Scalable Feedback for Satellite Broadcasts

Jörg Nonnenmacher, Ernst Biersack
Institut Eurecom
06904 Sophia-Antipolis Cedex France
{nonnen,erbi}@eurecom.fr

Abstract

We investigate the scalability of feedback for satellite broadcasts. We propose a new method of probabilistic feedback based on random exponentially distributed timers.

By analysis and simulation is shown that feedback implosion is avoided with good latency performance for up to 10^6 receivers. The mechanism is robust against the loss of feedback messages. We apply the feedback mechanism to reliable broadcast and compare it to existing timer-based feedback schemes. The results show that our mechanism leads to faster NAKs for the same performance in NAK suppression. Our mechanism is scalable, the amount of state at every group member is independent of the number of receivers. Our mechanism adapts to the number of receivers and leads therefore to a constant performance for implosion avoidance and feedback latency.

Keywords: Feedback, Broadcast, Multicast, Medium Access Control, Reliable Multicast, Performance Evaluation, Group Size Estimation, Extreme Value Theory

1 Introduction

Satellites allow the distribution of video, audio and data via broadcast. The scale of a satellite system and the high number of connected end systems (receive antennas) can lead for popular transmissions to a high number of receivers for the same broadcast.

In this paper we investigate feedback of groups from 1 up to 10^6 receivers to a single sender, as needed for:

- *Reliable Broadcast*: Reliable Broadcast 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 reception of data (ACK).
- *Estimation of the number of receivers*: Feedback is needed to estimate the number of receivers. The number of receivers allows to (i) advertise when the number of receivers is high (ii) stop transmission, when no receivers are listening (iii) adapt scalable protocols to the number of receivers, e.g. by adjusting the amount of FEC [12].

The major problem for broadcast feedback is the amount of feedback returned to the sender, resulting in a *feedback implosion*. Feedback implosion leads to a high traffic concentration at the sender, wasted bandwidth and high processing requirements. The amount of potential feedback increases linearly with the number of receivers and imposes high requirements to the mechanism for *feedback implosion* avoidance. Several solutions for implosion avoidance exist based on hierarchies, timers, tokens and probing (see section 6 on related work).

Very little work [7, 17] was done on the analysis of timer-based schemes for broadcast feedback. We give the analytical foundation of timer-based feedback, where the timer choice can be modeled by an arbitrary distribution. The analysis allows to compute:

- The expected number $E(X)$ of feedback messages returned to the sender.
- The expected feedback delay $E(M)$ due to the timers.

We propose a new probabilistic feedback method for broadcast based on exponentially distributed timers and show by analysis and simulation that feedback implosion is avoided for up to 10^6 receivers.

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

Our mechanism requires very few state and has a low computational complexity at every receiver – independent of the group size.

By an estimate of the number R of receivers, our feedback mechanism allows to adjust the average number of returned FBMs to any value > 1 via a tradeoff with feedback latency.

The remaining part of the paper is organized as follows. In section 2 the analysis for timer-based feedback schemes is given. In section 3 the performance is evaluated for reliable broadcast feedback. Section 4 shows the robustness of timer-based feedback for loss. Section 5 shows how the number R of receivers can be estimated due to the feedback. Section 6 discusses the work in the context of related work and section 7 concludes the work.

2 Timer-based Feedback

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

A broadcast channel allows to avoid feedback implosion when every receiver delays its feedback sending by a random time. A receiver that receives a FBM of another receiver can suppress its own feedback sending, referred to as *feedback suppression*. Feedback suppression is possible when all receivers are connected to the sender via a multicast feedback channel, but feedback suppression is also possible in the case where receivers return feedback via unicast as long as the sender broadcasts the information about the received feedback to all receivers.

Our timer-based feedback mechanism works as follows:

1. The sender broadcasts a **request for feedback** (I, λ, T) to the R receivers. I is the identification for the feedback round.

2. Receiver i, receives the **request** at time d_i and schedules a **random exponentially distributed timer** z_i in the interval $[0, T]$. The parameter for the truncated exponential distribution is λ.

