Performance Modelling of Reliable Multicast Transmission

Jörg Nonnenmacher Ernst W. Biersack
Institut EURECOM
B.P. 193, 06904 Sophia Antipolis Cedex, FRANCE
{nonnen,erbi}@eurecom.fr

Abstract

Our aim is to investigate reliable transmission for multicast communication and explore its relationship to multicast routing. We derive two characterizations that enable the comparison of routing algorithms and error recovery mechanisms with respect to the multicast tree topology, namely the probability mass function of successful receptions and the expected number of retransmissions needed to deliver a packet from the source to all receivers. We also give a tight approximation of the computationally expensive expected number of retransmissions. These expressions allow to explore the relationship between routing and error recovery for multicast communication. We finally evaluate the impact of routing algorithms on the performance of reliable multicast transmission and give a realistic generic model for a multicast tree.

Keywords: Reliable Multicast, MBONE, Multicast Routing, Performance Evaluation, ARQ

1 Introduction

The MBONE [1] has given raise to a number of conferencing applications such as vat, ivs, or vic where timely delivery is most important and packet loss can be tolerated. However, there is another class of dissemination-oriented applications where *reliable* multicast delivery from one source to many receivers is required such as

- Information delivery e.g., newspaper excerpts, software updates and software distribution.
- Distributed Simulation where state information must be exchanged.
- Web caching and replication for cache hierarchies such as Harvest.

In order to get a handle for designing and evaluating reliable multicast transport protocols one needs to be able to compute performance measures such as delay or the number of retransmissions. We will derive the formulas for computing

- the probability mass function (pmf) for the number of receivers that successfully receive a packet that is emitted once.
- the mean number of retransmissions until all receivers have successfully received a packet.

Since the exact expression is difficult to compute we also give a simple approximation for the mean number of retransmissions.

Our aim is to investigate reliable transmission for multicast communication and explore its relationship to multicast routing. Very little work [2] was done in this area and the effect of the topology on reliable multicast is not well understood.

Recent multicast routing algorithms have been evaluated in terms of cost and delay [3, 4, 5], blocking probability [6, 7] and overhead [8]. The impact of the routing algorithm on reliable multicast transmission has not yet been studied. Our results enable us to study the impact of multicast routing algorithms on reliable transmission. We will demonstrate the impact for two multicast routing algorithms that are known to perform best in terms of cost and delay.

Nearly all the research on the performance of reliable multicast communication [9, 10, 11, 12, 13] assumes multicast trees where the loss on any link affects only a single receiver.

We will take this special case of a multicast tree, referred to as $MFAN$ (see figure 3) into consideration and compare it both, with trees that are the outcome of routing algorithms and with two other generic multicast trees. We will show that the full binary tree (see figure 1) is a more realistic model for a multicast tree than $MFAN$.

2 Multicast Trees

The formulas we derive are valid for all types of multicast trees, i.e. they are independent of the topology of the multicast trees. In order to evaluate the formulas we define three generic multicast trees and additionally use two of the most popular multicast routing algorithms to compute multicast trees for artificially generated networks.

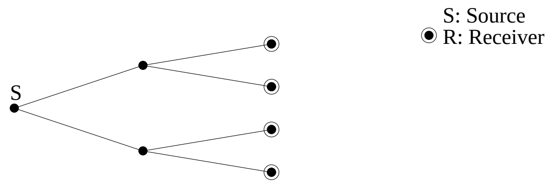

Figure 1: Full Binary Tree (FBT)

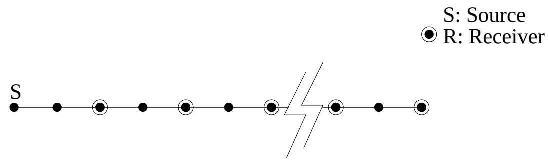

Figure 2: Linear Chain (LC)

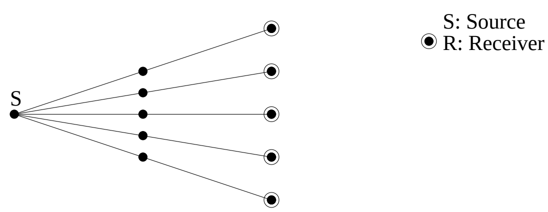

Figure 3: Multihop-Fanout ($MFAN$)

A 1:n – multicast connection forms a tree rooted at the source. The loss in a multicast tree is dependent on the topology. A tree topology has several parameters, each of them having a different influence on loss: (i) tree height, (ii) number of receivers (members in the multicast group), (iii) number of nodes in the tree[2], and (iv) the number of receivers affected by a loss over a single link.

We have chosen these particular three generic multicast trees because they behave very differently with respect to the impact of packet loss on a single link:

- For $MFAN$ (figure 3), always only a single receiver is affected.
- For the linear chain LC (figure 2), depending on what link the loss occurs, the number of affected receivers can range from one to all receivers.
- For the full binary tree FBT (figure 1), the impact of loss lies between the one for $MFAN$ and LC, affecting either a single receiver or a subgroup of all receivers.

By keeping the ratio of the number of receivers and the number of tree nodes for all three trees approximately at 0.5 (see Table 1) we collapse the two parameters (ii) and (iii) that influence loss into a single one.

