Scalable Feedback for Large Groups

Jörg Nonnenmacher
Bell Laboratories
Murray Hill, NJ 07974
nonnen@research.bell-labs.com

Ernst W. Biersack
Institut EURECOM
B.P. 193, 06904 Sophia Antipolis, FRANCE
erbi@eurecom.fr

To appear in IEEE/ACM Transactions on Networking 1999

Abstract

We investigate the scalability of feedback in multicast communication and propose a new method of probabilistic feedback based on exponentially distributed timers. By analysis and simulation for up to 10^6 receivers we show that feedback implosion is avoided while feedback latency is low. The mechanism is robust against the loss of feedback messages and works well in case of homogeneous and heterogeneous delays. We apply the feedback mechanism to reliable multicast and compare it to existing timer-based feedback schemes. Our mechanism achieves lower NAK latency for the same performance in terms of NAK suppression. No topological information of the network is used and data delivery is the only support required from the network. The mechanism adapts to a dynamic number of receivers and leads to a stable performance for implosion avoidance and feedback latency.

Keywords: Feedback, Multicast, Reliable Multicast, Performance Evaluation, Extreme Value Theory.

1 Introduction

Multicast communication is gaining in importance with the deployment of Multicast in the Internet and with the increasing number of satellites. A major challenge in multicast communication is the *feedback implosion* that occurs when a large number of receivers sends feedback to the sender.

In this paper we investigate feedback of groups of up to 10^6 receivers towards a single sender as needed for:

- *Reliable multicast*: Reliable multicast guarantees the delivery of data from the sender to every receiver. Feedback messages (FBMs) are needed in order to signal the loss (NAK), or the successful reception of data (ACK).
- *Estimation of the number of receivers*: is required to stop transmission, when no receivers are listening, to adapt scalable protocols to the number of receivers, e.g. by adjusting the amount of FEC [1], or to adjust the period of periodic control message emission.

The amount of potential feedback increases linearly with the number of receivers and may lead to a high traffic concentration at the sender, wasted bandwidth, and high processing requirements. Feedback implosion imposes high requirements to the mechanism for *feedback implosion* avoidance. Several solutions exist for implosion avoidance based on hierarchies, timers, tokens, and probing, see section 7 on related work.

Very little work [2, 3, 4] was done on the analysis of timer-based schemes for multicast feedback. We give an analytical foundation of timer-based feedback where the timer choice, the sender-receiver delays, and the delays between receivers can be modeled by arbitrary distributions. The analysis allows to compute:

- The expected number $E[X]$ of FBMs returned to the sender.
- The expected feedback delay $E[M]$ due to the timers.

We propose a new probabilistic feedback method for multicast based on exponentially distributed timers and show by analysis and simulation for up to 10^6 receivers that feedback implosion is avoided. We show the robustness of our mechanism to loss of FBMs, to homogeneous delays, and to heterogeneous delays.

We further evaluate our mechanism in the context of reliable multicast with respect to NAK implosion avoidance and to NAK latency. A comparison of our mechanism to existing timer-based feedback schemes shows that the feedback latency of our mechanism is lower for the same performance in NAK suppression.

Our mechanism requires very little state and has a low computational complexity at every receiver – independent of the group size. No knowledge about the network topology, nor support from the network is required to allow for implosion avoidance.

Using an estimate of the number R of receivers, our feedback mechanism allows to adjust the average number of FBMs returned to any value larger than one by trading off fewer FBMs for an increased feedback latency. Estimating the number of receivers is outside the scope of this paper; the interested reader is referred to [5, ch. 5], [6], and [7].

The paper is organized as follows. In section 2 we present an analysis for timer-based feedback schemes. In section 3 we evaluate the performance for reliable multicast feedback. Section 4 shows the robustness of timer-based feedback for loss and heterogeneous delays. The control of the amount of feedback is discussed in section 5. How to use the timer-based feedback scheme for networks providing only a unicast feedback channels is discussed in section 6. Section 7 discusses the work in the context of related work and section 8 concludes the work.

2 Timer-based Feedback

Consider the case where a sender needs to receive at least one FBM from R receivers and where the total number of FBMs returned should be small in order to avoid *feedback implosion*.

We consider a feedback mechanism with *feedback suppression*: A receiver that receives a FBM of another receiver will suppress its own feedback sending. FBMs are sent on a multicast feedback channel to be received at other receivers. If every receiver delays its multicast feedback sending by a random time, feedback implosion can be avoided. In section 6 the necessary modifications are given for the case where receivers return feedback via unicast.

Our timer-based feedback mechanism works as follows:

1. The sender multicasts a **request for feedback (I, λ, T)** to the R receivers. I is the identification for the feedback round, λ a parameter of the feedback algorithm, and T the interval size.

2. Receiver i receives the **request** (I, λ, T) after d_i time units and schedules a **exponentially distributed timer** z_i in the interval $[0, T]$. The parameter for the truncated exponential distribution is λ. When the timer z_i expires, receiver i:

   - sends the feedback message FBM(I, z_i) back to the sender if no other FBM(I, z_j) was received by i.
   - suppresses its feedback, if a FBM(I, z_j) of some other receiver j was received before (see figure 1 for an illustration of the suppression of i's feedback); this requires that j sends its feedback earlier than i and that the delay $d_{i,j}$ between receiver i and receiver j is such that:

   $$d_i + z_i > d_j + z_j + d_{i,j}$$

3. On the receipt of the FBMs, the sender computes an estimate $\hat{R}$ for the number of receivers, using the knowledge about the timer settings of all receivers i that returned feedback z_i, λ, T, see [5, ch. 5].

4. The sender computes T and λ for the next **request for feedback** based on $\hat{R}$ and its requirement for the feedback latency and the mean number of FBMs it wants to receive.

The SRM protocol [3] uses a similar mechanism for the sending of NAKs,with two major differences: First, SRM uses a **uniformly distributed timer choice** z_i from an interval that depends on the sender-receiver delay d_i. Second, SRM prevents loss of FBMs by scheduling a second request via an exponential back-off in a larger interval in the future.

In the following, we analyze the expected number $E[X]$ of FBMs returned to the sender from R receivers

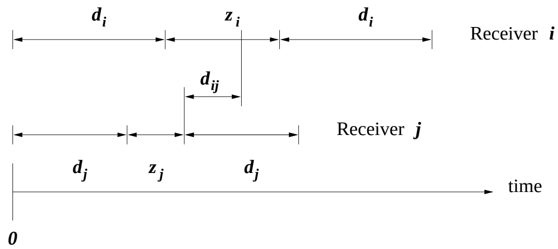

Figure 1: The timing for the feedback and the suppression of receiver i's FBM.

and the expected feedback latency $E[M]$ due to timers, when FBMs are not subject to loss. In section 4 we investigate the performance under loss of FBMs. First, we introduce the following random variables:

D_i	- one-way delay between the sender and receiver i. The delay paths are symmetric and D_i expresses also the one-way delay between receiver i and the sender.
Z_i	- time receiver i delays its feedback.
$V_i = D_i + Z_i$	- the time between the sending of the request for feedback and the time the timer expires at i.
$D_{i,j}$	- one–way delay between receiver i and receiver j. The delay paths are symmetric and $D_{i,j} = D_{j,i}$.
$W_{i,j} = V_j + D_{i,j}$	- time between the sending of the request for feedback and the reception of j's feedback at i.
X_i	- Bernoulli r.v., describes the number of FBMs from receiver i, either 0, or 1.
$X = \sum_{i=1}^{R} X_i$	- total number of FBMs received at the sender from the group of receivers.

