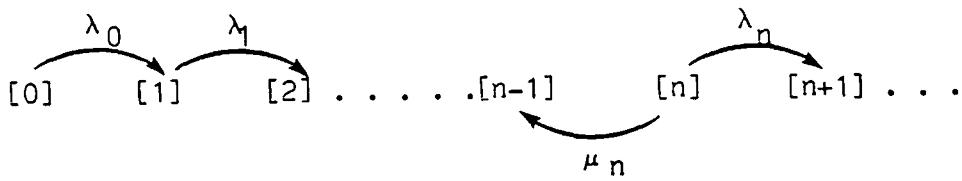

Figure 3-1: Transition behavior of the queueing process

Where the transition probabilities are given by:

$$\lambda_n = \begin{cases} \lambda & n=0 \\ \lambda x_n & n>0 \end{cases} \qquad \mu = \bar{\lambda}\,\bar{x}_n \tag{1}$$

(Here $\bar{x}$ denotes 1-x and $x_n \triangleq p_n \bar{s}$)

Let q_n denote the steady-state probability of n packets queued, then q_n is easily computed [KLIE 75] to be:

$$q_n = q_0 \left(\frac{\lambda}{\bar{\lambda}}\right)^n \prod_{i=1}^{n-1} \left(\frac{x_i}{\bar{x}_i}\right) \frac{1}{\bar{x}_n} \tag{2}$$

Where q_0 is the probability of an idle queue; this may be computed using the normalization condition: $1 = \sum_{n=0}^{\infty} q_n$.

The transform of the steady state distribution $Q(z) = \sum_{n=0}^{\infty} q_n z^n$, may be expressed in terms of the function $X(z) = \sum_{n=0}^{\infty} z^n \prod_{i=1}^{n} \left(\frac{x_i}{\bar{x}_i}\right)$ as follows:

$$Q(z) = q_0 \left(1 + \frac{\lambda z}{\bar{\lambda}}\right) X\left(\frac{\lambda z}{\bar{\lambda}}\right) \tag{3}$$

For example, when $p_n=1$ for all n (i.e., no packet is lost),

$$X(z) = \frac{1}{1 - \bar{s}z/s} \quad \text{and} \quad Q(z) = q_0 \frac{1 + \lambda z/\bar{\lambda}}{1 - \rho z} \quad \text{where } \rho \triangleq = \frac{\lambda \bar{s}}{\bar{\lambda} s}.$$

Given a choice of a loss policy, described by the sequence $\{p_n\}_0^\infty$ let L denote the expected loss and let T denote the expected delay. By definition, $L = \sum_{n=1}^{\infty} q_n \bar{p}_n \bar{s}$. Little's result [KLIE 75] implies $T=Q'(1)/\lambda$.

The loss policy (LP) problem to be solved may now be stated as follows:

<u>Given</u>: An acceptable bound on expected loss L

<u>Minimize</u>: The expected delay T

<u>With respect to</u>: The loss policy $\{p_n\}_0^\infty$.

The major result of this section is given by the following theorem.

<u>THEOREM 1</u>

Given a rate of arrival λ, a probability of successful transmission s and a bound on the acceptable expected loss L, there exists an integer $n_0=n_0(\lambda,s,L)$ such that the optimum loss policy $\{p_n\}_0^\infty$ is given by:

$$p_n = \begin{cases} 1 & n<n_0 \\ \\ 0 & n>n_0 \end{cases}$$

and p_{n_0} is a non-negative number smaller than 1.

This bang-bang principle for optimal loss-policy selection translates into a very simple loss-control protocol: no packets are discarded unless the queue of packets exceeds the threshold n_0; as soon as the queue of packets exceeds the threshold n_0, every new packet arrival causes the packet at the head of the queue to be discarded immediately following its transmission.

Let us proceed to prove theorem 1 and derive a closed form expression for the threshold $n_0(\lambda,s,L)$. Returning to equation (3), the normalization condition $Q(1)=1$ gives:

$$X\left(\frac{\lambda}{\bar{\lambda}}\right) = \frac{\bar{\lambda}}{q_0} \tag{4}$$

Equations (3) and (4) imply the following expression for the expected delay:

$$T = 1 + \frac{q_0}{\bar{\lambda}^2} X'\left(\frac{\lambda}{\bar{\lambda}}\right) \tag{5}$$

Therefore, minimization of the expected delay is equivalent to minimization of $q_0 X'\left(\frac{\lambda}{\bar{\lambda}}\right)$.

The expected loss L may be computed in terms of the loss transform $L(z) = \bar{s} \sum_{n=1}^{\infty} q_n \bar{p}_n z^n$, i.e., $L=L(1)$. A simple derivation yields:

$$L(z) = q_0 \left[s + X\left(\frac{\lambda z}{\bar{\lambda}}\right)\left(\frac{\lambda \bar{s}}{\bar{\lambda}} z - s\right) \right] \tag{6}$$

From which one may derive:

$$L = L(1) = s q_0 + \lambda - s \tag{7}$$

Finally, the probability of idle queue may be computed from equation (7) to be:

$$q_0 = \frac{1}{s}\left[L - \lambda + s\right] \tag{8}$$

Which, in turn, gives (using equation (4):

$$X\left(\frac{\lambda}{\bar{\lambda}}\right) = \frac{\bar{\lambda}s}{L-\lambda+s} \tag{9}$$

Let us define $\alpha \triangleq \frac{\bar{\lambda}s}{L-\lambda+s}$, the loss policy problem may now be re-stated as follows:

Minimize: $X'\left(\frac{\lambda}{\bar{\lambda}}\right)$

With respect to: $\{p_n\}_0^{\infty}$.

