

Reliable Multicast: Where to use FEC

Jörg Nonnenmacher¹ and Ernst W. Biersack²

Corporate Communications Department, Institut EURECOM
2229 route des Crêtes, B.P. 193, 06904, Sophia Antipolis Cedex, FRANCE

¹ Tel/Fax: +33 93002618/ 93002627, E-Mail: nonnen@eurecom.fr

² Tel/Fax: +33 93002611/ 93002627, E-Mail: erbi@eurecom.fr

Abstract

The MBONE provides an infrastructure for multicast communication on the internet based on IP. Several proposals have been made to create reliable multicast transport on top of the MBONE structure. Nearly all research on reliable multicast protocols for the MBONE focuses on ARQ error recovery. The counterpart of ARQ, FEC, does not guarantee 100% reliability but increases the reliability. The aim of this paper is to determine the best place in a multicast tree for the use of FEC. We develop a framework that allows us to model analytically the impact of FEC on the average number of transmissions necessary to transmit a packet to all members of the multicast group. We look on different multicast tree topologies, different degrees of correlated loss in the multicast tree, different multicast group sizes and show the effect of FEC in terms of transmissions needed to achieve 100% reliability. We find that the shared part of the multicast tree is not always the best part to employ FEC.

Keywords

Reliable Multicast, Forward Error Correction, Topology, Scalability

1 INTRODUCTION

Reliable multicast communication must handle the problem of transporting information from a sender to a number r of receivers via a network that is subject to packet loss. Two well-known methods exist to deal with loss: ARQ and FEC.

Adopting known ARQ schemes from point-to-point communication to $1 : r$ multicast communication is relatively straightforward and analytical results from early research [TM87], [Tow85] showed that ARQ with selective retransmission yields the best throughput performance.

After the first installation of the MBONE [Dee88] a few years ago, a growing part of the internet routers supports IP multicast. With the availability of the MBONE reliable multicast protocols have been implemented for the WAN environment and results have been reported in [RA95], [FJM⁺95], [YGS95]. The WAN environment provides a new challenge: reliable multicast communication had to face the scalability problem that arises for a large number

of receivers. The research community proposed two directions to deal with the scalability problem:

- *timer approaches* [Gro96], [FJM⁺95] that introduce a certain asynchrony among receivers to protect the source from a NAK-implosion and
- *cluster approaches* [YGS95], [LP96], [Hof96] that achieve the same goal by introducing a hierarchy into the multicast tree.

In this paper we are not following one of these ARQ approaches, but use Forward Error Correction (FEC) to decrease the packet loss experienced in the multicast tree. FEC adds a certain redundancy for the transport to the information resulting in a lower loss experienced by the receiver. FEC adds redundant information allowing a receiver to recover the original information even under partial loss. FEC in multicast connections [Car95] can have a big effect, since the global loss will increase with the number r of receivers. We evaluate the effect of the group size on the success of FEC.

We look at different WAN-LAN scenarios and quantify the effect achieved by FEC in terms of the average number of transmissions. We are also evaluating the effect that is achieved when FEC is applied either just on the LAN part or just on the WAN part of the tree and we want to know where FEC performs best. A counterpart for our research in FEC is [ZS91], where ARQ on a LAN-WAN-LAN environment is examined, with the conclusion that edge-to-edge error control performs very poorly and requires a lot of buffer. The advantage of FEC in an edge-to-edge case is that nearly no buffer is required and just a small encoding (decoding) delay is experienced. In [McA90] a VLSI coder and decoder for a burst erasure code is shown that can code at a speed of 1 Gigabit per second. The end-to-end semantic of the ARQ scheme on top of the FEC codec is not influenced. When we refer in the following to a LAN we mean a subnetwork connected to a WAN and refer to it also as a private network.

We demonstrate also that the effect of FEC on the multicast tree is strongly dependent on shared links. Nearly all the research on the performance of reliable multicast communication [TM87], [Tow85], [Den93], [ST87], [PTK94] assumes multicast trees with links where the loss over any link affects only a single receiver, referred to as independent loss. The assumption that the paths to different receivers have no shared links is typically not true for trees constructed by multicast routing algorithms as CBT [BFC93] or PIM [DEF⁺96]. We are evaluating FEC for different degrees of dependent loss by varying the number of shared links on the path from the source to the receivers.

Multicast routing algorithms establish the routes for a multicast group, that form a tree rooted at the sender. When we speak in the following of a tree, then we mean a multicast tree and vice versa.