3.
   - Receiver i sends the feedback message FBM(I, z_i) back to the sender, if its timer expires and no other FBM(I, z_j) was received yet.
   - Receiver i suppresses its feedback, if a FBM(I, z_j) of some other receiver j is received before its timer expires (see figure 1 for 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 small enough:

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

4. 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 section 5).

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

The SRM protocol [7] uses this mechanism for the sending of NAKs, with two differences: First, SRM uses a **uniform 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.

We consider the worst case for feedback implosion – the case, where the number of FBM sent is highest. This happens due to the timer settings in the first interval. After the exponential back-off, intervals are larger and the expected number of FBMs is smaller due to a sparser setting of timers and a higher number of suppressed FBMs. Our results will show the influence of the interval size on the number of FBMs.

In the following we analyze the expected number $E(X)$ of FBMs returned to the sender from R receivers and the expected feedback latency $E(M)$ due to timers, when FBMs are not subject to loss. In section 4 we remove the assumption of loss free conditions and investigate the performance

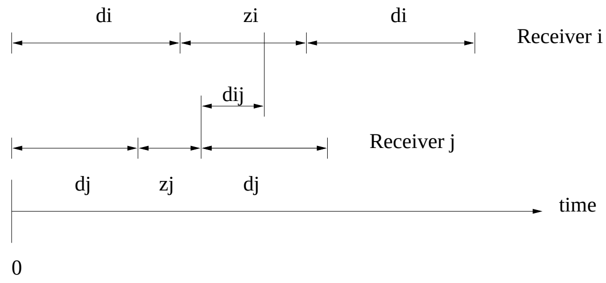

Figure 1: The timing for the feedback and the suppression of receiver i's FBM.

for loss of feedback messages. First we introduce the following random variables:

D_i	- the one–way delay between the sender and receiver i and vice versa.
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.
$D_{i,j}$	- the one–way delay between receiver i and receiver j and vice versa.
$W_{i,j} = V_j + D_{i,j}$	- the time between the sending of the request for feedback and the reception of j's feedback at i.
X_i	- Bernoulli, describes the number of feedback messages from receiver i.
X	- the total number of feedback messages received at the sender from the group of receivers.

Let the delay d_i of receiver i to the sender and the delay $d_{i,j}$ between two receivers i, j be given with the densities $f_{D_i}(d_i)$ and $f_{D_{i,j}}(d_{i,j})$. For a satellite communication the d_i and $d_{i,j}$ are homogeneous: All receivers $i = 1, \dots R$ have the same delay $d_i = c$ from the sender and the delay between any pair i, j of receivers $d_{i,j} = c$ is the same. For the case of homogeneous delays, the densities of D_i and $D_{i,j}$ are:

$$f_{D_i}(d_i) = \delta(d_i - c) \qquad f_{D_{i,j}}(d_{i,j}) = \delta(d_{i,j} - c) \quad (1)$$

Different timer choices can be compared on their performance, when the distribution F_{Z_i} for the timer choice is kept general.

We consider the case, where every receivers $i = 1, \dots R$ choose a timer out of an interval $[0, T]$. Let the timer delaying the potential feedback of receiver i by a time z_i be given by the density $f_{Z_i}(z_i)$ and the corresponding distribution:

$$F_{Z_i}(z_i) = \int_{-\infty}^{z_i} f_{Z_i}(x)dx \qquad , z_i \in [0, T] \qquad (2)$$

Then can the distribution of $V_i = D_i + Z_i$ be calculated by a change of variables, for details see [13, ch 6.3]. Since D_i and Z_i are independent is the joint density given by:

$$f_{D_i, Z_i}(d_i, z_i) = f_{D_i}(d_i) \cdot f_{Z_i}(z_i)$$

Such that the distribution of V_i using the transform in [13, ch 6.3] is given by:

$$f_{V_i}(v_i) = \int_{-\infty}^{\infty} f_{D_i}(s_i) \cdot f_{Z_i}(v_i - s_i) ds_i$$
$$F_{V_i}(v_i) = \int_{-\infty}^{v_i} f_{V_i}(x)\, dx \quad (3)$$

The same way the distribution of $W_{i,j} = D_{i,j} + V_i$ can be derived, resulting in:

$$f_{W_{i,j}}(w_{i,j}) = \int_{-\infty}^{\infty} f_{D_{i,j},V_i}(s_{i,j}, w_{i,j} - s_{i,j}) ds_{i,j}$$
$$F_{W_{i,j}}(w_{i,j}) = \int_{-\infty}^{w_{i,j}} f_{W_{i,j}}(x)\, dx \quad (4)$$

Since only the first timer setting is considered describes the Bernoulli random variable $X_i = 1$, if **the** feedback message from receiver i is sent or not.

