AN ADAPTIVE TIMEOUT ALGORITHM FOR RETRANSMISSION ACROSS A PACKET SWITCHING NETWORK

Stephen William Edge

Department of Computer Science

University College London

ABSTRACT

This paper presents and analyses a simple algorithm for setting an adaptive timeout value at a source Host for end-to-end retransmission on a packet-switched connection. The algorithm allows the recipient Host to acknowledge arriving data in either original transmission order or out-of-order. The timeout at the source Host is determined from current estimates - using exponentially weighted moving averages - of the mean and variance of successive acknowledgement delays. We show that when these delays are random variables forming certain stationary or non-stationary stochastic processes, the ensuing timeout gives a near-minimum retransmission delay, subject to some specified limit on the amount of unnecessary retransmission. This property is illustrated for a simulated sequence of acknowledgement delays obtained from loop delay measurements.

1. Introduction

End-to-end retransmission in a packet switching network usually occurs in a transport protocol, such as the Transmission Control Protocol (TCP) [Cerf 74, Postel 80] or the recent ISO protocol [ISO 82]. Its main purpose then is to provide a reliable communication path on a connection between two communicating processes.

The retransmission mechanism normally used in transport protocols - e.g. in TCP and the ISO protocol - has been called **Sequencing Positive Acknowledgement Retransmission (SPAR)** [Sunshine 75]. With this, the data stream for each direction on a connection is logically subdivided into data units - e.g. packets or bytes - each implicitly or explicitly carrying a sequence number. This sequencing allows arriving data to be correctly ordered at the receiver with any duplicates being removed. Reordered data is then acknowledged by the receiver by returning the sequence number of the next expected data unit. This implicitly acknowledges all preceding data. At the sender, unacknowledged data units are retransmitted following a certain timeout period since their previous transmission or retransmission. In some cases - e.g. the ISO transport protocol - only the earliest unacknowledged data unit would be retransmitted in this way.

Another retransmission mechanism, commonly used inside a network but also possible in theory for end-to-end use [Edge 83], is **unsequenced Positive Acknowledgement Retransmission (PAR).** This can behave identically to SPAR at a sender, but at a receiver each data unit must be acknowledged individually. However, this can occur on arrival rather than after data reordering.

The timeout algorithm presented here applies to both SPAR and PAR. Its general aims are described in the next section, and the algorithm itself is defined in section 3. Section 4 validates the algorithm analytically under both stationary and non-stationary conditions. The algorithm is then tested in section 5 against a sequence of simulated acknowledgement delays obtained from local area loop delay measurements.

2. The Aims of Adaptive Timeout Estimation

An end-to-end timeout has two possible aims: to ensure adequately low retransmission delay and to avoid most unnecessary retransmission. While the former requires a small timeout, the latter demands a timeout in excess of most acknowledgement delays.

When unnecessary retransmission alone is important, a timeout could be set to the largest possible acknowledgement delay, as obtained from the maximum combined lifetimes of any packet and its acknowledgement in transit and at a receiver [ISO 82]. However, such a timeout would almost certainly be several seconds or more, since even normal end-to-end response times can be around one second [Clipsham 76, Kleinrock 76 p.457, Rajaraman 78]. This delay could be unacceptable for interactive or real time traffic if uncorrected data loss occurred frequently - e.g. from transmission errors in a packet radio network or from buffer depletion in nodes, Gateways or the receiving Host.

A smaller retransmission timeout allowing a limited amount of unnecessary retransmission may thus be desirable in certain cases. However, such a timeout cannot necessarily be determined from advance calculation, since it would depend on the distribution of acknowledgement delay, requiring both foreknowledge of the latter and stability. To avoid these prerequisites, we define the following adaptive algorithm.

3. The Adaptive Timeout Algorithm

The timeout algorithm has two parts. In the first part, measurements of successive SPAR or PAR acknowledgement delays are used to update estimates, T_n and V_n, of the average and variance, respectively, of acknowledgement delay, where n ($n \geq 0$) gives the number of previous updates and T_0 and V_0 are the initial estimates. In the second part, a suitable retransmission timeout, R_n, is obtained from T_n and V_n.

The first part proceeds as follows. Each time that a returning acknowledgement acknowledges one or more outstanding data units (e.g. bytes, packets or letters), the time period, t_n, is obtained since the initial transmission of the first of these data units. This is used to update T_{n-1} and V_{n-1} as follows:

$$T_n = b\,T_{n-1} + t_n / a \qquad (a \geq 1,\ n \geq 1) \qquad (1)$$

$$V_n = d\,V_{n-1} + (t_n - T_{n-1})^2 / c \qquad (c \geq 1,\ n \geq 1) \qquad (2)$$

where

$$b = (a - 1) / a \qquad (3)$$

$$d = (c - 1) / c \qquad (4)$$

The second part updates the retransmission timeout, R_n, for some retransmission limit, Y, by:

$$R_n = T_n + e\,[V_n (1 - Y) / Y]^{1/2} \qquad (0 < Y < 1,\ e \leq 1) \qquad (5)$$

Later analysis shows the parameter **c** in Eq (2) should be large (e.g. between 10 and 20), whereas **a** in Eq (1) should be determined dynamically to minimise V_n. The parameter **Y** in Eq (5) defines the maximum probability of unnecessary retransmission for SPAR between two consecutive updates of R_n. Under certain conditions, it also defines the maximum probability of unnecessarily retransmitting the data associated with t_n for SPAR and PAR. To make this guarantee, the parameter **e** should equal one. However, for certain distributions of acknowledgement delay that can occur in practice, e = 0.4 will reduce R_n without significantly exceeding the retransmission limit.