The rest of the paper is organized as follows, first our performance measure is introduced and we define the topologies used to model the multicast trees. In the section Results we show the effect of FEC when applied on the whole multicast tree and we compare the cases, where FEC is just applied on the WAN part or just on the LAN part of the multicast tree. Finally we draw the conclusions.

2 PERFORMANCE MEASURE

We look at the number of transmissions that we need to transmit a single packet to **all** receivers. We believe that the expected number of transmissions captures the global packet loss in the tree and is the best way to evaluate the effect of FEC in a multicast tree.

We model FEC in a simple way by assuming that FEC will lower the link loss probability q by two orders of magnitude.* To evaluate the effect of FEC, we define the gain by employing FEC in terms of transmissions.

Let T be the average number of transmissions without FEC and T_{FEC} the average number of transmissions with FEC. Then, we can define:

$$gain_{FEC} = \frac{T}{T_{FEC}} \quad (1)$$

The gain performance measure depends on the number of receivers r , the link loss probability q and the topology of the multicast tree. To evaluate reliable multicast transmission, we calculate the expected number of transmissions needed to deliver a packet to **all** receivers. In [BMT94], the expected number of transmissions is given for loss at **nodes** in the multicast tree. We consider loss on a link, which is more appropriate due to two facts: loss at the source node is unlikely and link loss can be associated with loss in output buffers in routers. How to calculate the number of expected transmissions based on link loss and the multicast tree topology is given in the following.

Given that the packet is always successfully received by the predecessor (father) of node n , then let $T(n)$ be the number of transmissions of a packet until received by node n and all receivers in the subtree rooted at n and let $F_n(i) = P(T(n) \leq i)$ be the respective *CDF*. We assume a constant link loss probability q on the link leading to a node n and will calculate the expected number of transmissions $E(T(S))$ at the source S , using the *CDF* $F_S(i)$, as

$$E(T(S)) = \sum_{i=0}^{\infty} (1 - F_S(i)) \quad (2)$$

We can calculate $F_S(i)$ in a recursive fashion, starting at the leaves of the multicast tree. To obtain $F_n(i)$ for an arbitrary node n we must distinguish three cases.

- Node n is a leaf l , then the probability that fewer than $i + 1$ transmissions are needed to deliver the packet to n over the link leading to n is:

$$F_l(i) = P(T(l) \leq i) = 1 - q^i \quad (3)$$

- When n is an internal node, then exists one link leading to n and at least one child c . If there are i attempts to deliver the packet over the link leading to node n and it is lost exactly u times with the probability $q^u(1 - q)^{(i-u)}$, then a copy of the packet is forwarded $i - u$ times on every outgoing link to every child. The conditional probability that all

*Applying FEC end-to-end on a path with l links, instead of applying it on every link l does not change our analytical results, since for a small link loss probability q the approximation $1 - (1 - q)^l \approx ql$ holds.

children of n and the nodes in the subtrees rooted at the children are receiving the packet during these $i-u$ times is $\prod_{c \in \text{child}(n)} F_c(i-u)$. So we get $F_n(i)$ by summing over all possible u :

$$F_n(i) = \sum_{u=0}^{i-1} \binom{i}{u} q^u (1-q)^{(i-u)} \prod_{c \in \text{child}(n)} F_c(i-u) \quad (4)$$

- Node n is the source S , then there is no link leading to S and consequently only the loss experienced by its children c has to be considered:

$$F_S(i) = \prod_{c \in \text{child}(S)} F_c(i) \quad (5)$$

3 MODEL OF MULTICAST TREE AND LOSS

For our analysis, we define a generic multicast tree, referred to as tree 3 (see figure 1) that allows us to vary the essential parameters of a MC tree: link share, path length, number of receivers, and the link loss probability. In this multicast tree, we have r receivers, each at the same distance from the source. The number of links that each receiver shares with the other receivers is modifyable. The tree height is fixed to $h = 30$ and the number of unshared links on the path from the source to every receiver is $k = 1 \dots 30$, yielding $h - k = 0..29$ shared links. We have now two link loss probabilities: q_1 for a shared link and q_2 for the other links. This will help us to evaluate the loss in heterogeneous environments, where e.g. the shared links are in a WAN, and the other links are in a LAN. For tree 3 the loss probability on the path from the source to every single receiver is the same.