Receiver i sends feedback, only when no other receiver j suppresses the feedback from i – sends feedback before that is received by i. This is expressed in the condition:

$$\forall j \in \{1, \ldots, R\} : j \neq i : v_i < w_{i,j}$$

Therefore the distribution of X_i is given as:

$$P(X_i = 0) = 1 - P(X_i = 1)$$
$$P(X_i = 1) = \int_0^{\infty} f_{V_i}(v_i) \prod_{j=1, i\neq j}^{R} P(W_{i,j} > v_i)\, dv_i$$
$$= \int_0^{\infty} f_{V_i}(v_i)(1 - F_{W_{i,j}}(v_i))^{R-1}\, dv_i \quad (5)$$

We are especially interested in the minimal timer, 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 is the distribution of the minimal timer given by [4, ch 2]:

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

Performance measures

Our performance measures of the timer mechanisms are:

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

$$E(M) = \int_0^T (1 - F_M(m))\, dm = \int_0^T (1 - F_{Z_i}(m))^R dm \quad (6)$$

- **The expected number** $E(X)$ **of FBMs** at the sender in total is with (5) given as:

$$E(X) = E(\sum_{i=1}^{R} X_i) = \sum_{i=1}^{R} E(X_i) = RP(X_i = 1) \quad (7)$$

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

2.1 Uniform Distributed Timers

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

$$f_{Z_i}(z_i) = \begin{cases} \frac{1}{T} & ,0 \leq z_i \leq T, \\ 0 & ,otherwise \end{cases} \quad (8)$$

Using homogeneous delays (1) and the uniform distributed timer choice (8) in (2) the probability $P(X_i = 1)$ that a receiver i sends a FBM (5) can be calculated via f_{V_i} and $f_{W_{i,j}}$ (3), (4), such that the expected number $E(X)$ of FBMs (7) is:

$$E(X) = \begin{cases} R & ,c \geq T > 0 \\ 1 + \frac{c}{T}R - \left(\frac{c}{T}\right)^R & ,0 < c < T \end{cases} \quad (9)$$

The expected feedback latency $E(M)$ (6) due to the uniform distributed timer choice is:

$$E(M) = \frac{T}{R+1} \quad (10)$$

The interval size T is given as a multiple of the delay c between receivers. It can be seen that all R receivers send feedback, for the case $T = c$, as given by equation (9): The delay c between receivers is too large to allow for a suppression – no feedback message can reach a receiver before its timer expires. As the interval size T increases ($T = 2c, 5c, 10c$), suppression increases. All receivers set independently a timer 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.

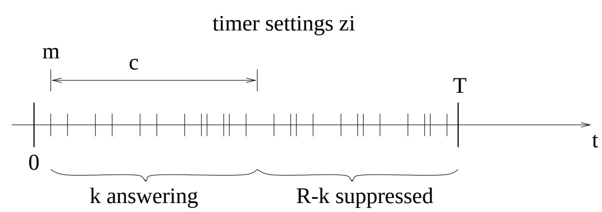

Figure 2: Timer Setting.

For large numbers $R > 10^2$ of receivers the expected number of FBMs is $E(X) \approx \frac{c}{T}R$ and thus increases linearly with the number of receivers. This indicates that feedback suppression based on a **uniform distributed timer choice** does not scale well with the number of receivers. The expected feedback latency $E(M)$ and therefore the minimal timer, decreases down to zero for an increasing number of receivers up to $R = 10^2$.

For the case, where suppression still works ($R < 10^2$ receivers) the tradeoff between suppression and latency is reported in [7] for **uniform distributed timer choice**.

The **only** way to improve suppression performance for **uniform distributed timer choice** is to increase the interval size T.

To provide good suppression also for groups up to 10^6 receivers a very large interval size T is required. The drawback of very large interval sizes T is that a high feedback latency may be encountered, due to a high variance of the minimal timer.