When retransmitted data is acknowledged following an earlier loss, the updates of T_{n-1}, V_{n-1} and R_{n-1} should be excluded since t_n could be unrepresentative of normal acknowledgement delay even when measured correctly from the time of retransmission. However, updates must be included for large acknowledgement delays causing unnecessary retransmission. To distinguish these alternatives, retransmissions and their acknowledgements could be labelled - e.g. using an extra data sequence bit. Otherwise the following assumptions can be made: assume unnecessary retransmission (and include retransmission updates) for a low frequency of loss; assume necessary retransmission (and exclude updates) for a high loss frequency (e.g. exceeding Y); otherwise assume unnecessary retransmission only when $t_n < RET_{n-1}$, where RET_{n-1} is an updated estimate for the minimum acknowledgement delay of a retransmission (e.g. $RET_{n-1} = R_{n-1} + MIN[t_1, t_2, \ldots t_{n-1}]$).

In the second and third cases above, T_n and V_n would be mean and variance estimates for the acknowledgement delay distribution truncated at approximately R_{n-1} and RET_{n-1}, respectively. This will reduce T_n and (most likely) V_n, thereby also reducing R_n (and RET_n) and allowing more unnecessary retransmission. This is illustrated later.

4. Analysis of the Algorithm

In this section, we validate the algorithm when the acknowledgement delays are derived from some stochastic process. We begin by analysing the mean and variance estimates, T_n and V_n, for a weakly stationary stochastic process of order two and a non-stationary random walk. We then derive the timeout formula in Eq (5). Throughout, **E**, **VAR** and **COV** will denote expectation, variance and covariance, respectively, and $\mathbf{E^c}$, $\mathbf{VAR^c}$ and $\mathbf{COV^c}$ will denote these terms conditioned at any time on all previous acknowledgement delays.

4.1 Weakly Stationary Acknowledgement Delays

We assume here that the acknowledgement delays t_n ($n \geq 1$) are random variables with the same positive and finite mean, **T**, and variance, **V**. We also assume that the correlation coefficient, $\mathbf{v}_j$, of t_n and t_{n-j} is independent of n ($1 \leq j \leq n-1$, $n \geq 2$). To analyse Eqs (1) and (2) under these conditions, we restate them as follows:

$$T_n = \sum_{j=0}^{n-1} b^j\, t_{n-j} / a + b^n T_0 \qquad (n \geq 1) \qquad (6)$$

$$V_n = \sum_{j=0}^{n-1} d^j (t_{n-j} - T_{n-j-1})^2 / c + d^n V_0 \qquad (n \geq 1) \qquad (7)$$

The right hand sides of the above equations are called **exponentially weighted moving averages (EWMAs)** [Box 76 p.145 Cox 65, p.303], due to their being weighted averages with geometrically decreasing weights. EWMAs are particularly useful for statistical forecasting, as shown later. Taking the expectation, E, and variance, VAR, of Eq (6) using ($0 \leq b < 1$) yields:

$$E(T_n) = \sum_{j=0}^{n-1} b^j T / a + b^n T_0 \qquad (n \ge 1)$$

$$= T + b^n (T_0 - T) \qquad (n \ge 0)$$

$$\rightarrow T \qquad (n \rightarrow \infty) \qquad (8)$$

$$VAR(T_n) = \sum_{j=0}^{n-1} b^{2j} V / a^2 + 2 \sum_{j=0}^{n-2} \sum_{k=j+1}^{n-1} b^{j+k} COV(t_{n-j}, t_{n-k}) / a^2 \qquad (n \ge 2) \qquad (9)$$

$$\le \sum_{j=0}^{n-1} b^j \sum_{k=0}^{n-1} b^k V / a^2 \qquad (n \ge 1)$$

$$< V \qquad (10)$$

Equation (8) shows that the expectation of T_n approaches the mean acknowledgement delay, T, for large enough n, while Eq (10) shows that T_n will fluctuate less than the acknowledgement delays, due to a smaller variance. This variance can be reduced further if the correlation coefficients, $COV(t_n, t_{n-j}) / V$, are bounded by some non-negative value, m, according to:

$$\sum_{j=1}^{r} COV(t_n, t_{n-j}) / V \le m \qquad (11)$$
$$(1 \le r \le n-1, \quad n \ge 2)$$