However, as the tree grows, the tree height will vary if we keep the ratio of receivers and nodes in the tree fixed (see Table 1).

	$MFAN$	FBT	LC
the ratio $\frac{receivers}{nodes}$	$\frac{1}{2+\frac{1}{r}}$	$\frac{1}{2-\frac{1}{r}}$	$\frac{1}{2+\frac{1}{r}}$
tree height	2	$\log_2(r)$	$2r$

Table 1: The characteristic of the three generic multicast trees with respect to the number r of receivers

To generate "real" multicast trees we use two different multicast routing algorithms that base their routing decision on the optimization of cost or delay:

Cost optimization tries to minimize the sum of the edge costs in the multicast tree. The Kou Markovsky Berman algorithm [14], referred to as KMB, is presently the most famous heuristic to approach the optimal cost solution for a multicast tree. It constructs a **Heuristic Steiner Tree** (HST) [15] based on the minimum spanning tree algorithm.

Delay optimization minimizes the delay from the source to every receiver. The Shortest Path Algorithm, analyzed by Doar [4] optimizes delay and constructs a **shortest path tree** (SPT) that connects every receiver to the source via the shortest path.

10 random networks with 200 nodes and an average outdegree of 3.0 were constructed following Waxman [16], with the modification of Doar in [4] that avoids the influence of the number of nodes on the average outdegree. The method of Waxman is commonly used by the Multicast Routing community [3, 4, 16, 17] to compare the performance of different Multicast Routing Algorithms on random networks.

On each of the 10 random nets, 100 multicast groups with varying group sizes (5...140) and random receiver locations had been routed by the two algorithms for Cost and Delay optimization. A sample SPT is shown in figure 4 and a sample HST for the same network and the same group of 5 receivers is shown in figure 5.

The characteristics of the two multicast trees is shown in figure 6 and figure 7. In figure 6 it can be seen that the ratio between receivers and nodes in the multicast tree is not a constant. For $r = 40$, the trees have a ratio receiver to nodes that is about 0.5 and are therefore comparable with the generic trees.

[2] The number of edges in a tree is not stated, since for a tree: $edges = nodes - 1$.

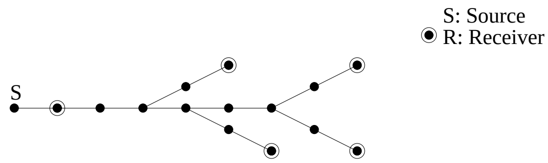

Figure 4: A SPT with 5 receivers computed on a random network

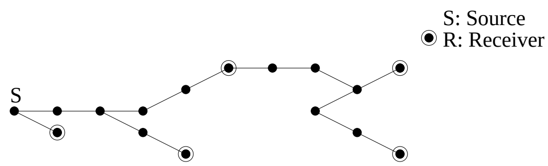

Figure 5: A HST for the same multicast group in the same random network

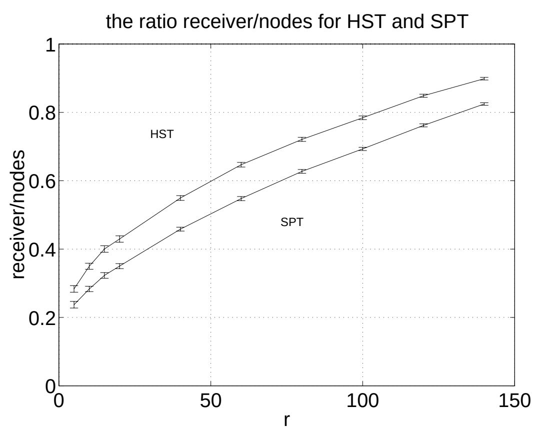

Figure 6: The ratio $\frac{receivers}{nodes}$ for HST and SPT

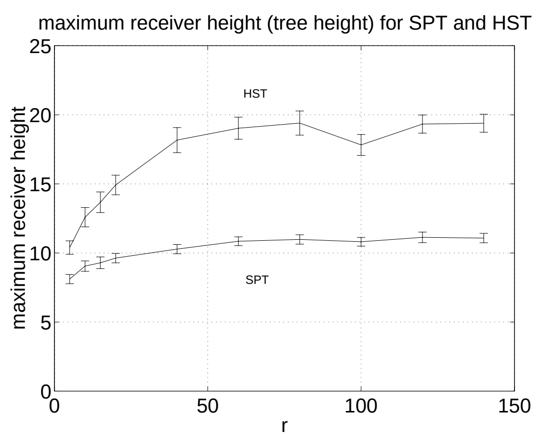

Figure 7: The tree height of HST and SPT

3 Loss characteristics of a multicast tree

Loss in a multicast tree affects several receivers, if it happens on a link that leads to several receivers. We will call such a link **shared link**.

Reliable multicast transmission has to deal with two major problems:

- **Feedback implosion**: Receivers in a reliable multicast communication have to provide the source with the status of the reception. Loss on shared links causes loss at several receivers and increases the amount of feedback.

- **High number of retransmissions**: The higher the number of receivers the higher becomes the number of links in the multicast tree and the average number of retransmissions.

