An Optimized Timer-based method for Feedback Control in Multicast Communication

Maziar Nekovee, Wen Haw Chong and Sverrir Olafsson
Complexity Research Group, BT Exact

We formulate a general method for finding the most optimal timer distribution function for stochastic feedback control in multicast systems by minimizing a user-defined cost functional, which incorporates both feedback implosion avoidance and latency constraints. We propose a novel shifted power-law timer and show through numerical analysis and simulations that it outperforms the previously proposed exponential timer (J. Nonnemacher and E.W. Biersack, IEE/ACM Trans. Networking, vol 7, pp 375-386, June 1999) function. We illustrate how our method can be used to design optimized timers for feedback control in system with specific latency constraints.

I. INTRODUCTION

IP multicast contrasts with more conventional unicast networking because data packets are only copied when necessary. Consider the scenario of a source computer broadcasting data to many receiver computers. Via unicast, multiple copies of the same data must be made and transmitted. This is not an efficient use of the network. Via multicast, copies of the data are made only when required. Only one copy of the data is sent along parts of the route that are the same for any subset of the receivers.

A key to scalability of multicast based systems is an effective mechanism for feedback control, in particular in order to alleviate feedback implosion, a situation in which the simultaneous feedback of receiver nodes overwhelms receiver's resources and cause congestion in the network [1]. Currently several solutions exist for alleviating feedback implosion, mainly in the context of reliable multicast [2,3]. Most of these solutions, however, are restricted to IP multicast over static networks and/or rely on a pre-existing network infrastructures for filtering and aggregation of feedbacks. The complexities of group communication technologies necessitate alternative strategies for alleviating the impact of feedback implosion in applications where, for various reasons, it is prohibitive to construct higher order delivery infrastructures on top of the basic multicast service, such as mobile and overlay multicast , or broad "fanout" communication is expected, for example in satellite communication.

Here we consider such protocols and focus on an end-to-end and generic feedback control mechanism based on stochastic timers. In their pioneering work Nonnenmacher and Biersack (NB) [4] provided a probabilistic model of this suppression mechanism , and showed that feedback implosion can be avoided using suitably chosen timer distribution functions (whose parameters could be dynamically adjusted to achieve optimal feedback suppression). In their approach NB were mainly concerned with reliable multicast to very large groups and primarily focused on minimizing the expected number of feedback messages. However, timer-based feedback control comes at a cost of increased delivery latency. In applications of multicast were timeliness is important it would be desirable to find an optimum trade-off between suppression and data delivery latency.
In this paper we develop a general approach to design of timer distribution functions for use with timer-based feedback control. Our approach allows optimizing the performance of timers for a general objective function, which can incorporate *both* feedback minimization requirement *and* latency constraints. Furthermore, we propose a novel shifted power-law timer distributions and show through numerical calculations and simulations for systems with up to a million receivers that it outperforms the truncated exponential timer which was proposed by Nonnenmacher and Biersack: It results in both a lower number of expected feedback messages (FBMs) and a lower latency for multicast groups with as many as one million receivers. We also provide closed-form expressions for the optimized parameters of these timers, which can be used in protocol design.

II. MATHEMATICAL MODEL

We consider the situation were a sender multicasts a request for feedback to a group of receivers. As the number of simultaneous feedback increases linearly with the number of receivers some kind of feedback control is required. Timer-based feedback control is based on the following algorithm [4]:

At each round of multicast the sender sends a request for feedback to all receivers together with a set of parameters specifying a timer probability distribution functions $f(t)$, $t \in [0,T]$, $i = 1,\ldots,R$ *, where R is the number of receivers. Upon receiving the request from the sender, a receiver* i *who detects a loss samples a waiting time* t_i *from* f *(we call this the backoff time) and waits. Upon expiry of this backoff time receiver* i *sends its feedback only if it hasn't already received a feedback message of any other receiver* j *, otherwise it suppresses its own feedback.*

By letting receivers sample their back-off time from an optimized probability distribution function the sender can dynamically control the number of feedback messages and achieve very efficient feedback suppression without any assistance from the network. Assume t_i is the times it takes a message to get from sender to receiver i and $t_{i,j}$ the time it takes a message to get from receiver *i* to receiver *j*. Receiver i will sends a feedback only if hasn't received feedback from any other receiver during the time $t_i + t_i$. Formally this condition corresponds to the following inequality

$$t_i + t_i < t_j + t_j + t_{ij} \qquad \text{for all } j \neq i . \qquad (1)$$

In the above equation all variables are stochastic. For the sake of simplicity we consider here the situation where the delay times between sender and the receivers and among receivers themselves are homogeneously distributed. We set delays equal to a constant c. The generalization to heterogeneous and/or stochastic delay times is straightforward and will be discussed elsewhere. The performance metrics which we consider are the expected number of simultaneous feedback messages $E[X]$ and the excess latency due to the feedback mechanism $E[M]$, which corresponds to the expected time for the expiry of the first timer. We consider a situation where all receivers in the multicast group, who all intend to send a feedback. Under the assumption of heterogeneous delays $E[X]$ and $E[M]$ can be written as functionals of the timer probability distribution functions $f(t)$ and the corresponding cumulative distribution function $F(t)$.