Equation (9) then gives for $(n \ge 2)$:

$$VAR(T_n) = \sum_{j=0}^{n-1} b^{2j} V / a^2 + 2 \sum_{j=0}^{n-2} b^{2j} \sum_{k=1}^{n-j-1} b^k COV(t_{n-j}, t_{n-j-k}) / a^2$$

$$< \sum_{j=0}^{\infty} b^{2j} V / a^2 + 2 \sum_{j=0}^{\infty} b^{2j} b\, m\, V / a^2$$

$$< (2m + 1) V / (2a - 1) \qquad (12)$$

Provided t_n is correlated strongly with no more than a few of the preceding delays, m in Eq (11) could be small. This is not unrealistic: ARPANET measurements in [Kleinrock 76, p.458] indicate $m \le 2$, while independent delays would give $m = 0$. A large value for the parameter a (e.g. $a \ge 20$) would then greatly reduce the variance of T_n in Eq (12), making T_n itself an accurate estimate for T.

We now obtain the expectation of V_n using $(0 \le b < 1)$, $(0 \le d < 1)$ and (for the first time) $v_j = COV(t_n, t_{n-j}) / V$ with $(-1 \le v_j \le 1)$. We begin from Eq (6) to obtain:

$$VAR(T_n) = \begin{cases} 0 & (n = 0) \\ V / a^2 & (n = 1) \\ \sum_{j=0}^{n-1} b^{2j} V / a^2 + 2 \sum_{j=1}^{n-1} b^j \sum_{k=0}^{j-1} b^k v_{j-k} V / a^2 & (n \ge 2) \end{cases}$$

thus for $(1 \le i \le n, \ n \ge 2)$:

$$|VAR(T_n) - VAR(T_{n-i})|$$

$$= \left| \sum_{j=n-i}^{n-1} b^{2j} V / a^2 + 2 \sum_{\substack{j=n-i \\ j \ne 0}}^{n-1} b^j \sum_{k=0}^{j-1} b^k v_{j-k} V / a^2 \right|$$

$$< b^{2n-2i} \sum_{j=0}^{i-1} b^{2j} V / a^2 + 2 b^{n-i} \sum_{j=0}^{i-1} b^j V / a$$

$$< b^{n-i} V / (2a - 1) + 2 b^{n-i} V$$

Also from Eq (6) for $(n \ge 1)$:

$$COV(t_{n+1}, T_n) = \sum_{k=0}^{n-1} b^k COV(t_{n+1}, t_{n-k}) / a$$

$$= \sum_{k=0}^{n-1} b^k v_{k+1} V / a$$

so for $(1 \le i \le n, \ n \ge 2)$:

$$|COV(t_{n+1}, T_n) - COV(t_{n-i+1}, T_{n-i})|$$

$$= \left| \sum_{k=n-i}^{n-1} b^k v_{k+1} V / a \right|$$

$$< b^{n-i} V$$

Finally for $(n \ge 2)$ in Eq (7):

$$E(V_n) = \sum_{i=1}^{n} d^{i-1} [VAR(t_{n-i+1} - T_{n-i}) + (E(t_{n-i+1} - T_{n-i}))^2] / c + d^n V_0$$

$$= \sum_{i=1}^{n} d^{i-1} [V + VAR(T_{n-i}) - 2 COV(t_{n-i+1}, T_{n-i}) + b^{2n-2i} (T - T_0)^2] / c + d^n V_0$$

giving:

$$|E(V_n) - \sum_{i=1}^{n} d^{i-1} [V + VAR(T_n) - 2\,COV(t_{n+1}, T_n)] / c|$$

$$< \sum_{i=1}^{n} d^{i-1} b^{n-i} [V / (2a - 1) + 4V + (T - T_0)^2] / c + d^n V_0$$

$$= \sum_{i=1}^{n} (d / b)^{i-1} O(b^{n-1}) + O(d^n)$$

and thus:

$$|E(V_n) - (1 - d^n) VAR(t_{n+1} - T_n)|$$

$$< \begin{cases} O(n\, b^{n-1}) + O(d^n) & (b = d) \\ O(|b^n - d^n|) + O(d^n) & (b \neq d) \end{cases}$$

$$E(V_n) \rightarrow VAR(t_{n+1} - T_n) \qquad (n \rightarrow \infty) \qquad (13)$$

Equation (13) shows that the expectation of V_n approaches the variance of $(t_{n+1} - T_n)$ for large enough n. In order for V_n to approximate $VAR(t_{n+1} - T_n)$ as well, $VAR(V_n)$ should be small. Provided $VAR[(t_n - T_{n-1})^2]$ $(n \geq 1)$ is either constant or replaced by some upper bound, the derivation of $VAR(V_n)$ is identical to that of $VAR(T_n)$ previously. This is due to the functional equivalence of Eqs (6) and (7), with a, b, T_0 and t_n $(n \geq 1)$ in the former corresponding to c, d, V_0 and $(t_n - T_{n-1})^2$, respectively, in the latter. The V_n-analogues of Eqs (11) and (12) then imply that $VAR(V_n)$ would be reduced by a large parameter, c, if the updates $(t_n - T_{n-1})^2$ $(n \geq 1)$ were correlated weakly by an analogue of Eq (11).