The densities $f_{D_i}(d_i)$ and $f_{D_{i,j}}(d_{i,j})$ describe the delay d_i of receiver i to the sender and the delay $d_{i,j}$ between two receivers i, j. Different timer choices and timer choices dependent on the source-receiver delay d_i can be compared in their performance when the density for the timer choice is kept general:

$$f_{Z_i|D_i}(z_i|d_i) \tag{1}$$

Then, the density of $V_i = D_i + Z_i$ can be calculated by a transform changing variables [8, ch. 6.3], resulting in:

$$f_{V_i}(v_i) = \int_{-\infty}^{\infty} f_{D_i}(s_i) \cdot f_{Z_i|D_i}(v_i - s_i|s_i) ds_i \tag{2}$$

The same way the density of $W_{i,j} = D_{i,j} + V_i$ can be derived. Since $D_{i,j}$ and V_i are independent, the joint density is given by:

$$f_{D_{i,j},V_i}(d_{i,j}, v_i) = f_{D_{i,j}}(d_{i,j}) \cdot f_{V_i}(v_i)$$

Such that the density of $W_{i,j}$ using the transform in [8, ch. 6.3] is given by:

$$f_{W_{i,j}}(w_{i,j}) = \int_{-\infty}^{\infty} f_{D_{i,j},V_i}(s_{i,j}, w_{i,j} - s_{i,j}) ds_{i,j} \tag{3}$$

We assume delays D_i, and $D_{i,j}$ to be independent among receivers. For a single request for feedback, the Bernoulli random variable X_i describes whether the FBM from receiver i is sent ($X_i = 1$) or not ($X_i = 0$). Receiver i sends feedback only when no other receiver j suppresses the feedback of i. The probability for receiver i sending feedback is:

$$P(X_i = 1) = \int_{0}^{\infty} f_{V_i}(v_i) \prod_{j=1, i \neq j}^{R} (1 - F_{W_{i,j}}(v_i))\, dv_i \tag{4}$$

The analysis of the timer settings given above is valid for arbitrary delay distributions of D_i and $D_{i,j}$.

For a better understanding of the timer mechanism and the feedback suppression we will first consider the case where the *delays are homogeneous*: All receivers $i = 1, \ldots, R$ have the same delay $d_i = c$ from the sender and the same delay $d_{i,j} = c$ to any other receiver j:

$$f_{D_i}(d_i) = \delta(d_i - c) \qquad f_{D_{i,j}}(d_{i,j}) = \delta(d_{i,j} - c) \tag{5}$$

In section 4.2 we analyze the timer mechanism for heterogeneous delays.

We consider the case where *all* receivers $i = 1, \ldots, R$ choose a timer out of an interval $[0, T]$ – independent of the delay d_i between sender and receiver:

$$f_{Z_i|D_i}(z_i|d_i) = f_{Z_i}(z_i) \qquad , z_i \in [0, T] \tag{6}$$

We are especially interested in the minimal timer, which is the one expiring first. Let $M = \min_{i=1}^{R}\{Z_i\}$ be the random variable describing the minimal timer. Since the Z_i are identically and independently distributed, the distribution of the minimal timer is given by [9, ch 2]:

$$F_M(m) = P(M \leq m) \quad = \quad 1 - (1 - F_{Z_i}(m))^R$$

Our performance measures for evaluating the timer mechanisms are:

- **The expected feedback latency** $E[M]$ **due to the timer mechanism**, given by the minimal timer:

$$E[M] = \int_0^T (1 - F_M(m))\, dm \qquad (7)$$

- **The expected number** $E[X]$ **of FBMs** at the sender given as:

$$E[X] = \sum_{i=1}^{R} E(X_i) = RP(X_i = 1) \qquad (8)$$

Using these two performance measures, three different distributions for the timer choice are examined in terms of feedback suppression and feedback latency: The **uniform** distribution, the **beta distribution**, and the **exponential** distribution.

2.1 Uniformly Distributed Timers

A uniformly distributed timer choice out of the interval $[0, T]$ for receiver i is given by the density:

$$f_{Z_i}(z_i) = \begin{cases} \frac{1}{T} & , 0 \leq z_i \leq T \\ 0 & , otherwise \end{cases} \qquad (9)$$

The expected number $E[X]$ of FBMs is:

$$E[X] = \begin{cases} R & , c \geq T > 0 \\ 1 + \frac{c}{T}R - \left(\frac{c}{T}\right)^R & , 0 < c < T \end{cases} \qquad (10)$$

The expected feedback latency $E[M]$ due to the uniform distributed timer choice is:

$$E[M] = \frac{T}{R+1} \qquad (11)$$

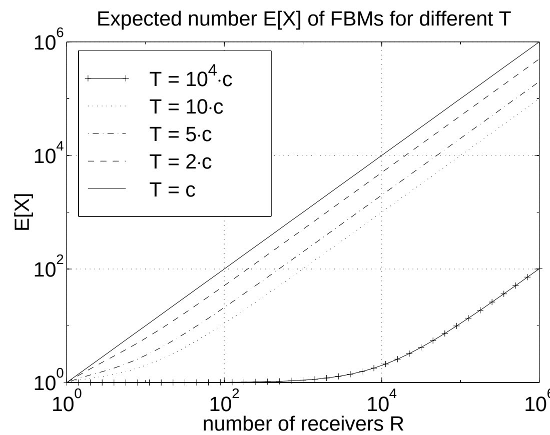

Figure 2: Expected number $E[X]$ of FBMs for **uniform distributed timer** choice from intervals of size $T = c, 2c, 5c, 10c, 10^4c$ for R receivers.

Let the interval size T be a multiple of the delay c between receivers. For a large number R of receivers, the expected number of FBMs is $E[X] \approx \frac{c}{T}R$ and thus increases linearly with the number of receivers, see figure 2. The feedback latency Eq. (11) on the other hand decreases with R, see figure 3. As already reported in [3], this means that there exists a tradeoff between suppression and latency. The approximation $\frac{c}{T}R$ for $E[X]$ and the feedback latency Eq. (11) show the occurrence of a reasonable tradeoff between the two considerations around $T = Rc$.

Figure 4 illustrates how suppression works: All receivers independently set their timers in the interval $[0, T]$. All k receivers that set their timer in the interval $[m, m+c]$ will send feedback. The other $R-k$ receivers with timers $z_i > m+c$ will suppress their feedback sending, since the FBM of the receiver with the minimum timer m reaches them before their timer expires.

For a uniform timer choice, the *only* way to adapt the feedback mechanism to the number R of receivers is to change the interval size T, which makes the performance of the scheme dependent on the accuracy of the receiver estimate:

- If the number R of receivers is overestimated, the interval size T will be chosen too large and a high feedback latency will be encountered.

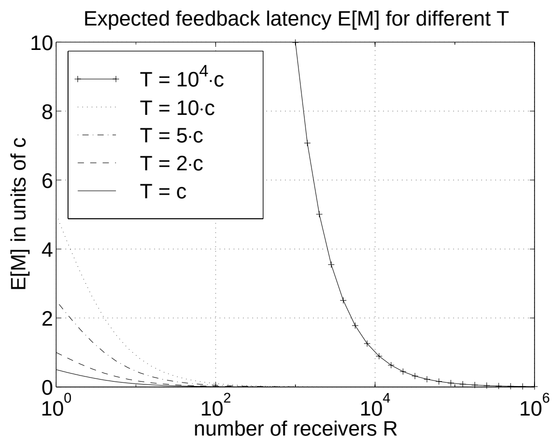

Figure 3: Expected feedback latency $E[M]$ for **uniform distributed timer** choice from intervals of size $T = c, 2c, 5c, 10c, 10^4 c$ for R receivers.

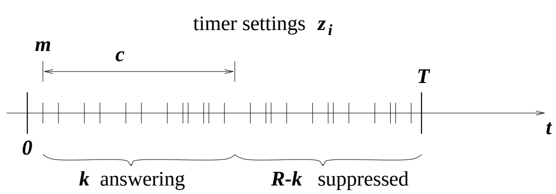

Figure 4: Timer Setting.

- If the number R of receivers is underestimated, the small interval size T will lead to a feedback implosion.

An alternative to the *uniform* distributed timer choice and to the difficulties arising from the need to carefully choose the interval size T is to change the shape of the distribution. Fixing the interval size gives a bound on the feedback delay. In order to also achieve a low number of FBMs, the minimal timer needs to be separated as far as possible from the mass of the timer settings. Therefore, the following properties are desirable for the density f_{Z_i} determining the timer choice:

- The minimal timer is separated from other timers by enabling a few timers to be set in a broad range and by grouping most timer settings in a small range.

- Feedback suppression is not sensitive to errors in the receiver estimate.

We investigate two other distributions f_{Z_i} for the timer choice: the **beta distribution** and the **exponential distribution**. Both distributions have parameters that allow us to change the shape of the distribution.

2.2 Beta Distributed Timers

The beta distribution [10] has two parameters a and b. For parameters $b = 1, a \geq 1$ a beta distributed timer choice on the interval $[0, T]$ is given by the density:

$$f_{Z_i}(z_i) = \begin{cases} \frac{a}{T}\left(\frac{z_i}{T}\right)^{a-1} & 0 \leq z_i \leq T, \\ 0, & otherwise \end{cases} \tag{12}$$

For $a = 1$ the beta distribution equals the uniform distribution. The weight of the density shifts towards T with an increasing a and results in a dense timer setting at high values.