$$E[X] = RF(c) + \int_C^T dt\, f(t)(1 - F(t-c))^{R-1} \qquad (2)$$

$$E[M] = \int_0^T dt\,(1 - \mathrm{F(t)})^R$$

We note that in real applications it is important that not only the expected number of simultaneous feedback messages is kept low but also fluctuations around the average value are small. Fortunately, the total number of feedback messages has a binomial distribution and minimization of $E[X]$ simultaneously results in minimization of the fluctuations in the number of feedback messages.

We now define a general objective functional Ω of the form

$$\Omega = E[X] + w\,\Theta(E[M]) \qquad (3)$$

where w is a weight and $\Theta(E[M])$ is a suitably chosen function of excess latency. With the choice $w = 0$ minimization of Ω with respect to the parameters in the timer distribution corresponds to NACK minimization without considering the impact on excess latency while choosing $w > 0$ and $\Theta(E[M]) = E[M]$ corresponds to simultaneous minimization of the number of NACKs and excess latency. Finally by choosing $w >> 0$ and

$$\Theta(E[M]) = 1 - \frac{1}{1 + \exp(g\,(E[M) - E^0[M])} \qquad (4)$$

Feedback minimization under the latency constraint $E[M] \leq E^0[M]$ can be achieved. In the above equation g is a parameter, which can be used to adjust the "softness" of the latency constraint.

III. PERFORMANCE EVALUATION OF TIMER DISTRIBUTIONS: EXPONENTIAL VERSUS SHIFTED POWER LAW

For uniformly distributed timers the number of simultaneous feedbacks scale linearly with the number of receivers and the stochastic suppression mechanism is unable to alleviate feedback implosion for large groups. Nonnenmacher and Biersack [4] showed that feedback implosion could be avoided by choosing the timer distribution function such that it separates a small number of early feedback senders from the bulk of receivers. The feedback messages from these suppresses Feedback from all others receivers. These authors recommended a truncated exponential timer (TE) of the form

$$
f(t)=\begin{cases} \dfrac{1}{\exp l-1}\left(\dfrac{l}{T}\right)\exp\left(\dfrac{l}{T}t\right) & ;\ 0\le t\le T \\ 0 & ;\ otherwise \end{cases} \tag{5}
$$

In the above equations l is a parameter, which can be dynamically adjusted by the sender in order to minimize the impact of feedback implosion.

We propose a shifted power law (SPL) timer of the form

$$
f(t)=\begin{cases} \dfrac{b}{T}+(1-b)\dfrac{a}{T}\left(\dfrac{t}{T}\right)^{a-1} & ;\ 0\le t\le T \\ 0 & ;\ otherwise \end{cases} \tag{6}
$$

which is able to achieve a more optimal separation of early feedback senders from the bulk of receivers. In the above equation both a and b are adjustable parameters. Figure 1 display our results for the performance of both timers in terms of feedback suppression and latency (results are shown here only for T=5c. Results for other values of T will be included in the full paper). It can be seen that in all cases both optimized timers result in very efficient feedback suppression. However, the SPL timer outperforms the TE timer: it results in *both* a lower number of feedback messages and a lower latency. Note that the improvement becomes very significant as we consider smaller timer periods. This is an important improvement since excess latency increases with the timer period, and it is therefore desirable to choose T as small as possible[1]. With a view on applications in network, such as NORM [5], which make use of timer-based feedback control protocols, it is desirable to either have tabulated values of the optimized timer parameters for a range of receivers' numbers and timer periods, or closed-form expressions for the optimized parameters as functions of R and T. We used simple analytical arguments to derive such expressions for general timers and use our numerical data to demonstrate their usefulness. The resulting expressions are: (we found that there results for SPL timer are not very sensitive to the choice of exponent a, as long as its value is chosen, roughly, larger than 25).

$$
l^{opt}\approx\frac{T}{T-c}\log(R) \qquad \mathrm{b}^{opt}=\frac{T}{c}\frac{1}{R} \qquad (c<T), \tag{7}
$$

