Optimizing Real-Time Transport Protocols

Technical Report No 115

Bruce Christianson

November 1990

Optimizing Real-Time Transport Protocols

Bruce Christianson

(comqbc@hatfield.ac.uk)

School of Information Sciences, Hatfield Polytechnic, Herts AL10 9AB, England, Europe

Abstract. A real-time transport protocol must trade delay against loss. Here we consider the problem of minimizing the mean transport delay subject to a given maximum acceptable loss rate. This form of the trade-off is appropriate for many voice, video and graphics applications.

Yechiam Yemini has shown that for this problem, and for a particular model of the transport mechanism, the optimal protocol is a combination of send and wait for ack with send and discard, where the decision to wait for ack or to discard is made by comparing the number of packets currently awaiting ack with a threshold value, which is itself calculated from the arrival and permitted loss rates. In this paper we generalize Yemini's result, and give an elementary proof.

1. The Transport Model. In our model, we assume that time is continuous. (A discrete time model is considered in [3]. Our arguments can readily be specialized to the discrete case.) Packets are assumed to have identical handling characteristics for the transport mechanism, and are distinguished only by their identity number.

The state of the transport mechanism is represented by a whole number n, which corresponds to the number of unacknowleged packets at the sender.

Events (state transitions) may occur at any time. Two events may occur arbitrarily closely in time but (we assume) cannot be simultaneous. There are two types of event, arrival events and departure events.

During any period of time, packets may arrive at the sender. Packet arrival occurs with rate λ (ie with mean frequency λ per unit elapsed time.) We assume that the arrival rate λ is known and fixed (in particular, is independent of the transport state.)

A packet which arrives when the transport state is n is discarded with probability p_n (where $0 \leq p_n \leq 1$) and queued for transmission with probability $\bar{p_n} = 1 - p_n$. In the latter case the transport state is incremented by one. Discarded packets contribute to the loss rate.

Also, during any period of time for which the transport state is non-zero, acknowlegements may be received for packets. This occurs with rate μ. At present, we suppose μ to be independent of the transport state (but see section 5.) When an acknowlegement is received, a packet is discarded (successfully) from the transport system, ie the transport state is decremented by one.

Packets awaiting acknowlegement may also be purged, ie dropped from the transport mechanism without an acknowlegement having been received. During the time for which the transport state is n where $n > 0$, packets are dropped at rate ν_n where $\nu_n \geq 0$. Discards of this type also contribute to the loss rate.

We assume that we are able to modify the protocol in such a way as to vary the p_n and ν_n without affecting μ (which we regard as fixed.)

This model of the transport system as a delay and loss mechanism is appropriate to ALOHA and TDMA networks, and to some token based ring protocols.

We have abstracted from the details of the transport protocol (re-transmission regime, timeout period etc) and regard the protocol as completely described for our purposes by specifying the numbers μ, $(p_n)_{n\geq 0}$ and $(\nu_n)_{n\geq 1}$ corresponding to the behaviour of the protocol for (at least) the given λ.

Our task is to determine the choice of $(p_n), (\nu_n)$ which minimizes the mean packet transit delay T subject to a maximum acceptable loss rate L (where L may in practice be specified as a fixed fraction of λ.)

To do this, we first characterize the variables L and T in terms of p_n, ν_n, μ and λ.

2. The Flow-Balance Equations. In this paper we adopt an operational approach to the transport model, rather than the more usual stochastic approach adopted in [3]. Although this simplifies our analysis, all our equations could be re-interpreted stochastically. Stochastic analysis generally requires introducing more assumptions of a type which is relatively difficult to verify. Our calculations concern actual observation sequences, rather than means over a number of sample paths. For a full discussion of the relative merits of operational analysis see [1].

The number of transitions of the transport state from n to $n+1$ and the number of transitions from $n+1$ to n over the same time period can differ by at most one. Consequently we have (over a sufficiently long period of time)

$$\lambda \bar{p_n} y_n = (\mu + \nu_{n+1}) y_{n+1}$$

where $y_n, 0 \leq y_n \leq 1$, is the proportion of time which the transport system spends in state n.

Let $x_n = \dfrac{\lambda \bar{p_n}}{\mu + \nu_{n+1}}$. Then $0 \leq \dfrac{y_{n+1}}{y_n} = x_n \leq \rho$, where $\rho = \lambda/\mu$, whence

$$y_n = y_0 \cdot \prod_{i=0}^{n-1} x_i$$

with y_0 being given by the constraint that $\sum_{n=0}^{\infty} y_n = 1$.

The arrival rate is λ and the success rate is $(1 - y_0)\mu$, since acknowlegements occur only when the transport state is non-zero. The loss rate is the difference, ie $\lambda - (1 - y_0)\mu$. Our constraint that the loss rate be at most L is therefore equivalent to demanding that

$$y_0 \leq \alpha \quad \text{where } \alpha = \frac{L - \lambda + \mu}{\mu}$$

Note that $\rho + \alpha > 1$.

It remains to find an expression for the mean transit delay.