The expected number $E[X]$ of FBMs for a beta distributed timer choice is:

$$\begin{aligned} E[X] &= R & , c \geq T > 0 \\ E[X] &= R\left(\frac{c}{T}\right)^a & , 0 < c < T \\ &+ Ra\int_{c/T}^{1} x^{a-1}\left(1-\left(x-\frac{c}{T}\right)^a\right)^{R-1} dx \end{aligned} \tag{13}$$

The feedback latency of Eq. (7) is given as:

$$E[M] = T\int_0^1 (1-m^a)^R\, dm \tag{14}$$

Figure 5 shows the suppression performance of the beta distribution with parameter $a = 10$ for different interval sizes $T = c, 2c, 5c, 10c$. First, we observe that suppression is achieved by beta distributed timers for a wide range of numbers of receivers R. Second, a moderate interval size $T = 10c$ is sufficient to keep the expected number of FBMs $E[X] < 15$ for up to 10^5 receivers. As a consequence, feedback suppression with beta distributed timer choice is, compared with uniform distributed timers, less sensitive to an error in the estimate of the number of receivers. Also the feedback latency of beta distributed timers, shown in figure 6, is relatively insensitive to an

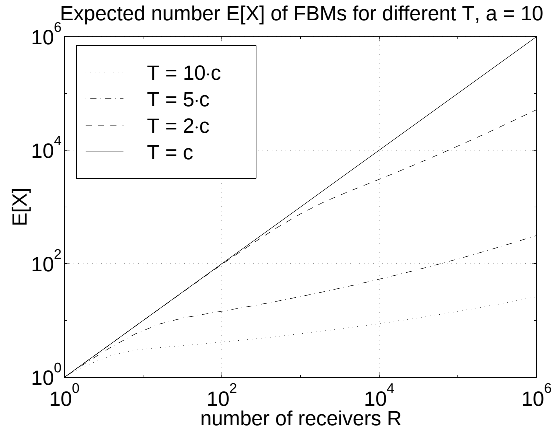

Figure 5: Expected number $E[X]$ of FBMs for **beta distributed timer** with parameter $a = 10$ from intervals of size $T = c, 2c, 5c, 10c$ for R receivers.

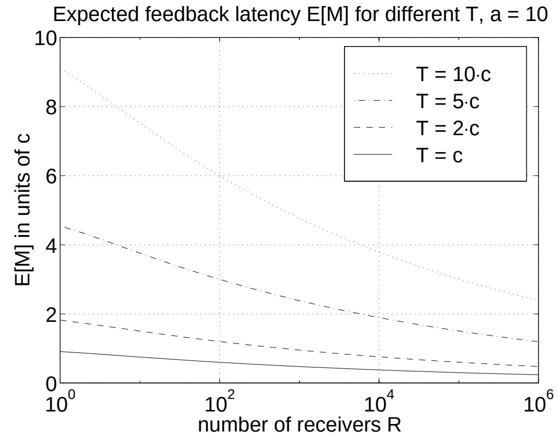

Figure 6: Expected feedback latency $E[M]$ for **beta distributed timer** with parameter $a = 10$ from intervals of size $T = c, 2c, 5c, 10c$ for R receivers.

error in the estimate of R: For $T = 10c$, the feedback latency varies only by $4c$ for the range from 100 to 10^6 receivers.

As in the case of uniformly distributed timers, a trade-off exists between the number $E[X]$ of FBMs and the feedback latency $E[M]$: the price to pay for good feedback suppression is an increase of the feedback latency.

Next, we study the performance of different beta distributions by varying parameter a. Figure 7 shows the impact of parameter a on suppression for $R = 10^4$ receivers, the corresponding feedback latency is shown in figure 8. We observe from figure 7 that the expected number $E[X]$ of FBMs is convex in a with a minimum at some a_o. For $a > a_o$ the number of FBMs is increasing with a, since the timer settings are forced on a narrow range close to T. The feedback latency $E[M]$ indicates that the minimal timer m also moves towards T with an increasing $a > a_o$. As a result, the timer settings of an increasing number of receivers fall in the interval $[m, m + c]$ and the number $E[X]$ of FBMs increases.

For $a < a_o$ the minimal timer is close to 0 and the other timers are not well separated from the minimal timer, resulting in feedback implosion.

We further observe from figure 7 that the minimal $E[X]$ at a_o does not depend on the interval size T, when T is large enough. Therefore optimal suppression is achieved by minimizing $E[X]$ for a given number R of receivers, not taking the interval size T into account. Once a_o is determined for optimal suppression, the interval size T can be used to trade-off feedback latency (Eq. (14)) against suppression (Eq. (13)).

In section 3 we will look at the question if a better suppression is achieved by beta distributed timers or by uniform distributed timers, given the feedback latency is the same in both cases. We now investigate the exponential distribution.

2.3 Exponentially Distributed Timers

The exponential distribution has one parameter λ and is defined from $-\infty$ to ∞. A truncated exponentially distributed timer choice in the interval $[0, T]$ is given by the density:

$$f_{Z_i}(z_i) \quad = \quad \begin{cases} \dfrac{1}{e^{\lambda} - 1} \cdot \dfrac{\lambda}{T} e^{\frac{\lambda}{T} z_i} & , 0 \le z_i \le T \\ 0, & , otherwise \end{cases} \tag{15}$$

As with the beta distribution, the weight of the density shifts towards T with an increasing λ and results in a dense timer setting at high values, see figure 9.

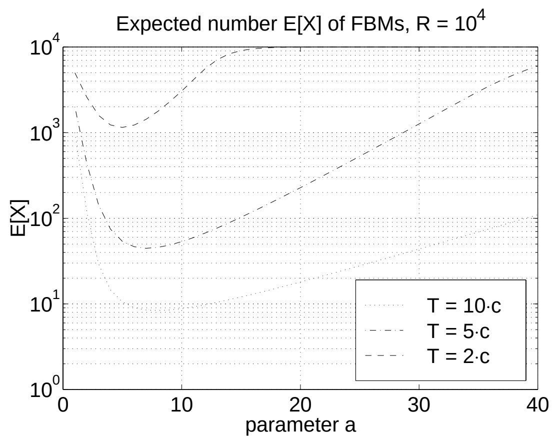

Figure 7: Expected number $E[X]$ of FBMs, dependence on parameter a for intervals of size $T = 2c, 5c, 10c$ for $R = 10^4$ receivers.

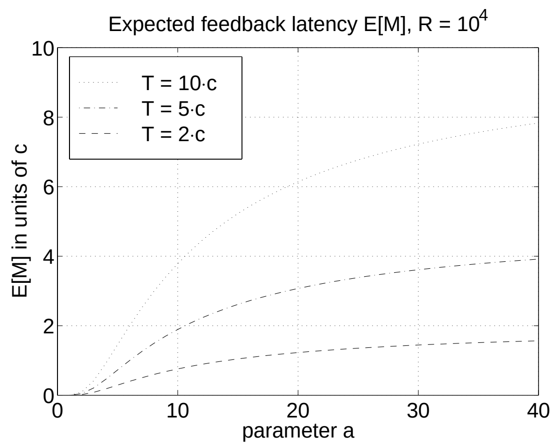

Figure 8: Expected feedback latency $E[M]$, dependent on parameter a from intervals of size $T = 2c, 5c, 10c$ for $R = 10^4$ receivers.

The expected number $E[X]$ of FBMs is:

$$
\begin{aligned}
E[X] &= R \qquad\qquad ,c \geq T > 0 \\
E[X] &= R\frac{e^{\lambda\frac{c}{T}} - 1}{e^{\lambda} - 1} \qquad ,0 < c < T \\
&- e^{\lambda\frac{c}{T}}\left(\left(\frac{1 - e^{-\lambda\frac{c}{T}}}{1 - e^{-\lambda}}\right)^{R} - 1\right)
\end{aligned} \tag{16}
$$

The feedback latency is:

$$
E[M] = T\int_0^1 \left(1 - \frac{e^{\lambda m} - 1}{e^{\lambda} - 1}\right)^{R} dm \tag{17}
$$

Figure 10 shows the suppression performance of an exponentially distributed timer choice with parameter $\lambda = 10$. We observe a *constant* suppression performance for a wide range of number of receivers. For an interval size $T = 10c$, suppression results in an expected number of FBMs $E[X] < 3.5$ for up to 10^4 receivers. Therefore, exponentially distributed timers outperform uniform and beta distributed timers: their suppression performance is less sensitive to a poor estimate of R. This can be seen by comparing figure 2 and figure 5.