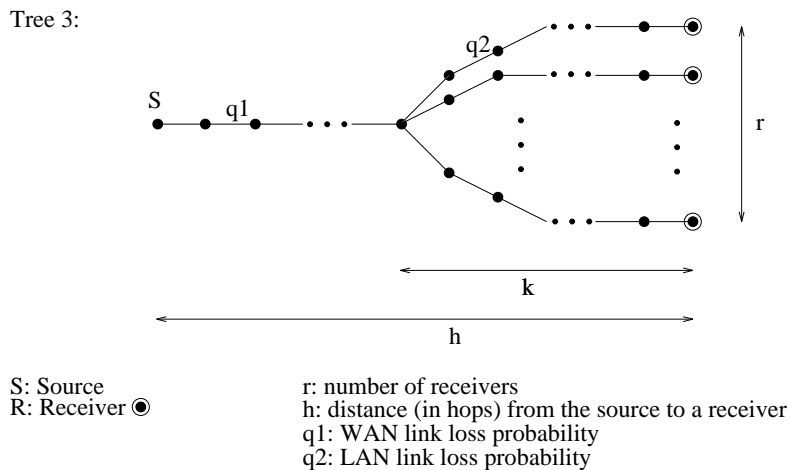


Figure 1: The multicast tree 3. Scalable in h , k and r .

The expected number of **transmissions** in Tree 3 can be evaluated according to (3), (4), (5) and (2). The calculation can be simplified by aggregating the $h - k$ shared links to one shared path sP with a packet loss probability of $q_{sP} = 1 - (1 - q_1)^{h-k}$. The k unshared links on the path to every receiver can be aggregated into one branch B with the loss probability $q_B = 1 - (1 - q_2)^k$, see figure 2.

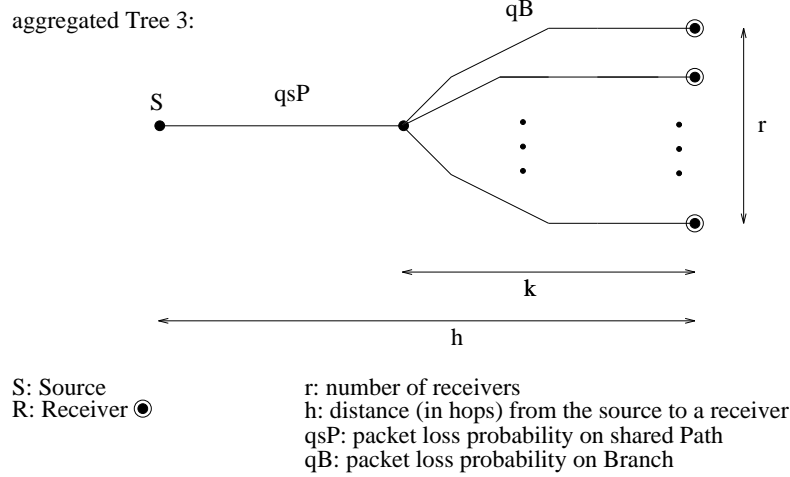


Figure 2: The aggregated multicast tree 3.

This aggregation yields:

For the single child c of the source using (4) we get:

$$F_c^{Tree3}(i) = \sum_{u=0}^{i-1} \binom{i}{u} (q_{sP})^u (1 - q_{sP})^{(i-u)} (1 - q_B^{i-u})^r \quad (6)$$

and for the source itself:

$$F_S^{Tree3}(i) = F_c^{Tree3}(i) \quad (7)$$

Having the expectation:

$$E_{Tree3}(T(S)) = \sum_{i=0}^{\infty} (1 - F_S^{Tree3}(i)) \quad (8)$$

Note that for a shared path of length $h - k = 0$, we have the case of r children at the source, which is different from the above case, described by (7) where we have just one child of the source. The above expressions for F_S^{Tree3} and F_c^{Tree3} are nevertheless valid and yield the same result in $E_{Tree3}(T(S))$ than the computation with r children at the source according

to (3), (4) and (5). This will be shown in the following. For $h - k = 0$ it is $q_{sP} = 0$ and $q_B = 1 - (1 - q_2)^h$, using the expressions (6) and (7) for one child at the source yields for F_S :