We derive a formula to analytically evaluate the **feedback implosion** at the source, by calculating the probability mass function (pmf) of successful and unsuccessful receptions for *a single packet emission*. We also give the expectation of the number of receptions and show its independence of shared links.

We give the expected number of **retransmissions** needed to deliver one packet to all receivers and propose a tight approximation that enables loss prediction for adaptive error control mechanisms.

3.1 The number of successful receptions in a multicast tree

Supposed that a packet is sent **once**, we are interested in the pmf of the number of receivers that successfully receive this packet.

Given is a multicast tree mct:

- with source S as the root

- r receivers placed at arbitrary nodes and at all leaves. We allow at most one receiver at a node in the tree, and we assume not to have a receiver at the source

- homogeneous link loss probability q of a packet.

Let X_S be the number of receivers out of the r receivers in the multicast tree rooted at S that receive the packet successfully when transmitted once from S. We will give a method to calculate the corresponding probability mass function (pmf) $p(X_S = k)$, which enables us to capture the loss characteristic of different multicast trees. First of all, some definitions:

n	A node in the *mct*.
$child(n)$	The set of children (successors) of n.
c_n	The number of children of n, $c_n = card(child(n))$.
r_n	The number of receivers in the subtree rooted at n. If n is a receiver, it is **not** included. The number of receivers in the whole tree is therefore $r = r_S$.
X_n	A random variable, describing the number of receivers out of the r_n receivers in the subtree rooted at n that successfully receive a packet, when transmitted from node n.
$p(X_n = k)$	The *pmf* of X_n, where $k = 0, \ldots, r_n$.
$s_n \in \{0,1\}^{c_n}$	Link success vector, indicating the success or loss of a packet transported via the links leading to the children of n.
$s_n(i)$	The i–th component of s_n. The success ($s_n(i) = 1$) indicator of the link leading from node n to its child i. If the packet is lost on the link leading from n to i, then $s_n(i) = 0$.
$x_n \in \{0,1\}^{c_n}$	The children receiver vector. Indicates which child of n is a receiver.
$x_n(i)$	The i–th component of x_n. $x_n(i) = 1$ indicates that the child i of node n is a receiver, otherwise $x_n(i) = 0$.
$a_n \in \times_{i=1}^{c_n}\{0, \ldots, r_i\}$	Behind child receptions vector. Gives the number of receivers behind the children of node n that received the packet successfully.
$a_n(i) \in \{0, \ldots, r_i\}$	The i–th component of a_n, gives the number of receivers in the subtree rooted at the child i of n that received the packet successfully.

The *pmf* can now be calculated in a recursive way, starting at the leaves of the multicast tree. We need to distinguish two cases:

Node n is a leaf. Then there are no receivers located behind node n and the probability that no receiver is receiving a packet is 1 and the *pmf* evaluates trivially to:

$$p(X_n = 0) = 1 \tag{1}$$

Node n is not a leaf. The *pmf* $p(X_n = k)$ is given by the sum of the probabilities of all different combinations of k successful receptions in the tree rooted at n. The recursive way of calculating the *pmf* allows the use of already known probabilities $p(X_i = a_n(i))$ at the children $i \in child(n)$ of n. For every node n we have therefore just to look at the adjacent links leading to the children.

We must sum over all the combinations of link success that allow in total k successful receiving receivers located at the children i of n and in the subtrees rooted at each of the children.

For one combination s_n of link success the number of successful receptions at the direct children, being also receivers, is given by the inner product $s_n^T x_n$. The number of receptions in the subtrees rooted at the children is given by $s_n^T a_n$.

To obtain the number k of successful receptions for a given s_n the following condition must hold:

$$k = s_n^T (a_n + x_n) \tag{2}$$

Since x_n is constant and s_n is given, equation (2) selects a subset of combinations of receptions in the subtrees rooted at the children of n: $A_n(s_n) = \{a_n | k = s_n^T(a_n + x_n)\} \subset \times_{i=1}^{c_n}\{0, \ldots, r_i\}$.

Different number of receptions in subtrees behind a failing link do not change the probability $p(X_n = k)$. $A_n(s_n)$ can therefore be reduced by masking the number of receptions in subtrees behind failing links.

$$A_n(s_n) = \{\, a_n \;\; | \;\; k = s_n^T(a_n + x_n) \;\wedge\; \forall i:\; s_n(i)a_n(i) = a_n(i)\} \tag{3}$$

The probability for one combination s_n of link success and one $a_n \in A_n(s_n)$ is then given by the product over the children:

$$p(a_n, s_n) = \prod_{i \in child(n)} \{\; s_n(i)(1-q)p(X_i = a_n(i)) + (1 - s_n(i))q \;\} \tag{4}$$

Since the link to child i is successful ($s_n(i) = 1$) with probability $(1-q)$ and the probability of $a_n(i)$ successful receptions in the subtree rooted at child i is $p(X_i = a_n(i))$. The packet gets lost ($(1 - s_n(i)) = 1$) on the link to child i with probability q and $a_n(i)$ has no contribution.