For more than 10^4 receivers, $\lambda = 10$ is too small to separate the minimal timer from all other timers. The feedback latency shown in figure 11 goes to zero, and an increasing number of receivers fall in the interval $[m, m+c]$,

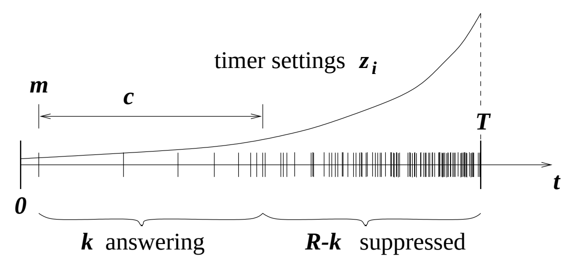

Figure 9: Timer Setting.

resulting in an increasing number of FBMs, as indicated in figure 10.

For the uniform and the beta distribution we observed a trade-off between suppression and feedback latency with the interval size T. This trade-off exists also for exponentially distributed timers, as shown in figures 10 and 11.

The impact of parameter λ on suppression is shown in figure 12. As for beta distributed timers, $E[X]$ is again a convex function with a minimum at some λ_o. We observe that the minimal number of FBMs with exponentially distributed timers is lower than the minimal number of FBMs with beta distributed timers for the same interval size T. This is seen by comparing figure 12 and figure 7.

As with beta distributed timers, the minimal $E[X]$ is nearly independent of the interval size T, if T is large enough (see figure 12). The feedback latency dependency

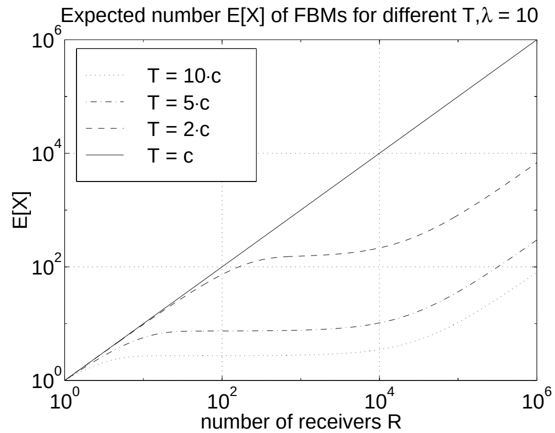

Figure 10: Expected number $E[X]$ of FBMs for **exponentially distributed timer** choice with parameter $\lambda = 10$ from intervals of size $T = c, 2c, 5c, 10c$ for R receivers.

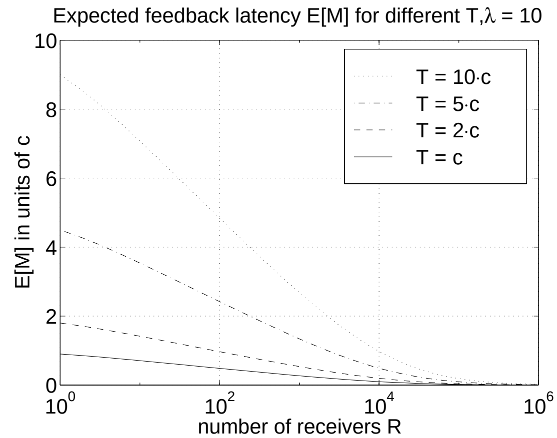

Figure 11: Expected feedback latency $E[M]$ for **exponentially distributed timer** choice with parameter $\lambda = 10$ from intervals of size $T = c, 2c, 5c, 10c$ for R receivers.

on λ, shown in figure 13 exhibits the same behavior: For different interval sizes, T, the feedback latency converges to 0 around the same λ. Therefore, λ_o for optimal suppression can be determined with the number of receivers, regardless of the interval size T. In section 5 the choice of the parameters λ and T is further investigated.

We can draw the following conclusions regarding feedback suppression for the three distributions evaluated:

- It is possible to avoid feedback implosion with probabilistic timers by a parametric distribution for the timer choice, while keeping the interval size T small. As a consequence, the feedback latency is small.
- The beta and exponential distribution are less sensitive to poor estimates of the number of receivers than is the uniform distribution:

  Dynamic changes in the number of receivers by orders of magnitude do not lead to feedback implosion and have only a minor effect on feedback latency with beta and exponential distributions.
- The parameter of the beta and exponential distribution can be adjusted for a desired suppression behavior in a tradeoff with feedback latency.
- Exponentially distributed timers outperform uniform and beta distributed timers for feedback suppression.

In the next section we evaluate the performance of the three timer schemes in the context of reliable multicast feedback and will take a close look on the tradeoff between latency and suppression for the three timer schemes.

3 Reliable Multicast Feedback

Different applications exist where feedback should be solicited fast from a subgroup of unknown size:

- A server selection process. From a large number R of servers only those being idle should respond to a request for a task assignment.
- Multicast flow control. From R receivers, only the R_l receivers that cannot keep up with the sending rate should respond.
- Access Control. A large number R of stations are connected to a medium that is limited in access. A monitor controls the access to the medium and polls all R stations for the interest in access. Only the subgroup of R_l stations wishing to access the medium responds.

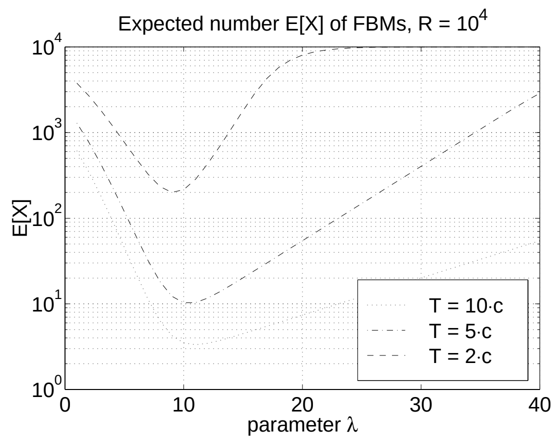

Figure 12: Expected number $E[X]$ of FBMs for **exponentially distributed timer** choice, dependent on parameter λ from intervals of size $T = 2c, 5c, 10c$ for $R = 10^4$ receivers.

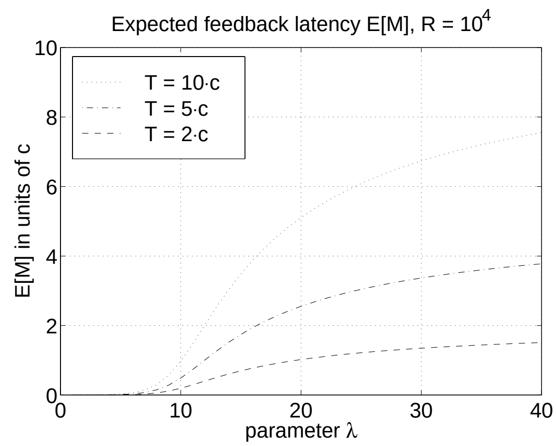

Figure 13: Expected feedback latency $E[M]$, dependent on parameter λ from intervals of size $T = 2c, 5c, 10c$ for $R = 10^4$ receivers.

We focus on reliable multicast feedback. In reliable multicast communication, negative acknowledgments (NAK) are shown to achieve a higher throughput performance than positive acknowledgments (ACK) [11]. Unfortunately, the meaning of an ACK is coupled with the receivers identity, and feedback suppression is therefore not possible for ACKs. NAKs on the other hand are redundant feedback and can be suppressed: a single NAK received by the sender is sufficient, given that the retransmissions are multicast.

The subgroup of receivers that are potential NAK senders depends on the loss of data packets. The subgroup consists of all receivers that detect a loss and subsequently want to send a NAK. Without a priori knowledge of loss, the number R_l of receivers in this subgroup is unknown and may vary from 0 to R. Feedback implosion must be avoided for the *worst case* where *all* R receivers want to send a NAK. Loss measurements [12] on the Internet have shown that this worst case is not unusual.

Let R_l be a fixed number of receivers out of all R receivers that lost data. In the following we evaluate feedback latency and choose R_l to be 1% of all R receivers, corresponding to a packet loss probability of $p = 10^{-2}$ and an average number of pR potential NAK senders out of R receivers.

We examine the timer distributions for:

- *NAK implosion* in the worst case: All R receivers are potential NAK senders.

- *NAK latency* in the average case: R_l receivers are potential NAK senders.