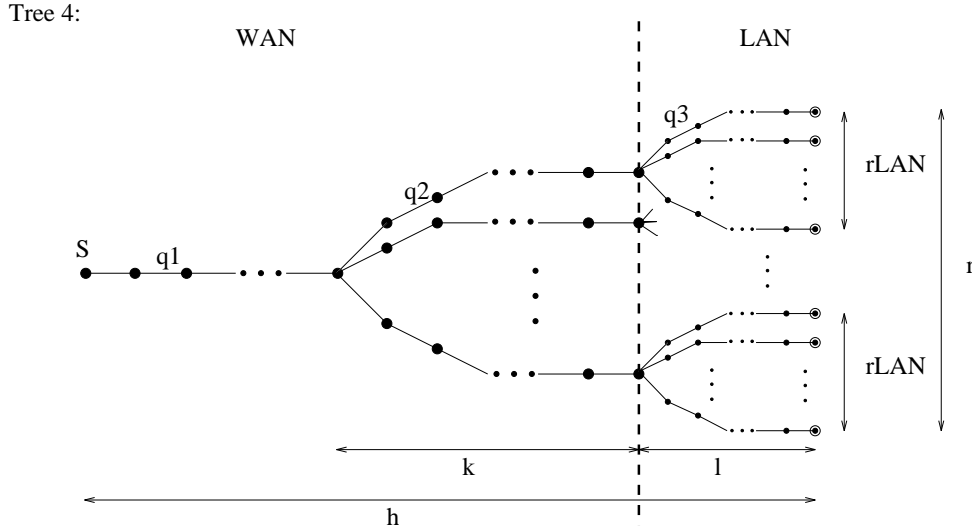
$$F_S^{Tree3}(i) = F_c^{Tree3}(i) = \sum_{u=0}^{i-1} \binom{i}{u} (0)^u (1)^{(i-u)} (1 - q_B^{i-u})^r = (1 - q_B^i)^r \quad (9)$$

This is then the same as having r children at the source, as expressed in:

$$F_S(i) = \prod_{c=1}^r F_c(i) = (F_c(i))^r = (1 - q_B^i)^r \quad (10)$$

The expected number of transmissions can therefore be calculated in any case with the expressions in (6) and (7).

A case to consider is big multicast groups with multiple receivers in several LANs. This case is modeled by multicast tree 4, see figure 3. Multicast tree 4 is based on tree 3 and has also a height of $h = 30$. Tree 3 models now the part of the Multicast tree, that resides in the WAN. The place of the receivers in tree 3 is now taken by LANs, introducing another hierarchy level. At the entry point into a LAN the tree splits another time into disjoint branches consisting out of $l = 5$ links and a low link loss probability q_3 . The part of the multicast tree in the WAN has therefore a height of $h - l = 25$. As in tree 3, there exists one common path of length $h - l - k = 25 - k$ hops of links with a link loss probability q_1 that splits into branches of length $k = 1 \dots 25$ consisting out of links with link loss probability q_2 , where $q_2 = q_1$. At the entry into the LAN each of these branches splits into $r_{LAN} = 5, 25, 100$ unshared branches of length $l = 5$ for every receiver. We will now have $r = 100 \dots 3000$ receivers in total, since we consider big multicast groups. The number of LANs is thereby given as $\frac{r}{r_{LAN}}$.



S: Source
 R: Receiver ●
 r: total number of receivers
 rLAN: number of receivers in a LAN

h: distance (in hops) from the source to a receiver
 k: length (in hops) of a WAN branch
 l: length (in hops) of a LAN branch
 q1: link loss probability on WAN shared path
 q2: link loss probability on WAN branch
 q3: link loss probability on LAN branch

Figure 3: The multicast tree 4. Scalable in h , k , l , r and r_{LAN} .

The expected number of transmissions in tree 4 is given in a similar way as the expected number of transmissions in tree 3. Let $q_{WANsP} = 1 - (1 - q_1)^{h-k-l}$ be the loss probability on the common path in the WAN, $q_{WANB} = 1 - (1 - q_2)^k$ the loss probability on a branch in the WAN and $q_{LANB} = 1 - (1 - q_3)^l$ the loss probability on the branch in the LAN, then the number of transmissions in tree 4 is given according to (3), (4), (5) and (2) by:

$$F_S^{Tree4}(i) = \sum_{u=0}^{i-1} \binom{i}{u} (q_{WANsP})^u (1 - q_{WANsP})^{(i-u)} \prod_{c=1}^{frac{rr}{LAN}} F_c^{WAN}(i-u) \quad (11)$$

$$F_c^{WAN}(i-u) = \sum_{j=0}^{i-u-1} \binom{i-u}{j} (q_{WANB})^j (1 - q_{WANB})^{(i-u-j)} (1 - q_{LANB}^{i-u-j})^{r_{LAN}} \quad (12)$$

4 RESULTS

Our analysis of the FEC gain consists of three parts.