The importance of the minimal timer for probabilistic timer-based feedback can be taken into account also in a different manner than by changing the interval size T:

- Important for good suppression is to separate the minimal timer as far as possible from all other timers. Keeping the interval size T constant, the distribution for the timer choice can be used to change the minimal timer.

A desired distribution function f_{Z_i} separates the minimal timer from other timers by grouping most timer settings on a small range and by enabling some few timers to be set on a broad range.

A desired distribution function f_{Z_i} is scalable to the number R of receivers via a parameter. The number R of receivers and the distribution f_{Z_i} for the timer choice determine together (6) the expected minimal timer $E(M)$.

A desired distribution function performs suppression for the maximal number of potential feedback senders R, but also for a subset of of potential feedback senders $R_l < R$. This situation is encountered in the case of NAKs in reliable broadcast – the number of potential feedback senders R_l is determined by the loss of packets and no assumption can be made about R_l.

We investigate two other distributions f_{Z_i} for the timer choice that have these properties: the **beta distribution** and the **exponential distribution**. Both have a parameter that allows to change the distribution.

Given a priori knowledge about the number of receivers R, the sender announces with the *request for feedback* the parameter for the distribution and the interval size T in order to adjust the distribution f_{Z_i} for optimal suppression.

2.2 Beta Distributed Timers

The beta distribution has two parameters a and b, (see [15] for the complete definition). For parameter $b = 1$ is a beta distributed timer choice on the interval $[0, T]$ given by the density:

$$f_{Z_i}(z_i) = \begin{cases} \frac{a}{T}\left(\frac{z_i}{T}\right)^{a-1} & 0 \leq z_i \leq T, \\ 0, & otherwise \end{cases} \tag{11}$$

For $a = 1$ the beta distribution equals the uniform distribution. For $a > 1$ the weight of the beta density shifts to the right and results in a dense timer setting at high values.

The expected number $E(X)$ of FBMs for beta distributed timer choice can be derived as in the case of the uniform distributed timer choice, yielding:

$$E(X) = \begin{cases} R & ,c \geq T > 0 \\ R\left(\frac{c}{T}\right)^a + Ra\int_{c/T}^{1} x^{a-1}\left(1-\left(x-\frac{c}{T}\right)^a\right)^{R-1} dx & ,0 < c < T \end{cases} \tag{12}$$

The feedback latency is due to (6) given by:

$$E(M) = T\int_0^1 (1-m^a)^R \, dm \tag{13}$$

The stringent question is now if it is possible to achieve better suppression by using optimal beta distributed timer choice, or by a uniform distributed timer choice with very large intervals T, when the same feedback latency for both distribution should be achieved. Before answering this question in section 3 the exponential distribution will be investigated.

2.3 Exponentially Distributed Timers

The exponential distribution has one parameter λ and is defined from $-\infty$ to ∞. By ensuring that the cumulative distribution function F of the density is 1, the exponentially distributed timer choice on the interval $[0, T]$ is given by the density:

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

The weight of the density shifts to the right with an increasing λ, the same as encountered for the beta distribution with an increasing a.

The expected number $E(X)$ of FBMs for exponentially distributed timer choice can be derived in the same way as for a uniform distributed timer choice, yielding:

$$E(X) = \begin{cases} R & ,c \geq T > 0 \\ R\frac{e^{\lambda\frac{c}{T}} - 1}{e^\lambda - 1} - e^{\lambda\frac{c}{T}}\left(\left(\frac{1-e^{-\lambda\frac{c}{T}}}{1-e^{-\lambda}}\right)^R - 1\right) & ,0 < c < T \end{cases} \tag{14}$$

The feedback latency for R receivers is due to (6) given by:

$$E(M) = T\int_0^1 \left(1 - \frac{e^{\lambda m} - 1}{e^\lambda - 1}\right)^R dm \tag{15}$$

In the next section we evaluate the three timer schemes on their performance for reliable broadcast feedback and will jointly take a close look on the trade-off between NAK latency and NAK suppression of the three distribution functions.

3 Comparison for Reliable Broadcast

For feedback in reliable multicast negative acknowledgments (NAK) are shown to achieve much higher throughput performance than positive acknowledgments (ACK) [19], when retransmissions are multicast.