Note that the feedback latency increases with a decreasing R_l. For this reason we examine latency for the *average case*, where $R_l < R$ receivers are potential NAK senders.

For each distribution, we evaluate the tradeoff between the expected number $E[X]$ of NAKs in the *worst case* where R receivers want to send a NAK and the expected feedback latency $E[M_p]$ in the *average case* where only R_l receivers want to send a NAK.

For both cases, the same interval size T is used. For the uniform distribution, $(E[M_p], E[X])$ is uniquely determined by T. The exponential and beta distribution have another parameter λ or a. This parameter is adjusted to the worst case, where all R receivers are willing to send a NAK: $E[X]$ is minimized for a group of R receivers, and the corresponding λ_o or a_o is used to evaluate the tradeoff in T.

The **expected NAK latency** $E[M_p]$ is the feedback latency in the average case. It is obtained by substituting R

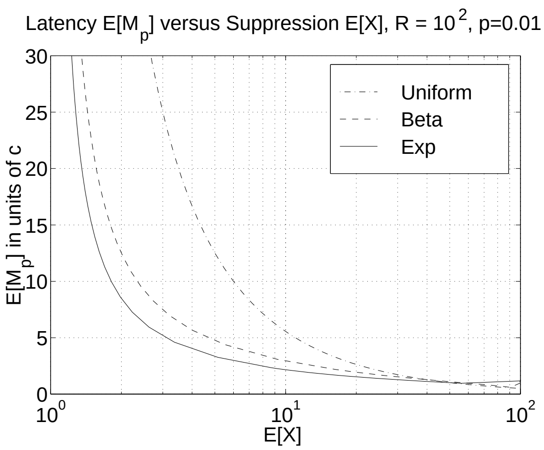

Figure 14: NAK latency $E[M_p]$ for optimal implosion avoidance with $R = 10^2$ receivers.

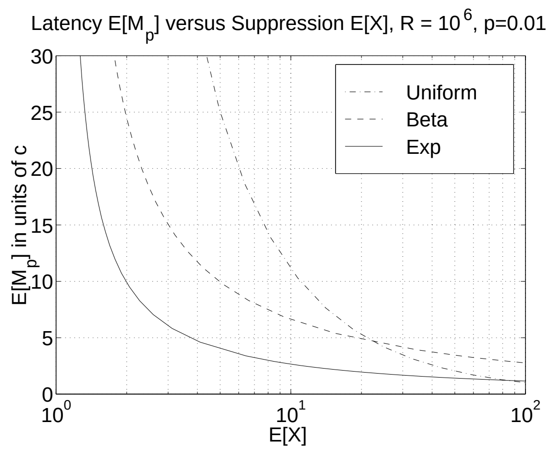

Figure 15: NAK latency $E[M_p]$ for optimal implosion avoidance with $R = 10^6$ receivers.

by R_l in $E[M]$. The expected number $E[X]$ of NAKs is given as before.

Figure 14 shows the expected NAK latency $E[M_p]$ versus the expected number $E[X]$ of NAKs for $R = 10^2$ receivers. This shows that on the average just one receiver will see a loss and send a NAK. The exponential distribution outperforms the other two distributions for up to $E[X] < 30$ NAKs in the worst case: For the same expected number $E[X]$ of NAKs in the worst case the first NAK of the average case is returned faster with the exponential distribution than with the other distributions. For a larger group of $R = 10^6$ receivers, the benefit of using the exponential distribution is even higher, compare figure 15.

Figure 15 shows that it is possible to adjust the exponential distribution for $R = 10^6$ receivers such that in the worst case an average of 4 NAKs are returned and in the average case, the first NAK is delayed by only 5 one-way delays c.

We adjusted the three timer distributions for the same performance in feedback suppression for the case where all R receivers want to send feedback and examined the feedback latency for the case where only a subgroup of $R_l < R$ receivers want to send feedback. Exponentially distributed timers result in faster feedback from the subgroup than with the other two timer distributions.

Due to the superior performance of exponentially distributed timers we will henceforth just consider those. In the following section we investigate the robustness of feedback suppression for exponentially distributed timers in case of loss and heterogeneous network delays.

4 Robustness of Exponentially Distributed Timers

4.1 Impact of Loss of FBMs

A lost FBM will not suppress the sending of FBMs by other receivers. While one might expect that loss of FBMs will result in feedback implosion, we show in the following that this is not the case.

We consider the worst case, where a FBM is lost directly at the feedback sender and is therefore not received by any of the other receivers. We simulated 100 feedback rounds and used parameters $\lambda = 10$ and $T = 10c$ in order to achieve simulation results that correspond to the former analytical results (see figure 10). FBMs were lost with different probabilities $p_{FBM} = 1\%, 10\%, 50\%$ and compared to the case of loss-free conditions. Figure 16 shows that the suppression performance of the timer mechanism is not sensitive to loss of FBMs for loss rates up to $p_{FBM} = 10\%$. We experienced a similar robustness

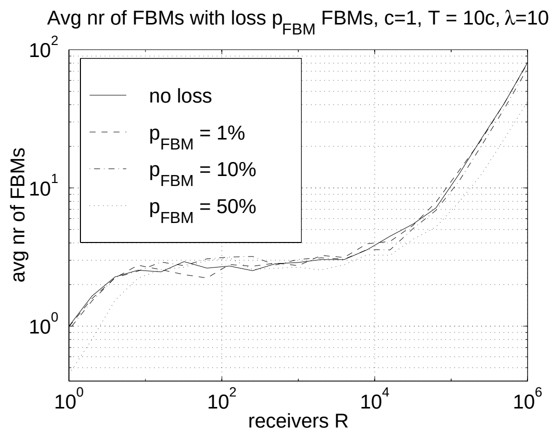

Figure 16: Average number of FBMs with loss p_{FBM}, $\lambda = 10, T = 10c$.

also for the average feedback delay. For the very high loss rate of $p_{FBM} = 50\%$, the average number of FBMs is decreased compared to loss free conditions and the average feedback latency is slightly increased. The reason for this behavior is twofold. First, the FBM due to the minimal timer m is lost with a probability of only p_{FBM}. Second, if the FBM due to the minimal timer is lost, the FBM of the next smallest timer $m' > m$ that is not lost jumps in and performs suppression. The number of timers expiring in $[m', m' + c]$ is higher than in $[m, m + c]$ due to the exponential distribution. However, feedback implosion does not happen since these un-suppressed FBMs themselves are subject to loss.

The investigated feedback mechanism results in the sending of a few FBMs. These un-suppressed FBMs constitute a natural redundancy useful in a lossy environment. The results show that feedback suppression using exponentially distributed timers is very robust with respect to the loss of FBMs.

4.2 Impact of Heterogeneous Delays

In a real network, receivers have different delays to the sender and different delays between each other. In order to understand the influence of heterogeneous delays on the timer mechanism, we examine the following two cases:

- Heterogeneous sender-receiver delays d_i, but homogeneous delays $d_{i,j} = c$ between receivers.
- Homogeneous sender-receiver delays $d_i = c$, but heterogeneous delays $d_{i,j}$ between receivers.

Both cases are compared to the case where the delays between sender and receivers and between receivers are homogeneous, i.e., $d_{i,j} = d_i = c$.

Heterogeneous delays d_i, or $d_{i,j}$ are in both cases beta distributed (see [10]) on the interval $[0, 2c]$ with parameters $a = 2$ and $b = 2$. This means that the average heterogeneous delay equals the homogeneous delay c (i.e. either $\bar{d}_i = c$, or $\bar{d}_{i,j} = c$). The given beta distribution models a realistic delay distribution. Wei showed [13] for different routing algorithms executed on random networks that the delay distribution follows roughly the beta distribution. Intuitively, this can be explained as follows: Starting from an origin in the network, the number of nodes reachable within a certain delay will increase as the delay increases. Since networks are limited in diameter, the number of nodes reachable within a certain delay D will, however, go to zero as D approaches the maximum delay from the origin to any node.

We simulated the FBM suppression by exponentially distributed timers with $\lambda = 10$ for this heterogeneous case for $R = 1, \ldots, 10^3$ receivers and used 95% confidence intervals. The interval size for the timer choice is $T = 10c$.

Heterogeneous delays to the sender

Let us consider the case where the delays between the sender and the receivers are heterogeneous and the delay between any pair of receivers i, j is homogeneous, $d_{i,j} = c$.