The first part shows for small and for big multicast groups the effect of FEC when applied from the source to every receiver.

The second part shows the effect of FEC for a small multicast group and compares the FEC gain when FEC is applied either on the LAN part of the tree or when it is applied on the WAN part.

The last part shows the effect of partial FEC for big multicast groups with multiple receivers in several LANs.

4.1 FEC on the whole multicast tree

Using tree 3 we investigated the FEC gain for a different number of shared links and a different number of receivers as shown in figure 4 and figure 5. We assume a homogeneous link loss probability in tree 3 of $q = 10^{-3}$ without FEC and a link loss probability of $q = 10^{-5}$ with FEC is assumed.

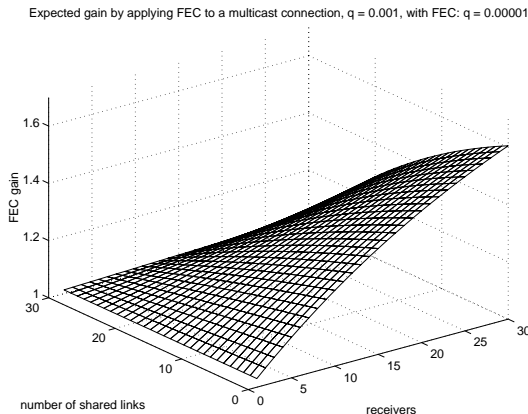


Figure 4: The gain of applying FEC to a multicast connection for a small multicast group and different degrees of link share.




Figure 5: The gain of applying FEC to a multicast connection for a big multicast group and different degrees of link share.

Figure 4 shows the FEC gain for a small multicast group. FEC yields an expected number of transmissions that is close to one, since there are few links in the multicast tree and the link loss probability $q = 10^{-5}$ is very low, compared to $q = 10^{-3}$ without FEC. The shape of the FEC gain in figure 4 is therefore caused by the expected number of transmissions without FEC. The reason for the increase of FEC gain with the number of receivers is that every new receiver causes k new links, and a packet is subject to loss on every link. The reason for the decrease of the FEC gain with the number of shared links is also due to the smaller number of links in the tree. Two major conclusions can be drawn: FEC is more efficient for a high number of receivers and FEC is more efficient for independent loss. In the case of independent packet loss (no shared links) and 30 receivers 60% more transmissions have to be made at average to deliver a packet to all receivers than in the case where FEC is used. Having a high number of shared links yields nearly no gain by applying FEC.

For a big multicast group (figure 5) the FEC gain is nearly constant and reduces the average number of transmissions by 50% independent of the group size and the number

of shared links. The reason for the constant FEC gain ($\frac{T}{T_{FEC}}$) is given by the number of transmissions with FEC (figure 6) and without FEC (figure 7).

Figure 6 and figure 7 show also that the typically used assumption of independent loss in analytical evaluations of reliable multicast communication is a worst case assumption, since the number of transmissions decreases with the number of shared links.




Figure 6: Expected number of transmissions for a big multicast group in tree 3 without FEC.

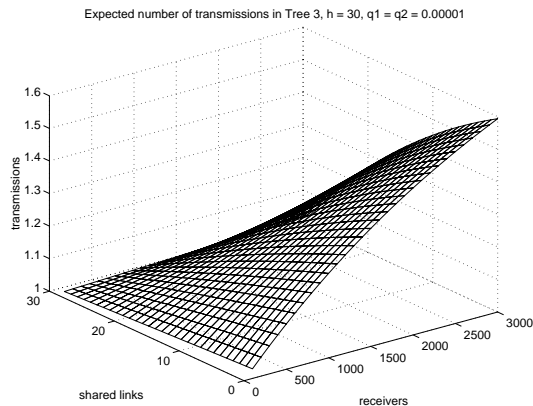


Figure 7: Expected number of transmissions for a big multicast group in tree 3 with FEC.

The number of links in the multicast tree is very high and the expected number of transmissions T_{FEC} in the case with FEC (figure 7) is despite a low link loss probability not close to one and influences the FEC gain ($\frac{T}{T_{FEC}}$ in figure 5) in contrast to the case for a small multicast group.