For the case of reliable multicast/broadcast feedback, where a FBM is a NAK, the number R_l of receivers that are potential feedback senders out of a group of R receivers depends on the loss of data packets.

Considering NAKs the *request for feedback*from the sender corresponds to the case where a packet is emitted by the sender and lost for all receivers. This happens for broadcast via satellite, where data is corrupted on the uplink and subsequently lost for all receivers. For NAK suppression we therefore need to consider this worst case, where all R receivers lose the packet and are all potential senders of a NAK.

Reliable broadcast also requires fast feedback. NAKs should be received fast, in order to perform a fast retransmission. The delay of a retransmission has impact on (i) the per-packet delay until it is received successfully, (ii) the throughput and (iii) the buffer requirements [5].

Our goal is to avoid NAK implosion first and second to optimize the NAK latency. NAK implosion needs to be avoided for the worst case, where **all** R receivers lose the packet. For NAK latency, however, the average case is more important. In the average case, $R_l < R$ receivers lose a packet. We consider this case for our second performance measure: Given independent packet loss with probability p at each of the R receivers, the average number of receivers losing the packet is $R_l = pR$.

The **expected NAK latency** $E(M_p)$ caused by timers is given as the expected feedback latency (6), where R is substituted by $R_l = pR$.

The expected feedback latency $E(M)$ given by (6) decreases with an increasing number R of receivers. Thus, the **expected NAK latency** is higher than the expected feedback latency of the worst case for NAK implosion, where R receivers are potential NAK senders.

For reliable broadcast we examine the three timer distributions for:

- The worst case for NAK implosion.
- The average case for NAK latency.

Each of the three distributions depends in different ways on the parameters R, T and λ, or a. In order to make a fair comparison are we looking at two different numbers R of receivers: $R = 10^2$ and $R = 10^6$. Then a broad range of interval distances T was defined.

For every T we calculated for each distribution the pair of performance measures (expected NAK latency, expected number of feedback messages) = $(E(M_p), E(X))$. The expected NAK latency $E(M_p)$ was calculated for a packet loss probability $p = 10^{-2}$.

For given R, T the expected number $E(X)$ is a convex function in a ((λ) for the beta distribution, (exponential distribution). For the beta and the exponential distribution we minimized $E(X)$ (12), (14) and used the corresponding a_o, λ_o for the calculation of the expected NAK latency. This means that the outcome of the minimization of the expected number of feedback messages, tuned for R receivers, is not optimal for the NAK feedback latency $E(M_p)$, calculated for $R_l = pR$ receivers.

In figure 3 and figure 4 the pair of performance measure is shown as the average NAK latency with respect to the worst case performance for NAK suppression. The tradeoff between good suppression and fast feedback is illustrated – the latency performance for the optimal suppression suffers for all distributions, when only very few feedback messages in the worst case are permitted (R receivers). The exponential distribution clearly outperforms beta and uniform distribution.

It can be seen that for the same performance for NAK suppression (minimal $E(X)$) the exponential distribution results in a lower feedback latency than the beta distribution and the uniform distribution for all expected numbers of returned FBMs.

The results show that it is possible to adjust the exponential distribution to an expected number of only 4 NAKs in the worst case for 10^6 receivers, with an expected NAK latency of 5 one-way delays= 2.5 round trip times due to the timers, experienced for the first NAK of only 1% of all receivers ($R_l = 10000$) that are potential NAK senders.

We showed that the exponential distribution performs best for the tradeoff between feedback suppression and feedback latency.

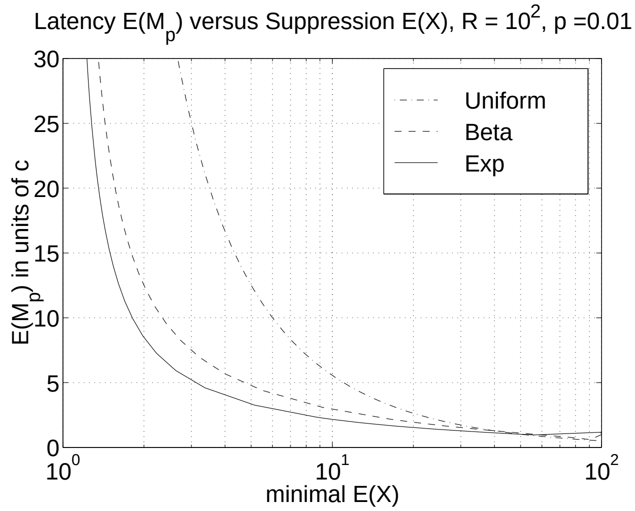

Figure 3: NAK latency for optimal implosion avoidance: $R = 10^2$ receivers.