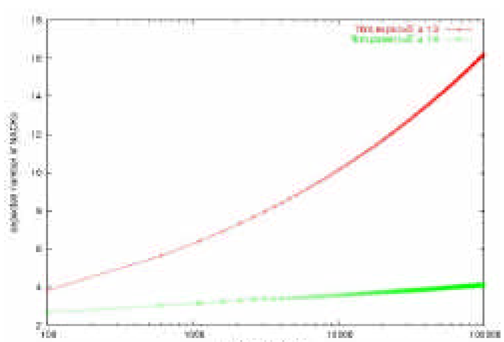

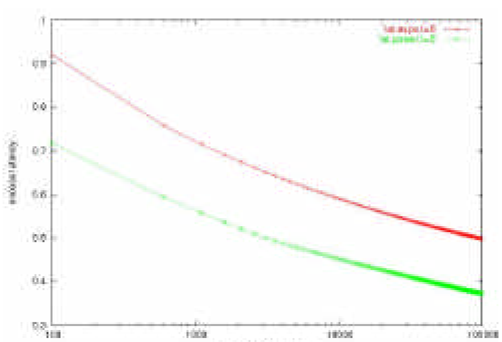

Fig.1 Number of expected feedback messages and excess latency as obtained from exponential (red lines) and shifted power law (green lines) timer distributions are shown for multicast groups with up to 10^6 receivers. Results are for timer periods $T=5$c.

VI FEEDBACK SUPPRESSION WITH LATENCY CONSTRAINTS

[1] For the exponential timer the latency increases linearly with T. In case of the shifted power-law timer it actually saturates.

Another situation that we consider is a scenario where suppression mechanism should operate under a user-specified latency constraint. Once again, our method can be used to optimize timers for this case. As an example, we show in figures 2 results for the case that a maximum latencies 0.25, 0.5, 1, 1.5 and 2.5 are imposed on the suppression mechanism (all latencies are in units of network delay c). We considered a system with up to

$R = 10^6$ Receivers and fixed $T = 10$, using optimized shifted power law timer family. It can be seen that the optimized timers achieve the required maximum latency while the number of expected feedbacks is kept very low.

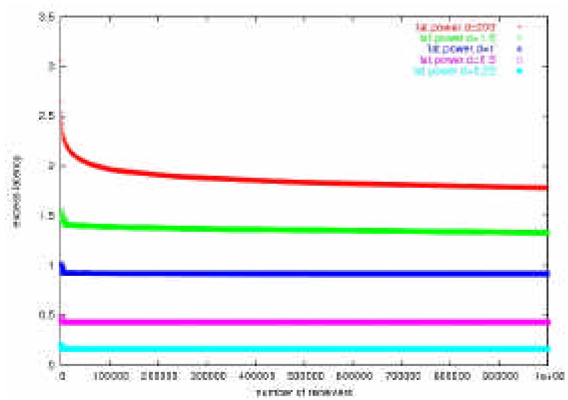

Fig. 2 Feedback suppression under latency constraint using optimised shifted power law timers, obtained from NACK minimization wi th the latency constraints E[M]< 0.25, 0.5,1, 1.5, 2.5, respectively. The corresponding number of NACKs remains below 8 for all latency constraints and all receiver numbers.

VII CONCLUSIONS AND OUTLOOK

We provided a general method for performance optimization of stochastic timers for feedback suppression in reliable multicast, which incorporates both low feedback rate and low latency requirements. We used our method to examine two families of timers. The exponential timer, which was previously suggested and a novel shifted power law timer.

We showed that the, when optimized, the modified power law timer outperforms the exponential timer, resulting both in a better feedback suppression and a lower excess latency. Timer-based feedback control comes at a cost of excess delivery latency. In some applications of reliable multicast it might be desirable to keep this latency below a certain threshold. We showed how this requirement could be met within our approach, without the need to adjust the timer period. Our future work will explore the application of the method for feedback control in reliable multicast transport scenarios in grid computing applications and, possibly, for mobile multicast.

REFERENCES

[1] See, e.g., M. Ammar, `Probabilistic Multicast: Generalizing the Multicast Paradigm to Improve Scalability'', Proceedings of IEEE INFOCOM 94, Toronto, Canada, 855, June 1994.

[2] S. Paul, K. Sabnani, J. C. Lin, and S. Bhattacharyya, `Reliable Multicast Transport Protocol (RMTP)'', IEEE Journal on Selected Areas in Communications, Vol. 15, 407, April 1997.

[3]. S. Floyd, V. Jacobson, C. Liu, S. McCanne, and L. Zhang, ``A Reliable Multicast Framework for Light-weight sessions and Application Level Framing'', IEEE/ACM Transactions on Networking, Vol 5, 784, Dec. 1997.

[4] J. Nonnemacher and E.W. Biersack, ``Scalable Feedback for Large Groups'', IEE/ACM Transactions on Networking, vol 7, 375, June 1999.

[5] R. B. Adamson, C. Bormann, M. Handey, J. Macker, ``NACK-Oriented Reliable Multicast (NORM) Protocol" IETF Internet-Draft, http://norm.pf.itd.nrl.navy.mil/draft-ietf-rmt-pi-norm-07.pdf (expires December 2003).