The probability $p(X_n = k)$ is then given by summing over all link success combinations s_n and all $a_n \in A_n(s_n)$:

$$p(X_n = k) = \sum_{s_n} \sum_{a_n \in A_n(s_n)} p(a_n, s_n) \tag{5}$$

We show $p(X_S = k)$ for the generic multicast trees with a link loss probability of $q = 0.03$ in figure 8 for $r = 64$ receivers and, for $r = 128$ receivers in figure 9.

We can see that the *pmf*s vary significantly for the three generic multicast trees. This is due to the fact that the number of receivers affected by a loss on a single link also differs widely for the three generic multicast trees.

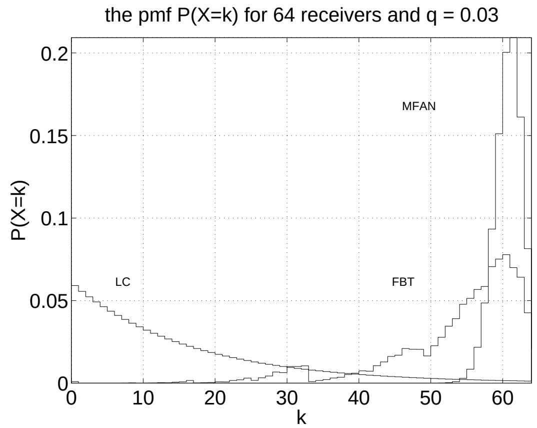

Figure 8: The probability mass function $P(X_S = k)$ for FBT, $MFAN$ and LC for 64 receivers and a link loss probability of $q = 0.03$

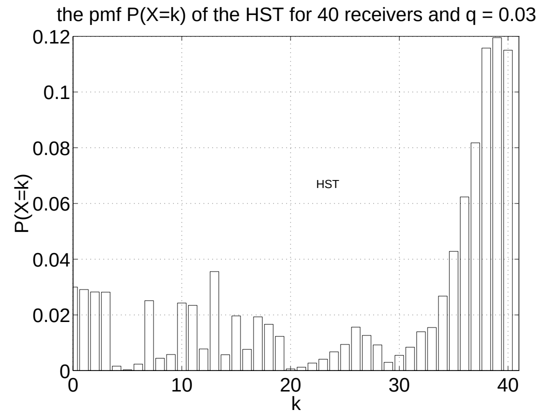

Figure 10: The probability mass function $P(X_S = k)$ for a HST with 40 receivers and a link loss probability of $q = 0.03$

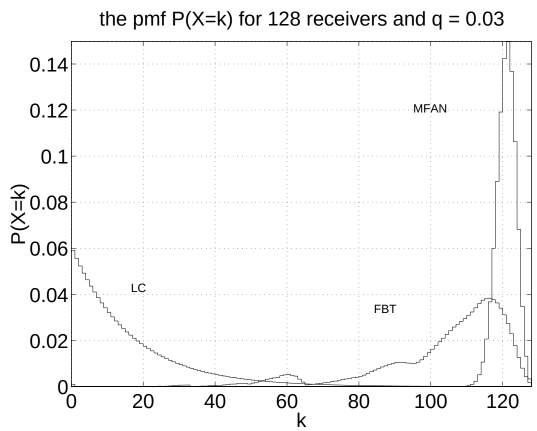

Figure 9: The probability mass function $P(X_S = k)$ for FBT, $MFAN$ and LC for 128 receivers and a link loss probability of $q = 0.03$

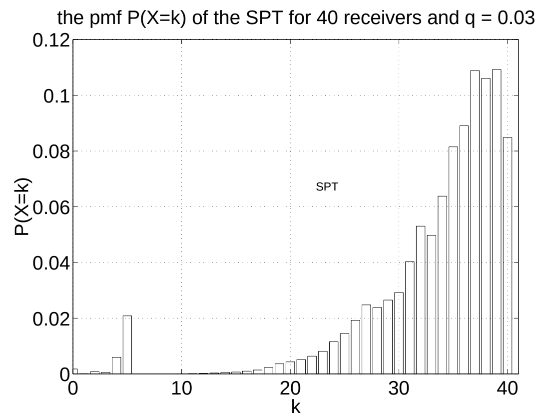

Figure 11: The probability mass function $P(X_S = k)$ for the SPT with 40 receivers and a link loss probability of $q = 0.03$

The pmf of the $MFAN$ is the binomial pmf, the pmf of the LC approximates the geometric pmf for a large number of receivers. The curve of the FBT is multimodal with peaks at $k = 2^{h-1}, 2^{h-1} + 2^{h-2}, \ldots$. These peaks are due to a high number of full binary subtrees with $2^{h-2}, 2^{h-3}, \ldots$ receivers and therefore a high number of possible combinations that lead to a sum of k successful receptions, whereas for $k+1$ successful receptions the number of possible combinations of full binary subtrees is much lower.

The $pmfs$ for the HST and the SPT for the same multicast group on the same network (figures 10 and 11) indicate that the variance of the number of successful receptions for the HST is higher than for the SPT. The high probabilities for low numbers of successful receivers are due to shared paths near the source. We observe that the $pmfs$ for the HST and the SPT resemble most closely the pmf for the FBT.

3.2 The number of responses

We are interested in the number of responses, which can be either positive or negative ACKs, we can expect from the r receivers in the multicast tree, when a packet is emitted once by the source. We make the assumption that the feedback reverse channel from the receivers to the source is loss–free, in which case the number of ACKs/NAKs is identical to the number of receivers that have received or have not received a packet.