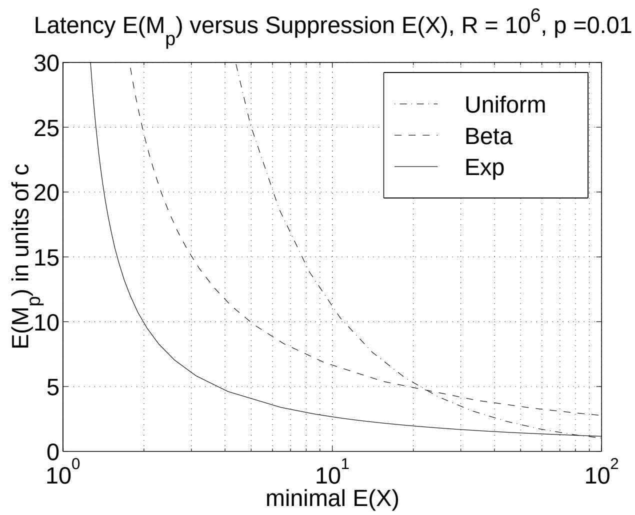

Figure 4: NAK latency for optimal implosion avoidance: $R = 10^6$ receivers.

In the following section the feedback mechanism based on the exponentially distributed timer choice is examined for its robustness to loss.

4 Robustness

Here we investigate the sensitivity of the exponentially distributed timer choice for its sensitivity to lost responses. The loss of responses may lead to:

- Feedback implosion, since a lost feedback message will not suppress other feedback sending.
- Increased feedback latency.
- A bad receiver estimate.

Again we consider the worst case, where a FBM is lost directly at the feedback sender and therefore lost for all other receivers. We simulated the feedback 100 times and used parameters $\lambda = 10$ and $T = 10c$. Feedback messages were lost with different probabilities $p_{FBM} = 1\%, 10\%, 50\%$ and compared to the case of loss free conditions. We experienced that the timer mechanism is not sensitive to loss of FBMs for loss rates up to $p_{FBM} = 10\%$.

For the very high loss rate of $p_{FBM} = 50\%$ the average number of FBMs is decreased compared to no loss and the average feedback latency is slightly increased.

We can conclude that feedback suppression by exponentially distributed timers is very robust with respect to the loss of feedback messages.

5 Estimating the group size R

Now the problem of estimating the group size R is investigated.

The group size can be estimated due to the distribution given for the feedback. From the parameters T, λ of the distribution given for the timer choice and the feedback received of the group for these parameters an estimate $\hat{R}$ on the group size can be given.

When a receiver i sends feedback, it includes in the feedback message:

- An identifier for the feedback round. This way the sender knows the interval size T and parameter λ associated with this feedback round.
- Its timer setting: $z_i \in [0, T]$.

On the reception of the feedback the sender can conclude on the number of concurrent feedback senders (receivers), via the setting of the minimal timer and the number of received feedback messages. In the interval $[m, m+c]$ from the minimal timer $m = z_i$ until the moment of suppression $m + c$, N timers expire. The information about the minimal timer m and the number N of responses can therefore be used to give an estimate $\hat{R}$ by evaluating the probability of N messages in the interval $[m, m+c]$ for the given parameters λ, T.

6 Discussion & Related Work

Ammar defined the feedback problem in a more general manner as response collection via several cost functions [1]. The most research on the the feedback implosion problem was driven by reliable multicast feedback.