Figure 17 illustrates that FBM suppression performs better for small groups, $R < 10$, in the case of heterogeneous sender-receiver delays than for homogeneous sender-receiver delays. This is caused by a wider spread of timer settings over $[0, 2c+T]$ due to the heterogeneous reception times d_i of the request for feedback, instead of a more narrow setting in $[c, c + T]$ with homogeneous sender-receiver delays $d_i = c$.

As the group size R increases, FBM suppression does not increasingly benefit anymore from heterogeneous sender-receiver delays, since the impact of the number

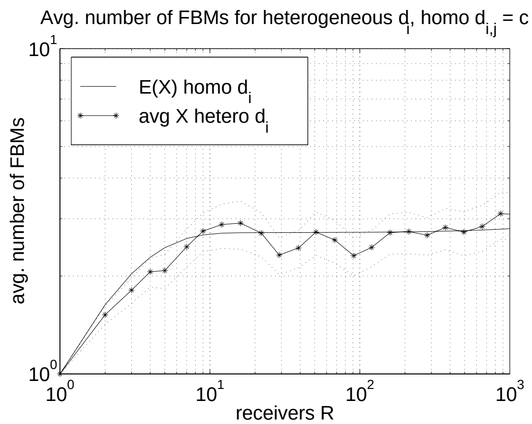

Figure 17: Expected number $E[X]$ of FBMs for heterogeneous sender-receiver delays $d_i \in (0, 2c), d_{i,j} = c$, interval size $T = 10c, \lambda = 10$.

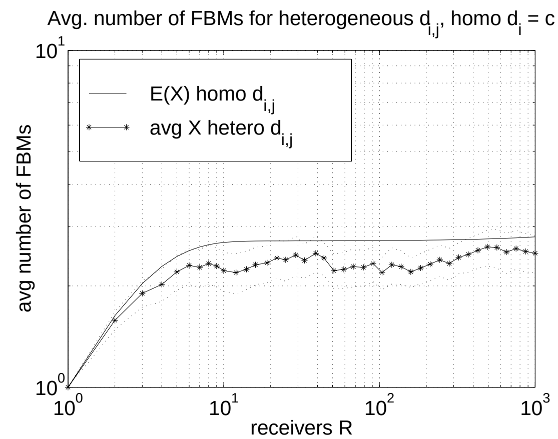

Figure 18: Expected number $E[X]$ of FBMs for heterogeneous inter-receiver delays $d_{i,j} \in (0, 2d)$, interval size $T = 10c, \lambda = 10$.

R of receivers on the density of the timer settings, and therefore on suppression, is higher than the difference in the delays.

Heterogeneous delays between receivers

Let us now consider a homogeneous sender-receiver delay, $d_i = c$, but heterogeneous delays $d_{i,j}$ between receivers, with $d_{i,j} \in [0, 2c]$. Therefore, the request for feedback is received at all receivers at the same time and all receivers set a timer in the interval $[0, T]$.

This is, for instance, the case for a forward channel via a satellite, where receivers are additionally connected among each other and to the sender via a terrestrial multicast feedback channel. The request for feedback is sent via the satellite (homogeneous $d_i = c$) while the delay $d_{i,j}$ between receivers via the terrestrial multicast feedback channel is heterogeneous.

Figure 18 shows that for all values of R, suppression benefits from heterogeneous delays between receivers. The reason is that not only does the minimal timer FBM perform suppression, but FBMs triggered by other small timers also perform suppression. For example, the FBM due to the 2nd smallest timer may suppress the feedback sending of the 3rd smallest timer. Heterogeneous delays between receivers therefore result in the suppression of FBMs that would have been sent in the homogeneous case.

From this section we can conclude that feedback suppression by **exponentially distributed timers** is:

- **not** sensitive to loss of feedback messages.
- **not** sensitive to heterogeneous delays between sender and receiver.
- **not** sensitive to heterogeneous delays between receivers.

Instead, these cases contribute to feedback suppression with probabilistic exponential timers and so lead to even better suppression performance.

5 Controlling the Feedback Bandwidth

Given limited network resources, the bandwidth available for feedback is limited. With the feedback mechanism from section 2, the feedback bandwidth is determined by the amount of feedback returned in the time between two successive feedback rounds. For a fixed FBM size of P bytes, the amount of feedback is given by $P \cdot X$, where X is the number of FBMs returned. Therefore, control

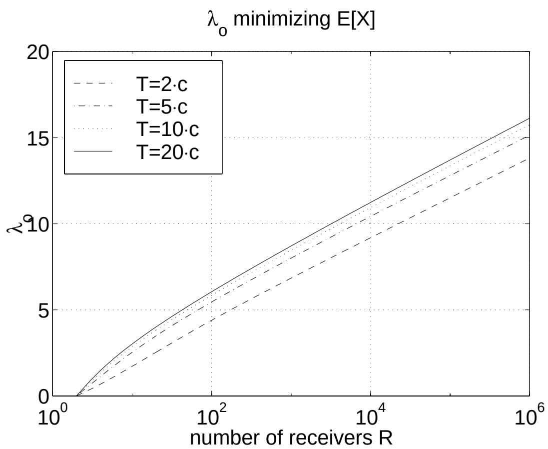

Figure 19: The λ_o minimizing the number $E[X]$ of FBMs dependent on R.

over the feedback bandwidth is provided, if the number of FBMs of all receivers can be adjusted.

In the following we consider a **desired number** N of feedback messages and show how the parameters λ and T can be tuned to obtain, on average, N feedback messages with low feedback latency. To keep the sender implementation simple, we give closed-form expressions for λ and T.

First, assume that the number R of receivers is known. For how to estimate the number of receivers see [5, ch. 5]. In section 2.3 it is shown that $E[X]$ is a convex function with a minimum at λ_o that is nearly independent of T. This allows us to determine a λ_o for optimal suppression - dependent only on the number of receivers: $R \mapsto \lambda_o$. Figure 19 shows the λ_o obtained for a given interval size, T, by minimization of $E[X]$ based on a golden section search and parabolic interpolation [14]. This is one way to adjust λ_o to the number of receivers. Another possibility is to approximate λ_o by a closed-form expression.

¿From figure 19 we observe that λ_o depends almost linearly on $\log_e(R)$. We further observe that the dependency of λ_o on the interval size T is minor. Taking this observations into account, λ_o is approximated by λ_o' for a given R:

$$\lambda_o' \sim a \cdot \log_e(R) + b$$

Parameters a and b are found by numerically fitting the polynomial $\lambda_o'(x) = a \cdot x + b$ to $\lambda_o(x)$ for $e^x = R = 10, \ldots, 10^6$ receivers.

	$T = 2c$	$T = 5c$	$T = 10c$	$T = 20c$
a	1.0383	1.0740	1.1000	1.185
b	-0.4214	0.4651	0.7326	0.8563

Table 1: Polynomial fitting of λ_o' to λ_o.

Table 1 shows the fitted parameters a and b for different interval sizes $T = 2c, 5c, 10c, 20c$. The value of a is stable between $1.0383 < a < 1.185$, while b deviates for a small interval size $T = 2c$ from the other values of b. Such small interval sizes do not allow for good suppression for most of the numbers of receivers used in the fitting process – with a, b for $T = 2c$ in case of 10^3 receivers already 62.3 FBMs are expected. Therefore, the deviation for small interval sizes is ignored and the parameters are chosen as $a = 1.1$ and $b = 0.8$. The adjustment of λ_o is then given by:

$$\lambda_o = 1.1 \cdot \log_e(R) + 0.8 \tag{18}$$

Given λ_o, the tradeoff between the expected number of FBMs Eq. (16) and the feedback latency Eq. (17) is determined solely by the interval size T. For increasing T, the expected number $E[X]$ of FBMs is decreases and the expected feedback latency due to timers increases linearly. Therefore, T is chosen as the smallest value for which $E[X] = N$, where N is the desired number of FBMs for R receivers. The expected number of FBMs can be approximated, since a large number R of receivers makes the following term converge to 0 for $T > c$:

$$\lim_{R \longrightarrow \infty} \left(\frac{1 - e^{-\lambda \frac{c}{T}}}{1 - e^{-\lambda}} \right)^R = 0$$

Thus, $E[X]$ is approximated by:

$$E(X) \approx R \frac{e^{\lambda \frac{c}{T}} - 1}{e^{\lambda} - 1} + e^{\lambda \frac{c}{T}} \tag{19}$$

If more FBMs are desired than there are receivers ($N \geq R$), the interval size is set to $T = 0$ and every receiver sends feedback immediately. If suppression is needed ($N < R$), λ_o is used and T set such that the minimum of $E[X]$ equals the desired number, N, of FBMs. By

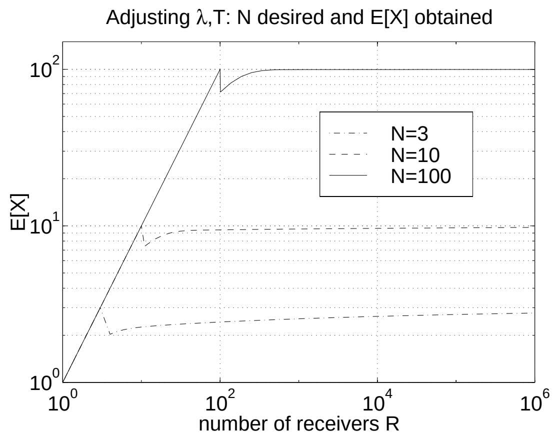

Figure 20: Error in adjustment of parameters λ and T to a desired mean number of FBMs $N = 3, 10, 100$.