Subject to constraints: $X\left(\frac{\lambda}{\bar{\lambda}}\right) = \alpha$

In order to solve the LP problem another transformation is required to linearize the problem. Let $\underline{X}_0=1$ and $\underline{X}_n = \prod_{i=1}^{n} \left(\frac{\lambda x_i}{\bar{\lambda}\bar{x}_i}\right)$ when n>0, then $\underline{X}\left(\frac{\lambda}{\bar{\lambda}}z\right) = \sum_{n=0}^{\infty} X_n z^n$.

For every $n \geq 0$, $0 \leq p_n \leq 1$ implies $0 \leq \underline{X}_n \leq \rho\underline{X}_{n+1}$, where $\rho = \frac{\lambda\bar{s}}{\bar{\lambda}s}$; and vice versa: given numbers X_n such that $\underline{X}_0=1$ and $0 \leq X_n \leq \rho X_{n+1}$, there exist unique numbers $0 \leq p_n \leq 1$ such that $X_n = \prod_{i=1}^{n} \left(\frac{\lambda x_i}{\bar{\lambda}\bar{x}_i}\right)$ where $x_i \triangleq p_i \bar{s}$.

Therefore, rather than solving the LP problem by minimization with respect to $\{p_n\}_0^{\infty}$, we may minimize the delay with respect to $\{X_n\}_0^{\infty}$ subject to the constraints $0 \leq X_n \leq \rho X_{n+1}$.

Consider the space ℓ^1 [YOSI 71] of absolutely summable infinite sequences. Let us define a subset of ℓ^1

$$\mathscr{A} = \left\{ \{X_n\}_0 \,\middle|\, X_0 = 1,\ 0 \le X_n \le \rho X_{n+1} \text{ and } \sum_{n=0}^{\infty} X_n = \alpha \right\}$$

The set $\mathscr{A}$ is a convex and weakly compact [YOSI 71] subset of ℓ^1. The optimum loss policy problem may finally be stated as follows:

Minimize: $\sum_{n=0}^{\infty} nX_n$

With respect to: $\{X_n\}_0 \in \mathscr{A}$.

The LP problem is thus reduced to a minimization of a linear functional over a convex and compact set. Convex analysis ([HOLM 72], Lemma 5.d, page 10) shows that the optimum sequence $\{X_n\}_0^{\infty}$ is an extreme point of the set $\mathscr{A}$. Therefore, it is sufficient to compute the extreme points of $\mathscr{A}$ and select those points which minimize $\sum_{n=0}^{\infty} nX_n$.

LEMMA 1

The extreme points of the set $\mathscr{A}$ have the form:

$E_n = (1, \rho, \rho^2 \ldots\ldots \rho^{n-1}, a_n\rho^n, 0, 0\ldots\ldots)$

Where n and a_n satisfy the equality:

$$\sum_{i=0}^{n-1} \rho^i + a_n \rho^n = \frac{\lambda s}{L-\lambda+s} \tag{10}$$

and $a_n < 1$.

The proof of this lemma is straightforward and of little interest which is why we omit it. The characterization of the extreme points of $\mathscr{A}$ leads to the conclusion of the proof of theorem 1. All that is left is the computation of the optimum extreme point(s) $\underline{E}^n$ (i.e., the one minimizing $X'(\frac{\lambda}{\bar{\lambda}})$). The value of $X'(\frac{\lambda}{\bar{\lambda}})$ at $\underline{E}^n$ is:

$$\sum_{i=1}^{n-1} i\,\rho^{i-1} + n a_n \rho^n \tag{11}$$

Clearly this function is minimized when n is chosen to be minimal. Therefore let n_0 be the minimal value of n for which:

$$\frac{\bar{\lambda} s}{\bar{s}L - \lambda + s} - \sum_{i=1}^{n-1} \rho_i < \rho^n \tag{12}$$

this is the minimal value of n for which $\underline{E}^n \in \mathcal{S}$. The optimum solution of the LP problem is thus $\{X_n\}_0^{\infty} = \underline{E}^{n_0}$. The expressions $X_i = E_i^{n_0} = \rho^i$ for $i < n_0$ imply $x_i = \bar{s}$ for $i < n_0$ or, in other words, $p_i = 1$ for $i < n_0$. For $i > n_0$ $X_i = 0$, implying $x_i = 0$ or $p_i = 0$. This concludes the proof of theorem 1.