X_S is a random variable that describes the number of successful receptions in the whole multicast tree. X_S is the sum of random variables $X_{S,i} \in \{0, 1\}$, each describing the reception of a single receivers i: $X_S = \sum_{i=1}^{r} X_{S,i}$. Since we assume uniform link loss q on all links, the probability of a successful reception for receiver i, which lies h_i hops away from the source, is $P(X_{S,i}) = (1-q)^{h_i}$. The expected number of ACKs for every single receiver is therefore

$E(X_{S,i}) = P(X_{S,i})$. The expected number of successful receptions $E(X_S)$ in a tree with r receivers is then:

$$E(X_S) = E(\sum_{i=1}^{r} X_{S,i}) = \sum_{i=1}^{r} E(X_{S,i}) = \sum_{i=1}^{r}(1-q)^{h_i} \tag{6}$$

We can also express $E(X_S)$ dependent on the receiver distribution over the tree levels h, by accumulating receivers that have the same distance from the source. Let n_h be the number of receivers that lie in tree level h, e.g. h hops from the source, then:

$$E(X_S) = \sum_{h=1}^{h_{max}} n_h(1-q)^h \tag{7}$$

gives the expected number of ACKs. Please note that $E(X_S)$ is **not** dependent on the number of shared links, since in (6) the path from the source to every receiver accounts by its full length.

The expected number $E(X_S)$ of ACK-packets at the source is shown in figure 12 as a function of the number of receivers in the multicast group for a link loss probability $q = 0.03$. For HST, the number of ACKs is slightly lower than for SPT, accounting for the fact that the number of links traversed between the source and a receiver is higher for HST than for SPT.

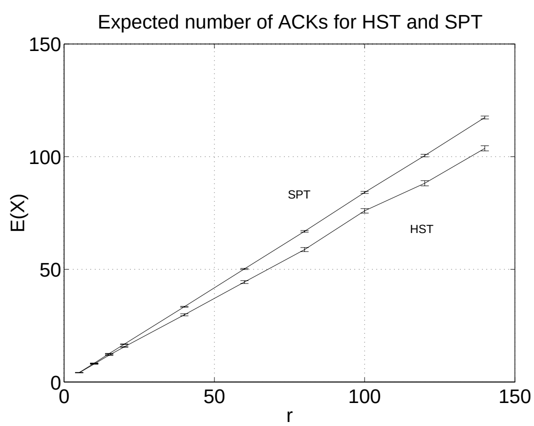

Figure 12: Expected number of ACK-packets at the source for a link loss probability of $q = 0.03$.

The error control scheme may use positive ACKs or negative ACKs (NAKs). Let $Y_S = r - X_S$ be the random variable that describes the number of unsuccessful receptions, then the pmf of Y_S is:

$$p(Y_S = k) = p(X_S = r - k)$$

and the expected number of NAKs for r receivers is given as:

$$E(Y_S) = r - E(X_S)$$

.

3.3 The expected number of transmissions for reliable delivery

The expected number of multicast transmissions to deliver a packet to **all** receivers is an important measure in reliable multicast communication. The expected number of transmissions captures the global packet loss behaviour in the tree and the cost and the time of a reliable multicast delivery. The expected number of multicast transmissions depends on the link loss probability q and the topology of the multicast tree.

In [2], the expected number of multicasted transmissions is given for the case of loss at **nodes** in the multicast tree. It is more appropriate to consider loss on a link due to two reasons: loss at the source node is unlikely and link loss can be associated with loss in output buffers in routers. In [18] the expected number of multicasted transmissions for link loss is given by a slight modification of the formula given in [2].

The Cumulative Distribution Function (CDF) $F_n(i)$ of the number of transmissions for the link leading to node n and the tree rooted at n is calculated in a recursive fashion starting at the leaves: It has to be distinguished, if node n is a leaf l, an internal node, or the source S (for details see [18]):

$$\begin{aligned} F_l(i) &= 1 - q^i \\ F_n(i) &= \sum_{u=0}^{i-1} \binom{i}{u} q^u (1-q)^{(i-u)} \prod_{c \in child(n)} F_c(i-u) \\ F_S(i) &= \prod_{c \in child(S)} F_c(i) \end{aligned} \tag{8}$$

Using $F_S(i)$, the expected number of multicasted transmissions $E(T(S))$ from the source S is:

$$E(T(S)) = \sum_{i=0}^{\infty}(1 - F_S(i)) \tag{9}$$

The expected number of retransmissions $E(R(S))$ is:

$$\begin{aligned} E(R(S)) &= E(T(S) - 1) = E(T(S)) - 1 \\ &\overset{(9)}{=} \sum_{i=1}^{\infty}(1 - F_S(i)) \end{aligned} \tag{10}$$

3.4 A useful approximation for $E(R(S))$

Reliable multicast protocols need to know the expected number of retransmissions. However, the exact calculation as derived above is not practical:

- The expected number of retransmissions is hard to calculate, since the calculation of the recursive CDF in (8) is computationally intensive for arbitrary topologies.
- Adaptive transport protocols need simple but effective mechanisms to decide.