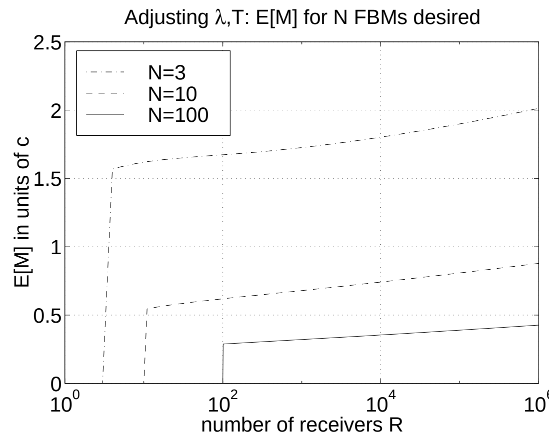

Figure 21: Feedback latency for the adjustment of parameters λ and T to a desired mean number of FBMs $N = 3, 10, 100$.

solving Eq. (19) for T we obtain the expression for the adjustment of the interval size T:

$$T = \begin{cases} 0 & , N \geq R \\ \dfrac{\lambda_o \cdot c}{\log_e\left(N + R\frac{1}{e^{\lambda_o}-1}\right) - \log_e\left(1 + R\frac{1}{e^{\lambda_o}-1}\right)} & \\ & , N < R \end{cases} \tag{20}$$

The error incurred by the approximation via Eqs. (18) and (20) is evaluated in the following. Figure 20 shows the expected number $E[X]$ of FBMs for $N = 3, 10, 100$ desired FBMs. We observe that the desired number N of FBMs is approached very fast. The discontinuity in the curves comes from the fact that for $N \geq R$ all receivers send immediately feedback. It can be observed that the adjustment of λ and T in the given fashion works well for widely differing N.

The corresponding feedback latency shown in figure 21 is low and does not vary significantly with the number of receivers. Even in the case where $N = 3$ FBMs are desired from $R = 10^6$ receivers ($999, 997$ suppressions) on average, the first FBM is delayed only for $2c$, which corresponds to one round trip time.

We gave closed-form expressions in Eqs. (18) and (20) for the adjustment of parameters λ and T to achieve a desired mean number N of FBM. Parameter λ is chosen such that the number of feedback messages is minimized for a given number of receivers. Parameter T is chosen such that the desired number N of FBMs equals this minimum. Due to the tradeoff between number of FBMs and feedback latency this adjustment yields low feedback latency.

Throughout this section we assumed that the number, R, of receivers is either known exactly, or that there exists an estimate $\hat{R}$ for the number of receivers. In the following we investigate the robustness of the parameter adjustment in case of an error in the receiver estimate.

Erroneous Receiver Estimate

The number of receivers might change, or the estimate of the number of receivers might be erroneous. We examine the danger of feedback implosion if the actual number R of receivers is different from the estimate $\hat{R}$. Parameters λ and T are adjusted via Eqs. (18) and (20) for $N = 10$ desired feedback messages and for estimates of the number of receivers $\hat{R} = 10^2, 10^3, 10^4$. From figure 22 we observe that the parameter adjustment results in the desired number of FBMs obtained just at the end of the flat segment of $E[X]$, right before the expected number $E[X]$ of FBMs starts slowly to increase when the actual num-

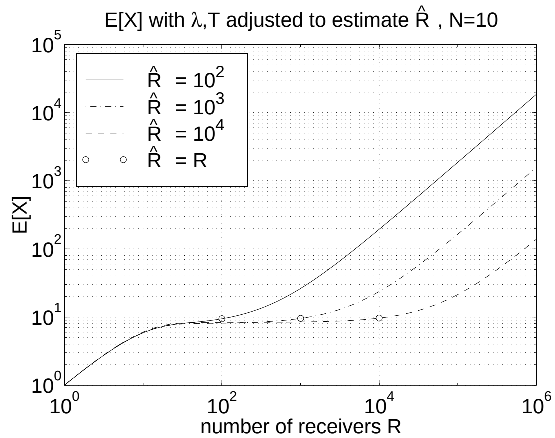

Figure 22: Impact of a wrong receiver estimate $\hat{R}$.

ber of receivers is higher than the estimate $\hat{R}$. First, at this point the feedback latency is low, compare figure 10 and figure 11. Second, we observe that the parameter adjustment is robust against a poor receiver estimate. If the real number of receivers is one order of magnitude higher than estimated, the number of FBMs only doubles or triples. If the real number of receivers is one order of magnitude lower than estimated, the number of FBMs stays roughly constant.

To assure that feedback implosion is avoided, we propose to adjust the parameters λ and T using a worst case receiver estimate $\hat{R}_{max} > \hat{R}$. If the receiver group is known to be stable, $\hat{R}_{max}$ can be chosen close to $\hat{R}$ to decrease feedback latency.

6 Unicast Feedback

Feedback suppression as introduced in section 2 requires a multicast feedback channel for *every* receiver. In this section we show how the same mechanism can be made to work in the presence of unicast feedback channels from the receivers back to the sender.

Unicast feedback has several advantages:

- The state in the routers is reduced for multicast routing algorithms, such as DVMRP [15], that keep state not per multicast group but per multicast sender. In such a case a separate multicast tree is built for every multicast sender, even if the senders belong to the same group. Receivers that multicast feedback *are* senders and the state in the network is therefore proportional to the number of receivers.
- Feedback suppression is possible for satellite networks, using a terrestrial unicast feedback channel.

Feedback suppression requires receivers to be aware of feedback sent by other receivers.

For unicast feedback, the missing multicast feedback channel is *emulated*. On the receipt of the first unicast feedback message, the sender multicasts information to all receivers to indicate that the feedback round is closed. On the reception of this message, receivers suppress feedback for this round.

For a multicast feedback channel, a natural robustness against FBM loss exists that assures feedback suppression, see section 4, since multiple FBMs are sent, each of which is able to suppress other feedback sending. To achieve a similar robustness to FBM loss for unicast feedback, the sender must indicate to the receivers the end of the feedback round several times. Several possibilities exist:

- The sender forwards every FBM received.
- The sender indicates several times, using the forward multicast channel, the end of feedback round I.
- The sender starts a new feedback round $I+1$. Receivers that have pending feedback for round $J < I+1$, then suppress this feedback.

The advantages of unicast feedback are offset by a larger feedback delay. This larger feedback delay, in turn, must be taken into account, when determining timer intervals. The round trip of the feedback via the sender results in a delay $d_{i,j}$ between two receivers i and j that is given by the sum of the symmetric delays d_i and d_j to the sender:

$$d_{i,j} = d_i + d_j$$

For unicast feedback and homogeneous delays $d_i = d_j = c$, the distance $d_{i,j}$ between receivers becomes $d_{i,j} = 2c$, as opposed to the case of multicast feedback, where $d_{i,j} = c$. The interval size T adjusted with Eq. (20) in proportion to the distance between receivers; therefore T also doubles. Since the feedback latency (Eq. (17)) is proportional to T, it will also double. The expected number $E[X]$ of FBMs in Eq. (16) will not change, since it is determined by the ratio of the distance between receivers and the interval size: c/T for the case of multicast feedback and $2c/2T = c/T$ for unicast feedback. As a consequence, the results from previous sections hold also for the case of unicast feedback, except that the expected feedback delay $E[M]$ due to timers will double.

7 Discussion and Related Work

Ammar has defined the feedback problem as response collection via several cost functions [16]. Most research on the feedback implosion problem has been driven by reliable multicast feedback.