To obtain weakly correlated updates for V_n we argue as follows. If the acknowledgement delays t_n are extended by fictitious additions $\{t_0, t_{-1}, t_{-2}, ..\}$ and derived from a strictly stationary process (where the joint probability distribution of any set $\{t_{n1+j}, .. t_{nr+j}\}$ is independent of j), then any delay t_{n+1} will be uniquely represented by an infinite linear regression, $(k_1 t_n + k_2 t_{n-1} + k_3 t_{n-2} + ..)$, plus an uncorrelated residual, X_{n+1}, derived from another strictly stationary process with zero mean and no autocorrelation [Cox 65 p.286, Wold 54 p.84 p.89]. The infinite regression $(k_1 t_n + k_2 t_{n-1} + \dots)$ gives the best linear forecast for t_{n+1} given $\{t_n, t_{n-1}, \dots\}$, since the mean square of the forecast error, X_{n+1}, is less than for any other linear forecast [Cox 65, p.302]. Although the EWMA formula for T_n in Eq (6) cannot provide this optimum linear forecast for a stationary stochastic process, examples show that it can sometimes produce a similar mean square forecast error [Cox 61]. This would occur by setting the parameter, a, in Eq (6) to minimise $E[(t_{n+1} - T_n)^2]$ and thus $VAR(t_{n+1} - T_n) = E(V_n)$. Prediction errors, $(t_{n+1} - T_n)$, close enough to X_{n+1} would then produce the required uncorrelated updates, $(t_{n+1} - T_n)^2 \simeq X_{n+1}^2$, for V_n, if the residuals, X_n, were independent of one another.

Since obtaining V_n accurately turns out later to be more important than reducing fluctuation in T_n, we are lead to propose a large value for the parameter, c, in Eq (2), and a dynamically adjusted value for the parameter, a, in Eq (1) to minimise V_n. However, to ensure rapid convergence of $E(T_n)$ and $E(V_n)$ in Eqs (8) and (13), neither $b = 1 - 1/a$ nor $d = 1 - 1/c$ can be too close to one. As an example, to obtain b^n or $d^n < 0.1$, $a \leq 9$ or $c \leq 9$ is required for n = 20 updates, and $a \leq 22$ or $c \leq 22$ is required for n = 50 updates. These indicate suitable maximums for a and c.

4.2 Non-Stationary Acknowledgement Delays

Here, we consider acknowledgement delays with no steady mean or variance. One simple case occurs when the previous weakly stationary delays change suddenly to a new stationary sequence after the mth $(m \geq 1)$ update. The values of T_m and V_m immediately preceding the change then become the initial estimates for T_n and V_n $(n > m)$ for the new sequence. The subsequent behaviour of T_n and V_n is then covered by the previous subsection.

Another case occurs when the acknowledgement delays are formed from a random walk according to:

$$t_n = A_0 + h\, z_n + \sum_{j=0}^{n-1} z_{n-j} / g \qquad (14)$$
$$(g \geq 1, \quad n \geq 1)$$

where

$$h = (g - 1) / g \qquad (15)$$

A_0 is an arbitrary positive number

and $\{z_1, z_2 .. \}$ is a sequence of independent random variables with mean 0 and variance U

Equation (14) allows the acknowledgement delays to drift from an initial value A_0 by incremental additions, z_n / g. Transitory "shocks", $h\, z_n$, also produce temporary fluctuation. We show in the Appendix that T_n and V_n will then satisfy:

$$E(T_n) \rightarrow E(t_{n+1}) \qquad (n \rightarrow \infty) \qquad (16)$$

$$E(V_n) \rightarrow VAR(t_{n+1} - T_n) \qquad (n \rightarrow \infty) \qquad (17)$$

and

$$T_n \rightarrow E^c(t_{n+1}) \qquad (n \rightarrow \infty,\ b = h) \quad (18)$$

$$E(V_n) \rightarrow VAR^c(t_{n+1} - T_n) \quad (n \rightarrow \infty,\ b = h) \quad (19)$$

Equations (16) and (17) match Eqs (8) and (13) in the previous subsection, with convergence also depending on terms of order b^n and d^n. When the parameters b and h are equal, Eqs (18) and (19) show that T_n and V_n become much more reliable conditional mean and variance estimates for the particular preceding delays $\{t_1, t_2, \ldots t_n\}$. At the same time, we show in the Appendix that the mean square forecast error $E[(t_{n+1} - T_n)^2]$, which equals $E(V_n)$ by Eqs (16) and (17), is minimised at $VAR(z_{n+1}) = U$. This optimum forecast by the EWMA formula for T_n is well known for the stochastic process of Eq (14) [Box 76 p.144, Cox 65 p.303, Muth 60]. It gives $(t_{n+1} - T_n) = z_{n+1}$ for large enough n, so the updates $(t_{n+1} - T_n)^2$ for V_n become independent random variables with constant variance provided $E(z_n^4)$ is also constant. This allows a large parameter c to give $V_n \approx E(V_n)$, as discussed in the previous subsection. The prior condition for this, b = h, can then be ensured by dynamically setting the parameter, a, to minimise $E(V_n)$ and thus (approximately) to minimise V_n. This treatment of the parameters a and c corresponds to that proposed previously.