4.2 Small Multicast Groups: Partial FEC

A class of applications including multimedia conferences or distance teaching can profit from multicast communication. We model the multicast tree for this class by a tree where several participants are located in a private network and a path traversing a WAN connects the source with the receivers at the remote site.

We consider that the node where the MC tree makes the transition from the shared links to the dedicated links is the gateway between the world of the WAN (shared links) and the world of the LAN.

We will now use tree 3 to model a multicast tree that splits up at the entry into a private network (LAN). In the private network, we will have a limited number of receivers $r = 1 \dots 30$ and the number of links we have from the branch node to every receiver will also be relatively small: $k = 1 \dots 5$. Also we assume that the link loss probability in the WAN is higher ($q_1 = 10^{-3}$), than in the private network ($q_2 = 10^{-4}$).

In the previous section we looked on the impact of FEC when applied from the source to the receivers of the multicast tree. However, FEC can also be applied selectively, covering

a part of the multicast tree. We consider the two cases where FEC is applied just on the shared links, or just on the independent links close to the receivers. Again we assume that FEC lowers the loss probability by two orders of magnitude. We will compare the last two cases in terms of gain of transmissions. Figure 8 shows that protecting the shared path in the WAN is more efficient in this case. The gain achieved by FEC is in both cases very low (under 1.03). For FEC in the WAN this is due to the low number $h - k = 25 \dots 29$ of links that are affected by FEC, for FEC in the LAN more links are affected, but the already low link loss probability does not lead to a big FEC gain, since the number of transmissions without FEC is already close to 1.

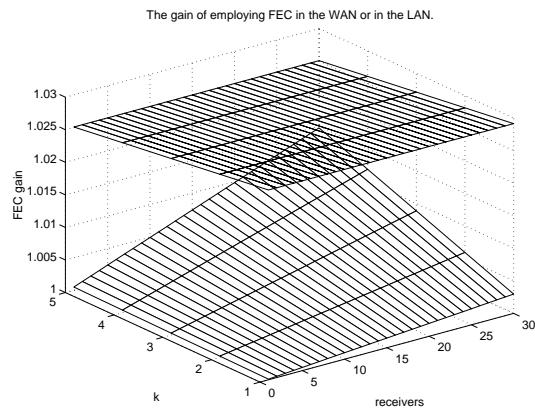


Figure 8: The gain for employing FEC just in the WAN (upper plane), or just in the LAN (lower plane).

We will model a small multicast group connected by wireless links in the private network, yielding a higher packet loss probability of $q_2 = 10^{-2}$ than in the WAN ($q_1 = 10^{-3}$). The other parameters of the model stay the same as in the LAN-WAN case. Figure 9 shows that the protection of the wireless links in the LAN is the thing to do and that the application of FEC on the WAN part results in no FEC gain. This is mainly due to the fact that the $h - k$ links in the WAN have already a small link loss probability, so that a further lowering of the packet loss probability in the WAN does not result in a decrease of number of transmissions, since the high link loss probability in the wireless LAN stays constant. The FEC gain of applying FEC on the wireless links is shown in figure 9 with the upper plane, where the gain increases with the number of receivers up to 1.8. As already stated in the results for FEC on the whole multicast tree the FEC gain increases with the number of independent LAN links k .

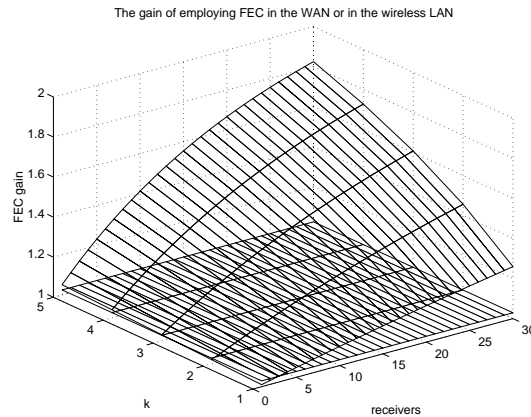


Figure 9: The gain for employing FEC just over the shared WAN links (lower plane), or just on the tree inside the wireless LAN (upper plane).

The results in this section give rise to the assumption that FEC applied on the links with the highest packet loss probability is always the thing to do. In the next section we will see that this does not hold for big multicast groups and we will give the reasons.