We give a tight and very simple approximation. The expected number of retransmissions is approximately the product of the link loss probability q and the number of links L in the multicast tree:

$$E(R(S)) \approx qL \tag{11}$$

This approximation is tight for $qL \leq 1$.

For space reasons we limit ourselves to a sketch of the derivation of the approximation for $E(R(S))$ (for details see [19]). By induction over the children is shown that every $F_n(i)$ can be expressed in the form

$$F_n(i) = 1 - \sum_{j_n^-}(Q_{j_n^-})^i + \sum_{j_n^+}(Q_{j_n^+})^i$$

where $Q_{j_n^-}$ and $Q_{j_n^+}$ are polynoms in q: $Q = \sum_k \zeta_k q^k$, with a minimal exponent $k_{min} \geq 1$. The difference between the sum $\sum_{j_n^+} \zeta_{1_{j_n^+}}$ of the ζ_1 of all the polynoms $Q_{j_n^+}$ with $k_{min} = 1$ and the sum $\sum_{j_n^-} \zeta_{1_{j_n^-}}$ of the ζ_1 of all the polynoms $Q_{j_n^-}$ with $k_{min} = 1$ equals the number of links in the subtree rooted at n. If there is a link leading to node n the difference is one greater than the number of links in the subtree rooted at n.

Afterwards, the expectation is calculated by:

$$E(R(S)) = \sum_{i=1}^{\infty}(1 - F_S(i))$$

Which results in:

$$E(R(S)) = \sum_{j_S^-} \frac{Q_{j_S^-}}{1 - Q_{j_S^-}} - \sum_{j_S^+} \frac{Q_{j_S^+}}{1 - Q_{j_S^+}}$$

Then, the ratios $\frac{Q}{1-Q}$ are approximated by Q, yielding

$$E(R(S)) \approx \sum_{j_S^-} Q_{j_S^-} - \sum_{j_n^+} Q_{j_n^+}$$

Finally, are we interested in the term q of the polynom Q, due to its relevance compared with the terms $q^2, q^3, \ldots$. Every polynom $Q = \sum_k \zeta_k q^k$ is approximated by $\zeta_1 q$, resulting in an approximation of the expected number of retransmissions as:

$$E(R(S)) \approx q(\sum_{j_S^-} \zeta_{1_{j_S^-}} - \sum_{j_n^+} \zeta_{1_{j_S^+}}) = qL$$

The last approximation, where higher order terms are suppressed, also gives us the condition for which the whole approximation of the expected number of retransmissions (11) is valid:

$$qL \leq 1$$

since for $qL \geq 1$ a second order term q^2 has an impact of one or more links on the approximative expectation: $q^2 \cdot L = q(qL) \geq q \cdot 1$.

We compare the two most extreme cases of multicast topologies. The first one is called linear chain (LC) and is just a chain of L links. The other extreme is called $MFAN$ and is the so beloved, frequently used, model for a multicast tree in performance evaluation of reliable multicast communication. The $MFAN$[3] has one separate link from the source to every of the L receivers. In both cases, we have L links. LC is the deepest, $MFAN$ the broadest multicast tree that can be built with L links. In figure 14 we compare the exact expected numbers of retransmissions for $MFAN$ and LC with the approximation qL as a function of the number of links. We previously saw that the loss characteristic and the pmf for the number of successful receptions of the FBT is similar to the one for real multicast trees. We therefore compare separately the expected number of retransmissions in the FBT with the approximation qL for $L = 30$ links in figure 13 and observe that the approximation is very tight.

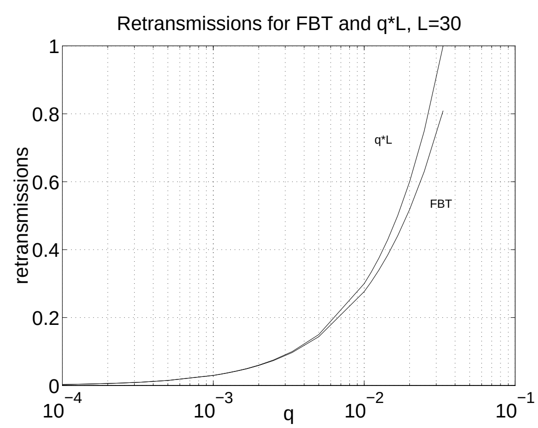

Figure 13: The expected number of retransmissions for the FBT and the approximation qL for different q.

[3]To get the most extreme multicast tree, we reduce the height of the $MFAN$ to $h = 1$, compared to the previous definition of MFAN, where we had chosen $h = 2$ (figure 3).

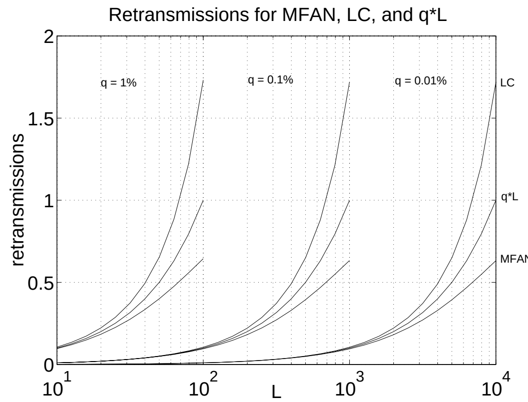

Figure 14: The expected number of retransmissions for the most extreme multicast trees and the approximation qL with respect to the number of links L.