In [Edge 84], we show that Eqs (16) and (17) also apply when the acknowledgement delays t_n are the sum of a weakly stationary stochastic process and an uncorrelated non-stationary process defined as in Eq (14). This may be the most realistic case, since the stationary component can include delays subject to rapid fluctuation - e.g. delay at the receiver - while the non-stationary component can include more gradual delay variation due to changes in network traffic.

4.3 The Timeout Formula

The two previous subsections have shown that the mean and variance estimates, T_n and V_n, satisfy $E(T_n) \rightarrow E(t_{n+1})$ and $E(V_n) \rightarrow VAR(t_{n+1} - T_n)$ for large enough n. A probability inequality in [Feller 66, p.150] - that is slightly tighter than Chebyshev's inequality - then gives for large n:

$$Pr(t_{n+1} - T_n > X) \le E(V_n) / [E(V_n) + X^2] \qquad (X > 0) \qquad (20)$$

giving

$$Pr(t_{n+1} > R_n) \le E(V_n) / [E(V_n) + (R_n - T_n)^2] \qquad (R_n > T_n) \qquad (21)$$

Now suppose that only the current timeout value initiates retransmission and that this occurs (unnecessarily) for R_n. Provided data is acknowledged sequentially, as for SPAR, and the measured delay t_{n+1} updates R_n immediately, then $t_{n+1} > R_n$ is implied, since t_{n+1} is then obtained for the earliest unacknowledged data unit during the lifetime of R_n. Thus if $Pr(t_{n+1} > R_n) \le Y$ $(0 < Y < 1)$, the probability of R_n causing unnecessary retransmission would not exceed Y. The latter is then guaranteed for large n from Eq (21) if:

$$R_n > T_n \quad \text{with} \quad E(V_n) / [E(V_n) + (R_n - T_n)^2] \le Y$$

requiring

$$R_n \ge T_n + [E(V_n)\,(1 - Y) / Y]^{1/2}$$

Thus, assuming V_n is an accurate estimate for $E(V_n)$, the minimum timeout, R_n, giving a maximum probability Y of unnecessary retransmission during its lifetime should be obtained as:

$$R_n = T_n + [V_n\,(1 - Y) / Y]^{1/2} \qquad (n \ge 1) \quad (22)$$

This is the formula given earlier in Eq (5) for e = 1. Notice that while V_n must approximate $E(V_n)$ accurately, T_n need not approximate $E(T_n)$, since both its variance and its covariance with t_{n+1} are included in V_n by Eq (13). This leads to our distinct treatment of the parameters, a and c.

The above retransmission guarantee for SPAR can be both improved and applied to PAR if the random duration, X_n say, of R_n satisfies $R_n - X_n \le R_{n-1}$ - e.g. due to infrequent acknowledgement. In this case, $(t_{n+1} - X_n - X_{n-1} \ldots - X_{i+1} > R_i)$ implies $(t_{n+1} > R_n)$ for $(i < n)$. Thus, unnecessary retransmission of the data associated with t_{n+1} on any previous timeout R_i $(i \le n)$ implies $(t_{n+1} > R_n)$. Such retransmission then has a maximum probability, Y, with Eq (22). This directly limits the amount of unnecessary retransmission.

In general, Eq (22) will give larger values of R_n than the exact value satisfying $Pr(t_{n+1} > R_n) = Y$ since it guarantees $Pr(t_{n+1} > R_n) \le Y$ for any distribution of $(t_{n+1} - T_n)$ in Eq (20). Illustrations of this phenomenon are given in Table (1) for an exponential and two Erlangian distributions (with mean T and variance V) of the acknowledgement delays, t_n. To obtain lower more accurate values of R_n in Table (1), the following alternative formula is also shown:

$$R_n = T_n + 0.4\,[V_n\,(1 - Y) / Y]^{1/2} \qquad (23)$$

Equation (23) corresponds to Eq (5) with e = 0.4. Although it has no widespread generality (the constant 0.4 being purely arbitrary), the tendency of measured acknowledgement delays to assume an Erlangian-like (bell shaped) distribution [Gien 78, Kleinrock 76 p.457] could make it useful.