4.3 Big Multicast Groups: Partial FEC

The case described in the previous sections was for a small group of receivers. We will now use tree 4 (see figure 3) to model big multicast groups where multiple receivers are located in several LANs.

We have $r_{LAN} = 5, 25, 100$ receivers in every LAN, the total number of receivers $r = 100 \dots 3000$ varies. The link loss probability is $q_1 = q_2 = 10^{-3}$ in the WAN and $q_3 = 10^{-4}$ in the LAN. The length of the common WAN path varied from $h - k - l = 0 \dots 24$. The results are shown in the figures 10, 11 and 12. We assume again that FEC lowers the link loss probability by two orders of magnitude. We compare the cases, where FEC is employed either on the WAN part or the LAN part of the multicast tree.

5 receivers in a LAN

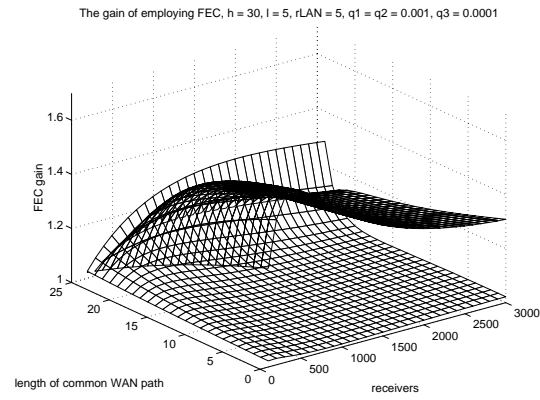


Figure 10: The gain of FEC in WAN (upper plane) and FEC in LAN (lower plane).

25 receivers in a LAN



Figure 11: The gain of FEC in WAN (plane up to 1.2) and FEC in LAN (plane up to 1.6).

100 receivers in a LAN

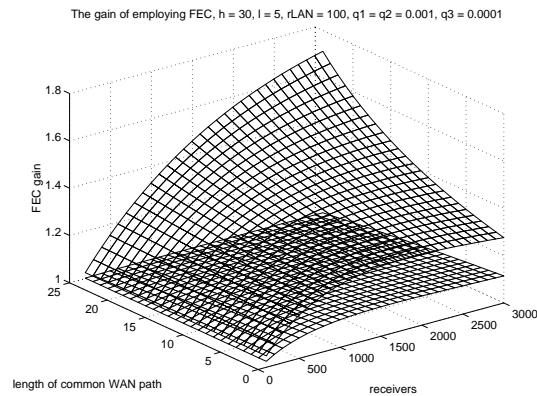


Figure 12: The gain of FEC in WAN (lower plane) and FEC in LAN (higher plane).

The results for small multicast groups indicated that the FEC gain is highest, when applied over the links that have the highest link loss probability. This does not hold for big groups. The high link loss probability can be found on WAN links, but the FEC gain is higher for FEC applied on the links in the LANs than on the links in the WAN, for the case of $r_{LAN} = 100$ receivers (figure 12). The FEC gain achieved by the application of FEC on the WAN is decreasing with an increasing r_{LAN} (figure 10 – figure 12) of receivers that can be found in a single LAN until the FEC gain for FEC on the WAN is lower than for FEC on the LAN (figure 12). The reason can be found in the high decrease of links in the WAN with **increasing** r_{LAN} , while the total number of links in all LANs is always constant ($5r$), whatever the number of receivers in one LAN r_{LAN} is.

5 CONCLUSION

We defined two generic multicast trees that allowed us to capture the essential characteristics of a multicast tree such as degree of link sharing or group size. In both cases we derived the analytical formula for the average number of transmissions necessary to reliably deliver a packet to all receivers of the group. The formula allowing to model for different link loss in different parts of the multicast tree. Our main results are as follows:

When FEC is applied to all links of the multicast tree we observed that

- for small groups the FEC gain increases as the number of receivers increases (figure 4) and the number of shared links decreases.
- for big groups the FEC gain is larger than for small groups and is independent of the number of receivers.

When FEC is applied either on the links of the LAN part or either on the links of the WAN part of the multicast tree, we observed that

- for small groups the highest FEC gain is achieved when applied on the part, where the links have a high loss probability.
- for big groups this is only true for a low number r_{LAN} of receivers that can be found in one LAN. For a large number of receivers in one LAN, FEC should be applied on the LAN links.