4 Implications of our work

We demonstrate the impact of our results in the following two domains:

- We show that multicast routing algorithms that optimize delay achieve better delay and throughput performance for reliable multicast communication, than algorithms that optimize cost.
- We show that the FBT is a good generic model of a multicast connection and that more realistic results are obtained than with the usual used $MFAN$.

4.1 Impact of Routing on Error Recovery

Multicast routing algorithms have been designed that take into account several metrics. However, the performance of an algorithm is the most time evaluated for cost and delay – the performance for reliable transmission is left aside. We give a tight approximation for the number of retransmissions needed to deliver one packet from the source to all receivers: $E(R(S)) \approx qL$ that enables performance evaluation of routing algorithms with respect to loss.

In the case where a unique link cost is chosen, the cost of the multicast tree is proportional to the number of links L in the multicast tree and therefore approximately proportional to the number of retransmissions.

For a given loss rate the performance of error recovery schemes for point to point connections is determined by the Round Trip Time (RTT) between the source and the receiver. We define the Round Trip Time as two times the sum of the propagation and transmission delays of the links on the path from the source to the receiver.

For a multicast connection, the receiver connected to the source via the longest path (in terms of delay) is the feedback bottleneck for the error recovery scheme. The RTT of a multicast connection is therefore defined as two times the sum of the propagation and transmission time on the links on this longest path and depends on the routing algorithm. The RTT for the HST is about two times higher than the RTT for the SPT (see figure 7). On the other hand, the difference between HST and SPT in terms of the expected number of retransmissions, using $E(R(S)) \approx qL$, is minor (compare figure 15), with the values for SPT being only slightly higher than for HST.

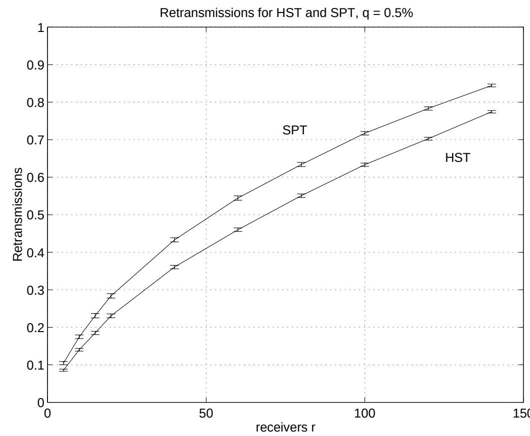

Figure 15: The expected number of retransmissions for HST and SPT for $q = 0.005$.

From these two observations, we can conclude that delay optimization (SPT) in multicast routing algorithms yields better delay and throughput performance for reliable transmission than does cost optimization (HST).

Applications with a stringent time–constraint profit also from routing algorithms that optimize delay. In recent years routing algorithms have been designed that optimize cost and try to meet a delay–constraint. However, most of the algorithms optimizing cost do not support dynamic multicast group membership changes – the SPT does.

We believe that SPT routing is the best solution for multicast routing. Due to its simplicity, it can use the routing of the underlying unicast algorithm, its support for dynamic membership changes and its good performance for reliable transmission as for applications that need timely delivery.

4.2 A better generic model for multicast trees: Full Binary Tree

We saw in previous sections that the loss characteristics of the FBT is very close to the loss characteristics of HST and SPT. To confirm that FBT is a good generic model for a multicast tree, we compare the **link share** in different trees, i.e. to what degree do receivers in a tree share common paths.

Let L be the number of links and r be the number of receivers in the multicast tree, then the link share of one link l_i, $i = 1, \ldots, L$ can be defined as the number of receivers $rd(l_i)$ that share the cost on link l_i divided by the total number of receivers: $ls(l_i) = \frac{rd(l_i)}{r}$. The link share ls for the entire tree mct is defined as the average link share of all links:

$$ls(mct) = \frac{1}{L}\sum_{i=1}^{L}\frac{rd(l_i)}{r} \qquad (12)$$

For a tree, there are several methods to define a measure of link share. We compared measures of link share and found that the definition given in (12) reflects well the degree to which receivers share links in a tree. For a further discussion on definitions of link share see [20].

The link share of the FBT is nearly identical with the link share of the SPT (see figure 16). The choice of the FBT as a model for a multicast tree is further confirmed by the degree to which receivers share the links in the multicast tree. The HST has a higher link share than the SPT, since the routing algorithm tries to connect the receiver set with a minimal cost, resulting in a high number of receivers that share an average single link in the multicast tree.