Ack. Delay Distribution	Y	Timeout R_n / T Exact	Eq(22)*	Eq(23)*
Exponential	.1	2.30	4.00	2.20
	.05	3.00	5.36	2.74
	.02	3.91	8.00	3.80
4-stage Erlangian	.1	1.67	2.50	1.60
	.05	1.94	3.18	1.87
	.02	2.27	4.50	2.40
25-stage Erlangian	.1	1.26	1.60	1.24
	.05	1.35	1.87	1.35
	.02	1.45	2.40	1.56

*assuming $T_n = T$ and $V_n = V$

Table (1) Normalised Timeout for a Probability Y of Unnecessary Retransmission

5. Illustration of the Algorithm

In this section, we apply the timeout algorithm to a sequence of simulated acknowledgement delays obtained from loop delay measurements over a Cambridge Ring. These measurements were obtained for the double loop journey shown in Fig (1) in which each packet and its (sequentially received) echoed reply execute a round trip [Alfano 83]. The delay distribution for 1980 packet transfers - spread over 52 minutes of heavy ring usage - is shown in Fig (2). This distribution has the Erlangian-like shape mentioned above. The average delay for the echos received over consecutive one minute intervals is also shown in Fig (3).

To apply the algorithm to these measurements, we take the nth delay measurement as t_n, and use initial estimates for T_n and V_n of $T_0 = 4$ seconds and $V_0 = 4$ seconds^2. Following an initial run up period of $n = 100$ updates, we obtain the average values of V_n and R_n ($101 \le n \le 1980$) and the minimum fraction of packets "retransmitted" due to ($t_n > R_{n-1}$). These statistics (in seconds or sec.^{2}) are shown in Table (2) for several values of the parameters a and c, and for both ($e = 1$) and ($e = 0.4$) in Eq (5). The maximum required retransmission probability is $Y = 0.1$.

Table (2) shows that with ($e = 1$), retransmission is well below the maximum level of $Y = 0.1$, whereas with ($e = 0.4$), it approaches or exceeds this limit due to smaller timeouts, R_n. This behaviour agrees with Table (1). Varying the parameter a as suggested earlier to minimise V_n maximises retransmission in Table (2) and minimises R_n. This is desirable when ($e = 1$), since retransmission is still less than 0.1. However, for ($e = 0.4$), there is less benefit, since retransmission can exceed the limit. Increasing the parameter c (with $a = 6$ minimising V_n) has no effect on V_n, but increases R_n and reduces retransmission. Our recommendation for large c thus makes the algorithm more conservative. As large c should reduce the variance, VAR, of V_n rather than its expectation, E, no change in the average of V_n is unsurprising. Conversely, the increase in R_n can be attributed to increased $V_n^{1/2}$ in Eq (5) (on average) due to less variation of V_n in the identity $E(V_n^{1/2})^2 = E(V_n) - VAR(V_n^{1/2})$.

Table (3) shows similar results for some other retransmission limits, Y. For the lower results, the acknowledgement delays are truncated at R_{n-1} (see section 3), with ($t_n > R_{n-1}$) then giving $T_n = T_{n-1}$, $V_n = V_{n-1}$ and $R_n = R_{n-1}$. It can be seen that without truncation, the retransmission limits, Y, are, at worst, only slightly exceeded. With truncation, T_n, V_n and R_n are reduced, causing retransmission to increase. This effect is greater with either ($e = 0.4$) or large Y, due to the reduction of the truncation point R_n in

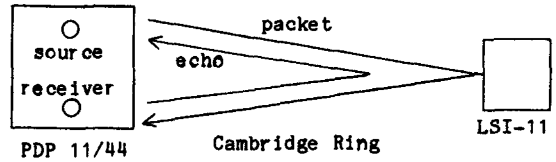

Figure (1) Double Loop Delay Measurements

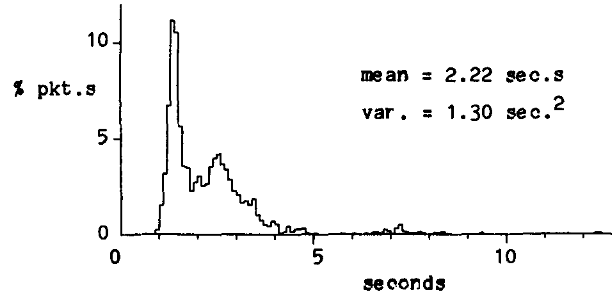

Figure (2) Delay Histogram (0.1 sec. bins)

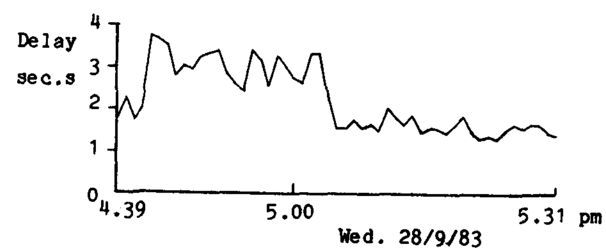

Figure (3) Average Delay at 1 Minute Intervals

c	a	Averages $101 \le n \le 1980$			Averages	
		V_n	R_n	Retran.	R_n	Retran.
		e = 1			e = 0.4	
10	1	0.96	4.62	.018	3.19	.079
	2	0.81	4.39	.023	3.10	.087
	4	0.77	4.33	.024	3.07	.093
	6	0.77	4.32	.023	3.07	.096
	10	0.77	4.34	.023	3.08	.103
	15	0.78	4.36	.022	3.09	.101
	30	0.81	4.42	.021	3.12	.096
2	6	0.77	3.89	.052	2.90	.140
5		0.77	4.14	.036	3.00	.120
10		0.77	4.32	.023	3.07	.096
20		0.77	4.50	.019	3.14	.083
30		0.77	4.58	.017	3.18	.073

Table (2) Algorithm Performance with 1980 updates
$[T_0 = V_0 = 4 \text{ secs.}^{(2)}, \quad Y = 0.1]$

Eq (5). This behaviour suggests that in practice, Y should be reduced or e increased (if e < 1) if truncation produces excessive retransmission. There should also be a short period without truncation in case T_0 and V_0 are initially too low.

