

Further Explanations on Adaptive Retransmission Timers in TP4

Prepared by O. Rose and A. Herber

Background

On the 10th meeting of ATNP-WG2, a working paper was presented entitled "On Adaptive Retransmission Timers in OSI TP4" (WP 345). That paper presents simulation results and recommends the use of an adaptive retransmission timer T1 for achieving a better usage of the available throughput. A successive working paper, WP 362, provided comments to it.

To re-emphasize to main points of WP 345: While it is very difficult to configure a T1 timer value that operates satisfactory under all circumstances, it is possible to completely remove that burden from the network operator, making the timer self-adjusting to whatever environment it is running in.

To demonstrate this, a certain scenario has been studied, first with a manually configured timer, than with two increasingly sophisticated timer value adaptation algorithms applied. The algorithm proposed by Van Jacobsen and currently implemented in TCP is recommended to be used, since it has proven its suitability under a wide range of circumstances.

Purpose

The purpose of this paper is to clarify some points and to provide response to the statements made in WG2 WP 362. These further explanations concerning the simulations performed and the results achieved should be taken into account in the further work of WG2.

Comments to WP 362

Ad 1.0, Discussion:

The algorithm presented in WP 345 makes no assumptions about the network, thus it will operate over ground-ground as well as air-ground links. Since the algorithm only measures the time span between transmission of a packet and reception of the associated acknowledgement, there is no distinction (except from the size of the time span) between ground-ground and air-ground networks.

Van Jacobsen's algorithms are designed to operate at the transport layer, i.e. in end-systems. Thus, it is not clear what is meant with the statement "can be implemented on an individual, router-by-router basis".

The starting value for T1 may be a fixed value (if it is large enough for all circumstances). However, to avoid having to configure another value when preparing the protocol stack, it can as well be derived from the time it takes to open the connection.

Ad 2.0, Some Difficulties

The algorithm as presented in section 4, "Change Proposal for SARP's text" of working paper 345, is exactly taken from Van Jacobsen's original paper (and corresponds actually to the implementation of the TCP protocol). The only difference is that he used shorthands (e.g. A for average), while the text contains words to increase readability. We see no source of confusion arising from this.

The techniques described in subsequent appendices of Van Jacobsen's paper deal with congestion control, in particular congestion recovery. These aspects are covered by the congestion avoidance algorithm in ATN, which has proven to produce far better results than congestion recovery.

Ad 3.0, Recommendation

Regarding the first statement in WP362, it is explained that in section 4 of the "Change Proposal for SARP's text", a recommendation is proposed to implement the Van Jacobsen algorithm. Thus the algorithm in deed is described to be optional, not (as assumed in WP 362) mandatory.

Regarding the statement of missing references to the published literature: While it is fully agreed that credits should be given to the original author (as done in the working paper), removing the description of the algorithm from the SARP's text has its problems. It is not desirable that someone who has to implement the protocol must first search for a particular paper. Since the algorithm's description takes only a few lines, we prefer to find it in the SARP's text. Furthermore, we included important hints about how the computation shall be performed (e.g. ignoring samples from retransmitted packets, how to derive the initial value of T1, and the like), **that is not found in the original paper by Van Jacobsen**. Finally, the algorithm involves three parameters. While these may be fixed for the ATN, it is difficult to unambiguously specify recommended values for these parameters, if the algorithm is described in a completely different place.

Regarding the last statement in WP 362, it is agreed that further simulation studies may have to be performed (at least for the sake of completion) to demonstrate the effectiveness of these techniques when used over air-ground subnetworks.

Recommendation

The WG is invited to note the further explanations to WP and to take them into account in its further work program.