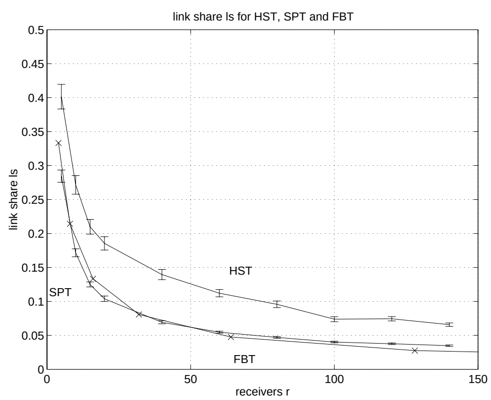

Figure 16: The link share ls for SPT, HST and FBT

5 Conclusion

We evaluated the impact of routing on reliable multicast and achieved two main results. First, multicast routing that optimizes delay achieves better throughput and delay performance for reliable multicast than cost optimal routing. Second, the full binary tree (FBT) is a good generic model for the loss characteristics of real multicast trees and provides more realistic results than the usual $MFAN$, in which a loss affects always only one receiver. We derived two characterizations that enable the comparison of routing algorithms and error recovery mechanisms with respect to the multicast tree topology, namely a pmf for the number of successful receptions when a packet is emitted once from the source and the expected number of retransmissions needed to deliver a packet from the source to all receivers. We show that the product of the link loss probability q and the number of links L in an arbitrary multicast tree tightly approximates the expected number of retransmissions $E(R)$ under the condition that $qL < 1$:

$$E(R) \approx qL$$

.

References

[1] H. Eriksson, "MBONE: The multicast backbone," *Communications of the ACM*, vol. 37, pp. 54–60, Aug. 1994.

[2] P. Bhagwat, P. P. Mishra, and S. K. Tripathi, "Effect of topology on performance of reliable multicast communication," in *Proceedings of INFOCOM'94*, vol. 2, (Toronto, Ontario, Canada), pp. 602–609, IEEE, June 1994.

[3] L. Wei and D. Estrin, "The trade-offs of multicast trees and algorithms," in *Proceedings of ICCCN'94*, (San Francisco, CA, USA), Sept 1994.

[4] M. Doar and I. Leslie, "How bad is naïve multicast routing," in *Proceedings of INFOCOM'93*, vol. 1, pp. 82–89, IEEE, 1993.

[5] V. P. Kompella, J. C. Pasquale, and G. C. Polyzos, "Multicasting for multimedia applications," in *Proceedings of INFOCOM'92*, pp. 2078–2085, IEEE, 1992.

[6] C. A. Noronha and F. A. Tobagi, "Evaluation of multicast routing algorithms for multimedia streams," in *Proceedings of IEEE ITS'94*, (Rio de Janeiro, Brazil), IEEE, August 1994.

[7] R. Widyono, "The design and evaluation of routing algorithms for real-time channels," Tech. Rep. TR-94-024, University of California at Berkeley, International Computer Science Institute, June 1994.

[8] T. Billhartz, J. B. Cain, E. Farrey-Goudreau, and D. Fieg, "Performance and resource cost comparisons for the cbt and pim multicast routing protocols in dis environments," in *Proceedings of INFOCOM'96*, vol. 1, (San Francisco, CA, USA), pp. 85–93, IEEE, March 1996.

[9] D. Towsley and S. Mithal, "A selective repeat arq protocol for a point to multipoint channel," in *Proceedings of INFOCOM'87*, (San Francisco, CA, USA), pp. 521–526, IEEE, April 1987.

[10] D. Towsley, "An analysis of a point-to-multipoint channel using a go-back-n error control protocol," *IEEE Transactions on Communications*, vol. 33, pp. 282–285, March 1985.

[11] R. H. Deng, "Hybrid arq schemes for point-to-multipoint communication over nonstationary broadcast channels," *IEEE Transactions on Communications*, vol. 41, pp. 1379–1387, September 1993.

[12] N. Shacham and D. Towsley, "Resequencing delay and buffer occupancy in selective repeat arq with multiple receivers," in *Proceedings of INFOCOM'87*, pp. 512–520, IEEE, 1987.

[13] S. Pingali, D. Towsley, and J. 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*, (Santa Clara, CA, USA), pp. 221–230, ACM, May 1994.

[14] L. Kou, G. Markowsky, and L. Berman, "A fast algorithm for steiner trees," *Acta Informatica*, vol. 15, pp. 141–145, 1981.

[15] P. Winter, "Steiner problem in networks: a survey," in *Networks*, vol. 17, pp. 129–167, 1987.

[16] B. M. Waxman, "Routing of multipoint connections," *IEEE JSAC*, vol. 6, pp. 1617–1622, December 1988.

[17] J. Kadirire, "Minimising packet copies in multicast routing by exploiting geographic spread," *ACM Computer Communication Review*, vol. 24, pp. 47–62, July 1994.

[18] J. Nonnenmacher and E. W. Biersack, "Reliable multicast: Where to use fec," in *Proceedings of IFIP 5th International Workshop on Protocols for High Speed Networks (PfHSN'96)*, (INRIA, Sophia Antipolis, FRANCE), IFIP, Chapman & Hall, October 1996.

[19] J. Nonnenmacher, "Reliable multicast: An approximation for the expected number of retransmissions," tech. rep., Institut EURECOM, 2229 route des Crêtes, B.P. 193, 06904 Sophia Antipolis Cedex, FRANCE, May 1996.

[20] J. Nonnenmacher, "Ressource sharing in multicast connections," tech. rep., Institut EURECOM, 2229 route des Crêtes, B.P. 193, 06904 Sophia Antipolis Cedex, FRANCE, February 1996.