We note finally that without the algorithm and a static timeout, R_n, retransmission (without the run up) would equal the proportion of delays in Fig (2) exceeding R_n. Limited retransmission with low R_n will thus require knowledge of the delay distribution, while the changes in Fig (3) would make this difficult to determine in advance.

ACKNOWLEDGEMENTS

I would like to thank Peter Higginson and N.P. Alfano for making available the loop delay measurements used in section 5. The referees also made some useful comments.

REFERENCES

[Alfano 83] "A Generalised Test Harness for Network Protocols" - N.P. Alfano, M.Sc Thesis and Indra Note 1482, Dept. of Computer Science, University College London, October 1983.

[Box 76] "TIME SERIES ANALYSIS forecasting and control - Revised Edition" - G.E.P. Box and G.M. Jenkins, publisher Holden-Day, 1976.

[Cerf 74] "A Protocol for Packet Network Intercommunication" - V.G. Cerf, R.E. Kahn, IEEE Trans. on Comm. Vol. COM-22 No.5, May 1974, pp.637-648.

[Clipsham 76] "Datapac Network Overview" - W.W. Clipsham, F.E. Glave, M.L. Narraway, ICCC Toronto 1976, pp.131-135.

[Cox 61] "Prediction by Exponentially Weighted Moving Averages and Related Methods" - D.R. Cox,

Y	Averages : e = 1				e = 0.4	
	T_n	V_n	R_n	Ret.	R_n	Ret.
	No Truncation				No Trunc	
.02	2.24	0.77	7.11	.007	4.19	.025
.05	2.24	0.77	5.27	.015	3.45	.059
.1	2.24	0.77	4.32	.023	3.07	.096
.2	2.24	0.77	3.63	.049	2.79	.153
	Truncation				Truncation	
.02	2.14	0.25	5.17	.018	2.94	.092
.05	2.11	0.18	3.79	.035	2.01	.470
.1	2.07	0.13	3.03	.082	1.48	.815
.2	1.87	0.07	2.32	.305	1.35	.943

Table (3) The Algorithm with Truncation at R_n
$[T_0 = V_0 = 4 \text{ sec}^{(2)}, \quad a = 6, \quad c = 10]$

Journal of the Royal Statistical Soc. Series B, Vol.23, 1961, pp.414-422.

[Cox 65] "The Theory of Stochastic Processes" - D.R. Cox, H.D. Miller, publisher Methuen, 1965.

[Edge 83] "Connection Flow Control is a Wide Area Packet Switching Network" - S.W. Edge, Ph.D Thesis and TR 88, Dept. of Computer Science, University College London, September 1983.

[Edge 84] "Analysis of an Adaptive Timeout Algorithm" - S.W. Edge, Indra Note 1523, Dept.of Computer Science, University College London, January 1984.

[Feller 66] "An Introduction to Probability Theory and Its Applications - Volume 2, 1st Edition" - W. Feller, published by Wiley, 1966.

[Gien 78] "Performance Evaluations in CYCLADES" - M. Gien, J.L. Grange, ICCC September 1978, Kyoto, Japan, pp.23-32.

[ISO 82] "Information Processing Systems - Open Systems Interconnection - Transport Protocol Specification" - ACM SIGCOMM Vol.12 No.s 3 & 4, July/October 1982, pp.24-67.

[Kleinrock 76] "Queueing Systems Volume 2" - L. Kleinrock, publisher John Wiley, 1976.

[Muth 60] "Optimal Properties of Exponentially Weighted Forecasts" - J.F. Muth, Journal of the American Stat. Assoc. Vol.55, 1960, pp.299-306.

[Postel 80] "DOD Standard Transmission Control Protocol" - J. Postel, IEN 129, Info. Sciences Institute, Univ. of S. Cal., January 1980.

[Rajaraman 78] "Routing in TYMNET" - A. Rajaraman, Eurocomp 78, Wembley, England, May 1978, publisher Online Ltd. Uxbridge, England, pp.9-21.

[Sunshine 75] "Interprocess Communication Protocols for Computer Networks" - C.A. Sunshine, Tech. Report 105 (Ph.D Thesis), Digital Systems Laboratory, Stanford University, December 1975.

[Wold 54] "A Study in the Analysis of Stationary Time Series - 2nd edition" - H. Wold, publisher Almqvist & Wiksell (Stockholm), 1954.