Two major classes of feedback mechanisms exist that provide a solution to the *feedback implosion* problem:

- *Hierarchical approaches* [17, 18, 19, 20, 21]: Are an inherent solution to the *feedback implosion* problem and ensure a limited number of FBMs by accumulation/filtering in subgroups.

- *Approaches based on MAC protocols* [22, 23, 3, 4]: The feedback problem in multicast communication is related to the problem of Medium Access Control [24]: The multicast channel constitutes the shared medium and messages sent on the multicast channel are seen by every connected group member. A token mechanism as in token ring is proposed in [22] and random timers with exponential back-off as in CSMA/CD [25] are used in XTP [23] or the SRM protocol [3, 4].

Both classes of solutions are not without disadvantages: Hierarchical approaches require the setup of the hierarchy of subgroups and can not be employed in a scenario like satellite distribution with unicast backward channels. Approaches based on MAC protocols suffer from scalability problems. Tokens lead to high feedback latencies and random timers in [3, 4] are based on a uniform distribution. The analysis in [2] compares multicast feedback with random uniform timers to unicast feedback with respect to the cost in terms of number of control packets per link. The authors conclude that unicast control packets outperform multicast control packets for a small number of receivers.

SRM [3] exploits heterogeneous delays for a deterministic suppression, but needs a delay estimate $\hat{d}_i$ to the sender. This involves at least one packet sending from every receiver i, resulting for large groups of R receivers in a high amount of control traffic proportional to R. The optimal deterministic timers setting of Grossglauser [26] ensures only one NAK. However, the scheme requires the knowledge of the delay and network support for the timer setting.

Our mechanism does not suffer from any of these drawbacks, since it is a pure end-to-end mechanism. It does not rely on a full table of delay estimates to all receivers and its complexity is independent of the number of receivers. It does not need any network support except for data delivery and it does not need topological information. It can be employed in any kind of multicast capable network, also in networks where the feedback channel is only unicast.

Another end-to-end solution based on probabilistic feedback with exponential steps is the probing method of Bolot [6] that proceeds in discrete rounds. Using discrete rounds leads to very good performance for suppression but incurs a higher feedback latency than our scheme that uses a single round.

8 Conclusions

We investigated probabilistic feedback for multicast groups of up to 10^6 receivers by analysis and simulation. Our main results are:

- Exponentially distributed timer settings lead to a lower feedback latency and better feedback suppression than existing schemes based on uniform distributed timer settings.

- Probabilistic feedback with exponential timers is scalable with the number of receivers and avoids feedback implosion while assuring moderate feedback latency.

Based on these results we proposed a new timer-based feedback scheme that requires very little state, does not need any network support other than data delivery, and adapts to the number of receivers:

- It avoids feedback implosion and assures low feedback latency.

- It is robust under loss of feedback messages.
- It works for heterogeneous and homogeneous delays between multicast group members and can therefore be employed on nearly any kind of network including satellite-based networks.
- It allows to control the feedback bandwidth by adjusting the parameters dependent on the trade-off between average number of feedback messages returned and the latency for the feedback.
- It allows to estimate the number of receivers (see [5, ch. 5]).
- It is robust against an erroneous receiver estimate.
- It can operate on networks that only provide unicast feedback channels.

Acknowledgment

The research reported in this paper was carried out while the first author was with Institut Eurecom. The authors are grateful to Martin Lacher for simulating the mechanism with heterogeneous delays. We also thank Raymond Knopp for fruitful discussion. Eurecom's research is partially supported by its industrial partners: Ascom, Cegetel, France Telecom, Hitachi, IBM France, Motorola, Swisscom, Texas Instruments, and Thomson CSF.

References

[1] J. Nonnenmacher, E. W. Biersack, and D. Towsley, "Parity-based loss recovery for reliable multicast transmission," in *SIGCOMM '97*, Cannes, France, Sept. 1997, pp. 289–300.

[2] Sassan Pejhan, Mischa Schwartz, and Dimitris Anastassiou, "Error control using retransmission schemes in multicast transport protocols for real-time media," *IEEE/ACM Transactions on Networking*, vol. 4(3), pp. 413–427, June 1996.

[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, no. 6, pp. 784–803, Dec. 1997.

[4] P. Sharma, D. Estrin, S. Floyd, and V. Jacobson, "Scalable timers for soft state protocols," in *Proc. INFOCOMM 97*, Apr. 1997.

[5] Jörg Nonnenmacher, *Reliable Multicast to Large Groups*, Ph.D. thesis, EPFL, Lausanne, Switzerland, July 1998.

[6] J.C. Bolot, T. Turletti, and I. Wakeman, "Scalable Feedback Control for Multicast Video Distribution in the Internet," in *Proceedings of SigComm*. ACM, Sept. 1994.

[7] T. Friedman and D. Towsley, "Multicast session membership session estimation," in *INFOCOM 99*, Mar. 1999.

[8] Athanasios Papoulis, *Probability, Random Variables, and Stochastic Processes*, McGraw-Hill, Inc., 3rd edition, 1991.

[9] Enrique Castillo, *Extreme Value Theory in Engineering*, Academic Press, Inc., 1988.

[10] Paul E. Pfeiffer, *Probability for Applications*, Springer Verlag, New York, Berlin, Heidelberg, 1990.

[11] D. Towsley, J. Kurose, and S. Pingali, "A comparison of sender-initiated and receiver-initiated reliable multicast protocols," *IEEE Journal on Selected Areas in Communications*, vol. 15, no. 3, pp. 398–406, 1997.

[12] Maya Yajnik, Jim Kurose, and Don Towsley, "Packet loss correlation in the mbone multicast network," in *Proceedings of IEEE Global Internet*, London, UK, November 1996.

[13] Liming Wei and Deborah Estrin, "The trade-offs of multicast trees and algorithms," in *Proceedings of ICCCN'94*, San Francisco, CA, USA, Sept 1994.

[14] G. E. Forsythe, M. A. Malcolm, and C. B. Moler, *Computer Methods for Mathematical Computations*, Prentice-Hall, 1976.

[15] D. Waitzman, C. Partridge, and S. Deering, "Distance vector multicast routing protocol; rfc 1075," *Internet Request for Comments*, , no. 1075, nov 1988.

[16] Mostafa H. Ammar and George N. Rouskas, "On the performance of protocols for collecting responses over a multiple-access channel," in *Proceedings of INFOCOM'91*. IEEE, 1991, vol. 3.

[17] Rajendra Yavatkar, James Griffoen, and Madhu Sudan, "A reliable dissemination protocol for interactive collaborative applications," in *Proceedings of ACM Multimedia*, San Francisco, CA USA, 1995, ACM, pp. 333–344.

[18] S. Paul, K. K. Sabnani, J. C. Lin, and S. Bhattacharyya, "Reliable multicast transport protocol (rmtp)," *IEEE Journal on Selected Areas in Communications, special issue on Network Support for Multipoint Communication*, vol. 15, no. 3, pp. 407 – 421, April 1997.

[19] M. Hofmann, "A generic concept for large-scale multicast," in *Proc. International Zuerich Seminar*, B. Plattner, Ed. Feb. 1996, vol. 1044 of *LNCS*, pp. 95–106, Springer Verlag.

[20] J.M. Bonnin and J.J. Pansiot, "A scalable collect," in *HIPPARCH Workshop*, Sweden, June 1997.

[21] Dante DeLucia and Katia Obraczka, "Multicast feedback suppression using representatives," in *IEEE Infocom97*, 1997.

[22] J.M.Chang and N.F.Maxemchuk, "A broadcast protocol for broadcast networks," in *Proceedings of GLOBECOM*, Dec. 1993.

[23] Timothy W. Strayer, Bert J. Dempsey, and Alfred C. Weaver, *XTP - THE XPRESS TRANSFER PROTOCOL*, Addison-Wesley, 1992.

[24] Israel Cidon and Moshe Sidon, "Conflict multiplicity estimation and batch resolution algorithms," *IEEE Transactions on Information Theory*, vol. 34, no. 1, pp. 101–110, January 1988.

[25] Institute of Electrical and Electronics Engineers, "Carrier sense multiple access with collision detection (csma/cd) access method and physical layer specifications," *IEEE Std 802.3-1985*, 1985.

[26] Matthias Grossglauser, "Optimal deterministic timeouts for reliable scalable multicast," *IEEE Journal on Selected Area in Communications*, vol. 15, no. 3, pp. 422–433, April 1997.