3. Little's Law. Little's Law has the following operational form. For any system, and over any period of time, $\bar{N} = XR$ where X is the mean throughput of the system, R is the mean residence time of an object in the system, and $\bar{N}$ is the mean number of items in the system. The law holds exactly provided the system is empty at the start t_0 and at the end t_1 of the time period $\Delta t = t_1 - t_0$, and holds approximately provided that the number of items in the system at t_0 and at t_1 is small relative to the number of completions C which occur during the time period Δt. See [1] for further details.

The law can be simply proved by considering the area A under a graph of system occupancy N against time t.

$$\bar{N} = \frac{A}{\Delta t} = \frac{A}{C} \cdot \frac{C}{\Delta t} = RX$$

In our transport model, Little's Law implies that

$$T = \frac{1}{\lambda} \sum_{n=0}^{\infty} n.y_n$$

4. The Optimization Problem. We can now re-state the protocol optimization problem as follows.

Given $\rho, \alpha > 0$ with $\alpha + \rho > 1$, choose $(y_n)_{n=0}^{\infty}$ subject to constraints

$$\sum_{n=0}^{\infty} y_n = 1 \; ; \quad 0 \leq y_0 \leq \alpha \; ; \quad 0 \leq y_{n+1} \leq \rho y_n$$

so as to minimize the linear function

$$\sum_{n=0}^{\infty} n y_n$$

In Functional Analytic terms, the feasible set is a non-empty convex subset of the space l^1 of absolutely summable sequences, and is compact under the weak-* topology induced by the space c_0 of null sequences. Hence by the Krein-Milman Theorem (see for example [2]) the target function obtains its minimum at an extreme point of the feasible set, which can then be explicitly calculated.

However, for this particular problem we do not need such elaborate machinery. It is intuitively clear that the target function is minimized by spending as much time as possible in the states corresponding to low numbers of outstanding packets.

Whichever argument is followed, the conclusion is the same. The optimum transport delay under our constraints is obtained by the protocol for which

$$y_n = \alpha \rho^n \text{ for } n < n_0 \; ; \quad y_n = 0 \text{ for } n > n_0$$

where n_0 is chosen to allow the y_n to sum to 1 (always possible even if $\rho < 1$ since $\alpha/(1-\rho) > 1$.)

This corresponds to the regime

$$p_n = 0 \text{ for } n < n_0 \; ; \quad p_n = 1 \text{ for } n > n_0 \; ; \quad \nu_n = 0 \text{ for all } n \geq 1.$$

In fact the value of ν_n for $n > n_0$ is irrelevent, since the system never enters such states.

The critical value n_0 must therefore satisfy

$$\alpha\,\frac{1-\rho^{n_0}}{1-\rho} \leq 1 < \alpha\,\frac{1-\rho^{n_0+1}}{1-\rho}$$

since we have a geometric series, and solving this for n_0 gives

$$n_0 = \left\lceil \frac{\log\frac{\alpha-1+\rho}{\alpha}}{\log\rho} \right\rceil = \left\lceil \frac{\log\frac{L}{L-\lambda+\mu}}{\log\lambda/\mu} \right\rceil$$

In case $\rho = 1$ we have $n_0 = \lceil 1/\alpha \rceil$.

5. Conclusions. The optimum transport protocol requires only a finite buffer capacity of length n_0. Packets which arrive when all n_0 buffers are full should be discarded. Once buffered for transmission, packets should always be retained until acknowleged. This is in fact a commonly used "dirty" protocol, but the crucial choice of n_0 is not usually made.

It is worth noting that the parameter T which we have minimized is in fact the mean length of time for which a packet occupies a buffer, which is usually the quantity of interest for the network manager. This time is (in general) not quite the same as the mean packet transit delay, which is the quantity of interest for network users.

Allowing the acknowlegement rate μ_n to depend on the transport state n leads (via a slightly modified calculation) to the same optimal protocol, provided that we retain the *homogeneity* assumption [1] that the μ_n's do not depend upon the p_n's or ν_n's.

In this more general case, n_0 is the least n for which we have

$$1 + \sum_{i=1}^{n} \prod_{j=1}^{i} \frac{\lambda}{\mu_j} > \frac{1}{\alpha}$$

Although the value of n_0 is not now given by an explicit formula, it is still easy to calculate n_0 numerically for given μ_n. It would be of interest to discover a generalization of this result appropriate for modelling a sliding window protocol which allows μ_n to depend directly upon the y_n, or even upon the detailed state transition regime.

References

[1] P.J. Denning and J.P. Buzen, 1978, *The Operational Analysis of Queueing Network Models*, ACM Computing Surveys **10**(3) 225–261

[2] J. Lindenstrauss, 1988, *Some Useful Facts about Banach Spaces*, Geometric Aspects of Functional Analysis, Springer Lecture Notes in Mathematics **1317** 185–200

[3] Y. Yemini, 1983, *A Bang-Bang Principle for Real-Time Transport Protocols*, ACM SIGCOM 1983 Symposium, Computer Communication Review **13**(2) 262–268