APPENDIX

Analysis of Non-Stationary Acknowledgement Delays

We consider the acknowledgement delays, t_n, to be defined as in Eq (14) by:

$$t_n = A_0 + h\, z_n + \sum_{j=0}^{n-1} z_{n-j} \,/\, g \qquad (g \geq 1, \quad n \geq 1)$$

where

$h = (g - 1) / g$

A_0 is an arbitrary positive number

and $\{z_1, z_2 \, .. \}$ is a sequence of independent random variables with mean 0 and variance U

The mean estimate, T_n, in Eq (6) then becomes:

$$T_n = \sum_{j=0}^{n-1} b^j \,[A_0 + h\, z_{n-j} + \sum_{k=0}^{n-j-1} z_{n-j-k} / g] / a + b^n T_0 \qquad (n \geq 1)$$

$$= A_0 + \sum_{j=0}^{n-1} [b^j h + \sum_{k=0}^{j} b^k / g]\, z_{n-j} / a + b^n (T_0 - A_0)$$

$$= A_0 + \sum_{j=0}^{n-1} [b^j h\,(1 - b) + (1 - b^{j+1})(1 - h)]\, z_{n-j} + b^n (T_0 - A_0)$$

$$= A_0 + \sum_{j=0}^{n-1} z_{n-j} / g + \sum_{j=0}^{n-1} b^j (h - b)\, z_{n-j} + b^n (T_0 - A_0)$$

giving

$$(t_{n+1} - T_n) = \begin{cases} z_1 - (T_0 - A_0) & (n = 0) \\ z_{n+1} - \sum_{j=0}^{n-1} b^j (h - b)\, z_{n-j} - b^n (T_0 - A_0) & (n \geq 1) \end{cases} \qquad (24)$$

and thus

$$E(t_{n+1}) = E(T_n) - b^n (T_0 - A_0) \qquad (n \geq 0) \qquad (25)$$

$$\rightarrow E(T_n) \qquad (n \rightarrow \infty) \qquad (26)$$

$$E^c(t_{n+1}) = T_n - \sum_{j=0}^{n-1} b^j (h - b)\, z_{n-j} - b^n (T_0 - A_0) \qquad (n \geq 1)$$

$$= T_n - b^n (T_0 - A_0) \qquad (b = h)$$

$$\rightarrow T_n \qquad (n \rightarrow \infty) \qquad (27)$$

Equations (26) and (27) lead to Eqs (16) and (18) in section 4.2. The condition $(b = h)$ in Eq (27) also minimises the mean square forecast error $E[(t_{n+1} - T_n)^2]$, since Eq (24) gives:

$$E(t_{n+1} - T_n)^2 = \begin{cases} U + (T_0 - A_0)^2 & (n = 0) \\ U + \sum_{j=0}^{n-1} b^{2j} (h - b)^2 U + b^{2n} (T_0 - A_0)^2 & (n \geq 1) \end{cases}$$

$$= U + (1 - b^{2n})(h - b)^2 U / (1 - b^2) + b^{2n} (T_0 - A_0)^2 \qquad (n \geq 0) \qquad (28)$$

$$\begin{cases} \rightarrow U & (b = h, \quad n \rightarrow \infty) \\ \rightarrow U + (h - b)^2 U / (1 - b^2) & (b \neq h, \quad n \rightarrow \infty) \end{cases}$$

To show that $E(V_n)$ converges to either $VAR(t_{n+1} - T_n)$ or $VAR^c(t_{n+1} - T_n)$ if $b = h$, we use Eqs (25), (28) and (24) to obtain:

$$VAR(t_{n+1} - T_n) = U + (1 - b^{2n})(h - b)^2 U / (1 - b^2) \qquad (n \geq 0)$$

$$= U \qquad (b = h)$$

$$= VAR^c(t_{n+1} - T_n) \qquad (29)$$

Taking the expectation of Eq (7) using Eq (28) with $(0 \leq b < 1)$ and $(0 \leq d < 1)$ then gives:

$$E(V_n) = \sum_{j=0}^{n-1} d^j \,[U + (1 - b^{2n-2j-2})(h - b)^2 U / (1 - b^2) + b^{2n-2j-2} (T_0 - A_0)^2] / c + d^n V_0 \qquad (n \geq 1)$$

$$= (1 - d^n)\,[U + (h - b)^2 U / (1 - b^2)] + d^n V_0 + \sum_{j=0}^{n-1} b^{2n-2} (d / b^2)^j \,[(T_0 - A_0)^2 - (h - b)^2 U / (1 - b^2)] / c$$

$$= (1 - d^n)\, VAR(t_{n+1} - T_n) + O(b^{2n}) + O(d^n) \pm \begin{cases} O(n\, b^{2n-2}) & (d = b) \\ O(b^{2n} - d^n) & (d \neq b) \end{cases}$$

$$\rightarrow VAR(t_{n+1} - T_n) \qquad (n \rightarrow \infty) \qquad (30)$$

Equation (30) corresponds to Eq (17), while combining Eqs (29) and (30) gives Eq (19).