We also showed that the assumption of independent loss among receivers is a worst case assumption. In [BMT94] the cumulative loss probability experienced at the source was evaluated for different generic multicast trees for the same number of links in those trees. The authors conclude that minimizing the degree of overlap between paths in a multicast tree results in more efficient reliable multicast. This is true for multicast groups that are connected to the source with the **same** number of links either by independent paths or by overlapping paths. But, in general, overlapping paths result in a lower number of links in the multicast tree, compared to a multicast tree consisting of independent paths. Our multicast tree model (tree 3) indicates that the number of links has a higher impact on the number of transmissions than has loss on shared links. The number of transmissions decreases with a high overlap of paths.

REFERENCES

- [BFC93] Tony Ballardie, Paul Francis, and Jon Crowcroft. Core based trees (cbt). In *Proceedings of SIGCOMM'93*, pages 85–95, Ithacs, N.Y., USA, September 1993. ACM.
- [BMT94] Pravin Bhagwat, Partho P. Mishra, and Satish K. Tripathi. Effect of topology on performance of reliable multicast communication. In *Proceedings of INFOCOM'94*, volume 2, pages 602–609, Toronto, Ontario, Canada, June 1994. IEEE.
- [Car95] Georg Carle. Towards scaleable error control for reliable multipoint services in atm networks. In *12th International Conference on Computer Communication, ICC'95*, Seoul, Korea, August 1995.
- [Dee88] Stephen E. Deering. Multicast routing in internetworks and extended lans. *Computer Communications Review*, 18(4):55–64, 1988.
- [DEF+96] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C.G. Liu, and L. Wei. The pim architecture for wide-area multicast routing. *IEEE/ACM Transactions on Networking*, 4(2):153–162, April 1996.
- [Den93] Robert H. Deng. Hybrid arq schemes for point-to-multipoint communication over nonstationary broadcast channels. *IEEE Transactions on Communications*, 41(9):1379–1387, September 1993.
- [FJM+95] Sally Floyd, Van Jacobsen, Steven McCanne, Ching-Gung Liu, and Lixia Zhang. A reliable multicast framework for light-weight sessions and application level framing. In *Proceedings of SIGCOMM'95*, volume 25, pages 342–356. ACM, October 1995.
- [Gro96] Matthias Grossglauser. Optimal deterministic timeouts for reliable scalable multicast. In *Proceedings of INFOCOM'96*, pages 1425–1432, San Francisco, CA, USA, March 1996.

- IEEE.
- [Hof96] Markus Hofmann. A generic concept for large scale multicast. In *International Zürich Seminar*, Zürich, Switzerland, February 1996.
- [LP96] John C. Lin and Sanjoy Paul. Rmtp: A reliable multicast transport protocol. In *Proceedings of INFOCOM'96*, pages 1414–1424, San Francisco, CA, USA, March 1996. IEEE.
- [McA90] A. J. McAuley. Reliable broadband communications using a burst erasure correcting code. In *Proceedings of SIGCOMM'90*, pages 287–306, Philadelphia, PA, September 1990. ACM.
- [PTK94] Sridar Pingali, Don Towsley, and James F. Kurose. A comparison of sender-initiated and receiver-initiated reliable multicast protocols. In *Proceedings of Sigmetrics, Conference on measurement and modeling of Computer Systems*, pages 221–230, Santa Clara, CA, USA, May 1994. ACM.
- [RA95] Clark Russell and Mostafa Ammar. Providing scalable web service using multicast delivery. In *Proceedings of the IEEE Workshop on Services in Distributed and Networked Environments*, Whistler, Canada, June 1995.
- [ST87] Nachum Shacham and Don Towsley. Resequencing delay and buffer occupancy in selective repeat arq with multiple receivers. In *Proceedings of INFOCOM'87*, pages 512–520. IEEE, 1987.
- [TM87] D. Towsley and S. Mithal. A selective repeat arq protocol for a point to multipoint channel. In *Proceedings of INFOCOM'87*, pages 521–526, San Francisco, CA, USA, April 1987. IEEE.
- [Tow85] Don Towsley. An analysis of a point-to-multipoint channel using a go-back-n error control protocol. *IEEE Transactions on Communications*, 33(3):282–285, March 1985.
- [YGS95] Rajendra Yavatkar, James Griffioen, 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.
- [ZS91] M. El Zarki and N. Shroff. Performance analysis of error control schemes in interconnected lan-wan-lan networks. In *Proceedings of INFOCOM'91*, pages 1377–1385. IEEE, 1991.