As a final touch let us use equation (12) to compute the value of optimum threshold $n_0 = n_0(\lambda, s, L)$.

$$\frac{1-\rho^{n_0}}{1-\rho} \leq \frac{\lambda s}{L - \lambda + s} < \frac{1-\rho^{n_0+1}}{1-\rho}$$

This in turn implies:

$$n_0 = \left\lfloor \frac{\ln[\,L/L-\lambda+s]}{\ln\rho} \right\rfloor \tag{13}$$

Usually the loss rate L will be specified as a fraction β of the arrival rate, i.e., $L = \beta\lambda$. What happens when the utilization factor ρ approaches the value 1 (equivalently when the arrival rate β approaches the service rate s)?

It is easy to see that the value of n_0, given by equation (13), approaches $\overline{s}/\beta$. This last expression gives the maximum buffer size required by the optimum real-time protocol when the acceptable loss fraction is β.

What happens when the acceptable loss rate β converges to 0? The protocol converges to ordinary positive acknowledgment with retransmission transport protocol. The buffer size requirement n_0 converges to infinity.

4. GENERALIZATIONS

The question arises: to what extent are the results of the previous section dependent upon the particular choice of model and optimization problem? In this section we show that the bang-bang loss control law remains valid under significant generalizations of the loss-policy problem.

THEOREM 2:

Let f and g be two continuous monotonic non-decreasing functions. Consider the following modified loss-policy problem:

Minimize: the function $f(T)$, where T is the expected delay.

With respect to: choice of a loss policy $\{p_n\}_0^\infty$

Subject to constraints: $g(L) \leq D$, where D is a constant.

There exists a threshold number n_0 such that the optimal loss policy is given by the following rule:

$$p_n = \begin{cases} 1 & n<n_0 \\ \\ 0 & n>n_0 \end{cases}$$

and p_{n_0} is a non-negative number smaller than 1.

The functions f and g may be considered as the cost of delay and the cost of loss respectively. The theorem then states that the bang-bang behavior is optimal whatever continuous monotonic pricing of delay and loss is considered.

An alternative optimization of real-time protocols might consider a dual minimization of the loss cost function g(L), subject to delay constraints. Again a bang-bang principle governs the optimum loss policy.

THEOREM 3:

Let f and g be two continuous monotonic non-decreasing functions and consider the problem of minimizing the loss cost function g(L) with respect to the choice of a loss policy $\{p_n\}_0^\infty$, subject to the delay cost constraints $f(T) \leq D$. The optimum policy is a bang-bang policy.

The proofs of theorems 2 and 3 are generalizations of the ideas presented in the proof of theorem 1. Therefore we only sketch the differences. Define the subsets of ℓ^1:

$$\mathcal{S}_2 = \left\{ \{X_n\}_0 \;\middle|\; X_0 = 1,\ 0 \leq X_n \leq \rho X_{n+1} \text{ and } g(L) \leq D \right\}$$

$$\mathcal{S}_3 = \left\{ \{X_n\}_0 \;\middle|\; X_0 = 1,\ 0 \leq X_n \leq \rho X_{n+1} \text{ and } f(T) \leq D \right\}$$

Here L is given by equation (7).

The two sets $\mathcal{S}_2$ and $\mathcal{S}_3$ are convex and weakly compact in ℓ^1, as may be trivially demonstrated. Consider first the minimization problem of theorem 2. The function f(T) is minimized when T is minimal (due to its montonicity). Therefore the problem is reduced to a minimization of T over a convex and compact set, just as in theorem 1.

Now consider the optimization problem of theorem 3. The function g[L] is minimized when L is minimized. Using equations (7) and (4) L is minimized when $X(\frac{\lambda}{\lambda})$ is maximized, this happens at an extreme point of $\mathcal{S}_3$ since $X(\frac{\lambda}{\lambda})$ is a linear functional of $\{X_n\}_0^{\infty}$.

So both problems can be reduced to that of identifying optimum extreme points of the respective sets $\mathcal{S}_i$ (i=2,3). It is easy to show that these extreme points correspond to bang-bang policies.

5. CONCLUSIONS

We saw that under most general assumptions an optimum loss-control policy for a real-time protocol should process packets just like an ordinary transport protocol, as long as the size of the packet queue does not exceed some threshold. As soon as the threshold is exceeded, any newly arriving packet causes the first packet in the queue to be lost.

REFERENCES

[BIALLY 80] Bially T., McLaughlin A. J., and Weinstein C. J.
Voice Communication in Integrated Digital Voice and Data Networks.
IEEE Transactions on Communications com-28(9), September, 1980.

[COHEN 78] Cohen D. and Sproul R.
High Level Protocols.
Proceedings of the IEEE 66(11), November, 1978.
Special Issue on Packet Communication.

[DSN 78] Carnegie-Mellon University.
Distributed Sensors Networks, Workshop Proceedings, 1978.

[HOLM 72] Holmes R. B.
A Course on Optimization and Best Approximation.
Springer Verlag, 1972.
Lecture Notes in Mathematics #257.

[KLIE 75] Klienrock L.
Queueing Systems.
John Wiley & Sons Inc., 1975.

[YOSI 71] Yosida K.
Functional Analysis.
Springer Verlag, 1971.