There are two major classes of feedback mechanisms for multicast that are solutions to the *feedback implosion* problem:

- *Hierarchical approaches* [20, 14, 9, 3, 6]: Are an inherent solution to the *feedback implosion* and ensure a limited number of feedback messages by accumulation/filtering in subgroups. Representatives forward the feedback of a subgroup to the next hierarchy level, where the same process takes place, until the feedback is received at the sender.
- *Approaches based on MAC protocols* [11, 18, 7, 17]: The feedback problem in multicast communication is related to the problem of Medium Access Control: The broadcast channel constitutes the shared medium and messages sent on the broadcast channel are seen by every connected group member. This had been early recognized [11] and proposed solutions borrow mechanisms of medium access control. A token mechanism as in token ring is proposed in [11] and random timers with exponential back-off as in CSMA/CD [10] are used in XTP [18] or the SRM protocol [7, 17].

Both classes of solutions are not without disadvantages: Hierarchical approaches require the expensive 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 [7, 17] suffer up to now from state at every group member proportional to the group size.

Our mechanism does not suffer from any of these problems, since it is a pure end-to-end mechanism. It does not rely on delay estimates to other receivers and state and complexity are 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 broadcast network, the MBONE, as in satellite distribution. It works for unicast feedback channels or broadcast feedback channels, as long as the forward channel is broadcast – in the case of unicast feedback the sender forwards the information about the received feedback to the group of receivers for the purpose of suppression.

Another solution based on probabilistic feedback by probing with exponential steps is the probing method of Bolot [2], that proceeds in discrete rounds Using discrete rounds leads to very good performance for suppression, but to a higher feedback latency.

A related problem to the n:1 feedback is the one encountered in group communication, where regular periodic messages should be sent by every group member in order to exchange state. The problem are drastic changes of the group that may lead to long silence periods [16].

Other solution based on timers, other than the already mentioned ones, include a setting of optimal deterministic timers from Grossglauser [8], that ensures only one NAK based on the knowledge of the delay and on network support for the timer setting.

7 Conclusions

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

- Probabilistic feedback with exponential timers is scalable with the number of receivers and avoids feedback implosion up to 10^6 receivers for moderate feedback latency.

Based on this results we proposed a new timer based feedback scheme that requires very few state, that does not need any network support other than data delivery and that adapts to the number of receivers. Conclusions on the proposed feedback mechanisms are:

- It avoids feedback implosion and feedback is fast.
- The robustness to loss of feedback messages is shown.
- It allows to adjust the parameters dependent on the trade-off between average numbers of feedback messages returned and the latency for the feedback.
- It gives a good estimate of the number of receivers.

References

[1] 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*, volume 3. IEEE, 1991.

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

[3] J.M. Bonnin and J.J. Pansiot. A scalable collect. In *HIPPARCH Workshop*, Sweden, June 1997.

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

[5] Peter B. Danzig. Flow control for fast multicast. In *Proceedings of ACM SIGMETRICS*, pages 108–117, May 1989.

[6] Dante DeLucia and Katia Obraczka. Multicast feedback suppression using representatives. In *IEEE Infocom97*, 1997.

[7] S. Floyd, V. Jacobson, C. Liu, S. McCanne, and L. Zhang. A reliable multicast framework for light-weight sessions and application level framing. *Submitted to IEEE/ACM Transactions on Networking*, 1996.

[8] Matthias Grossglauser. Optimal deterministic timeouts for reliable scalable multicast. *IEEE Journal on Selected Area in Communications*, 15(3):422–433, April 1997.

[9] Markus Hofmann. A generic concept for large scale multicast. In *International Zürich Seminar*, Zürich, Switzerland, February 1996.

[10] 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.

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

[12] J. Nonnenmacher, E. W. Biersack, and Don Towsley. Parity-based loss recovery for reliable multicast transmission. In *To appear in SIGCOMM '97*, Cannes, France, September 1997.

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

[14] 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*, 15(3):407 – 421, April 1997.

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

[16] Jonathan Rosenberg and Henning Schulzrinne. Scaling feedback to very large groups. Technical report, Columbia University, New York, USA, February 1997. submitted to INFOCOM'98.

[17] P. Sharma, D. Estrin, S. Floyd, and V. Jacobson. Scalable timers for soft state protocols. In *Proc. INFOCOMM 97*, April 1997.

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

[19] 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*, 15(3):398–406, 1997.

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