On the Stability of End-Point-Based Multimedia Streaming^{*}

György Dán, Viktória Fodor, and Gunnar Karlsson

Department of Signals, Sensors and Systems,
KTH, Royal Institute of Technology
{gyuri, vfodor, gk}@s3.kth.se

Abstract. In this paper we propose an analytical model of a resilient, tree-based end-node multicast streaming architecture that employs path diversity and forward error correction for improved resilience to node churns and packet losses. Using the model and via simulations we study the performance of this architecture in the presence of packet losses and dynamic node behavior. We show that the overlay can distribute data to nodes arbitrarily far away from the root of the trees as long as the loss probability is lower than a certain threshold, but the probability of packet reception suddenly drops to zero once this threshold is exceeded. The value of the threshold depends on the ratio of redundancy and on the number of the distribution trees. Using the model and simulations we show that correlated and inhomogeneous losses slightly worsen the overlay's performance. We apply the model to study the effects of dynamic node behavior and compare its results to simulations.

1 Introduction

The delivery of streaming media over end-point overlays has received much attention recently ([1, 2] and references therein). Although current commercial content delivery networks are capable of supporting many simultaneous streams, end-node-based multicast could considerably decrease the cost of large scale streaming, while being resilient to sudden surges in the client population, such as flash crowds. In an end-point-based multicast distribution system end-points are organized or organize themselves into an application layer overlay and distribute the data among themselves. The main advantages are that such a system is easy to deploy and it reduces the load of the content provider, since the distribution cost in terms of bandwidth and processing power is shared by the nodes of the overlay.

Since the success of such schemes depends on the behavior of the participating nodes, several issues have to be dealt with, such as the effects of group dynamics, stability of the system or the incentives for nodes to collaborate. Furthermore, since nodes receive data from their peer nodes only, the performance

^{*} This work has been supported in part by E-NEXT and by the Swedish Foundation for Strategic Research under the program Affordable Wireless Services and Infrastructures.

of such a scheme in an error prone environment is unclear due to possible error propagation.

The first proposed architectures focused primarily on low overhead due to control traffic and on the efficiency of the data distribution. They were based on a mesh [3] or a single distribution tree [4]. Resilience to node failures and error prone transmission paths appeared as important criteria later.

Robustness to node churns, i.e. node departures that disturb the data flow, was considered in SRMS [5] by distributing packets to randomly chosen neighbors outside of the distribution tree. Though this scheme provides some resilience to losses, it is known that repeating information is less efficient than using error correcting codes. SplitStream [6] and CoopNet [1] introduce multiple distribution trees and employ priority encoding transmission (PET) [7] based on forward error correction (FEC) [8] to decrease the effects of node failures and to recover from packet losses. Simulations were used to show the resilience of these schemes under various scenarios showing that increasing the number of trees improves the resilience both to packet losses and node churns.

Feasibility issues of small overlays with less than 100 nodes were discussed in detail in [9] based on experimental broadcasts over the Internet, and showed promising results. The experiments showed that poor performance was due in a large extent to packet losses. The feasibility of larger deployments was studied for a CoopNet like end-node overlay via simulations based on measured traces of user behavior in [2]. The authors concluded that application layer multicast architectures have enough resources, are stable in spite of group dynamics and hence can support large scale streaming content distribution.

Albeit there is an extensive literature on end-point-based multicast streaming, previous work on the behavior of these systems was limited to simulations. In this paper we present a simple model for a CoopNet like overlay combined with FEC, and evaluate the performance of such a system for a large number of nodes. We consider correlated and inhomogeneous losses and the effects of group dynamics. Our results show that an arbitrarily high packet reception probability can be achieved independent of the number of nodes in the overlay by adding enough redundancy (while keeping the bitrate constant). The packet reception probability goes however to zero if there is not enough redundancy added. The transition between the stable and non-stable states of the system is ungraceful, which can raise problems in a dynamic environment.

The paper is organized as follows. In Section 2 we briefly describe the architecture of the considered end-point-based application overlay for multicast. In Section 3 we present the mathematical model and the main results. In Section 4 we discuss the performance of the system based on the analytical model and simulations. In Section 5 we conclude our work.

2 System Description

We consider an application overlay as the one described in [1, 2] consisting of a root node and N peer nodes. Peer nodes are organized in t distribution trees,

either by a distributed protocol or a central entity like in [1]. The nodes are members of all t trees, and in each tree they have a different parent node from which they receive data and a different child node to which they forward data. Child nodes of the root node can have the same parent (i.e. the root) in more than one tree. Upon construction of the distribution trees each node is at the same distance from the root node in all trees, and we will refer to nodes at distance i nodes from the root as members of layer i . In the presence of group dynamics it is the task of the tree building algorithm to ensure that all parent nodes of a node are in the same or almost the same layer. We denote the number of children of the root node in each tree by m , and we call it the multiplicity of the root node. The number of layers in the distribution tree is N/m . Typically the number of distribution trees is no more than the multiplicity of the root node $m \geq t$; we will consider this case in the analysis. We assume that nodes do not contribute more bandwidth towards their children as they use to download from their parents, so that the multiplicity of the peer nodes is one, i.e. they have one child in each distribution tree (See Fig. 1).

The root uses block based FEC, e.g. Reed-Solomon codes, so that nodes can recover from packet losses due to network congestion and node departures. To every k packets of information c packets of redundant information are added resulting in a block length of $n = k + c$. If a source would like to increase the ratio of redundancy while maintaining its bi-

trate unchanged, then it has to decrease its source rate. We denote this FEC scheme by FEC(n,k). Using this FEC scheme one can implement UXP, PET or the MDC scheme considered in [1]. Lost packets can be reconstructed as long as no more than c packets are lost out of n packets. The root sends every t^{th} packet to its children in a given tree. If $n \leq t$ then at most one packet of a block is distributed over the same distribution tree. Peer nodes relay the packets upon reception to their respective child nodes in the tree corresponding to the particular packets, and once they received at least k packets of a block of n packets they recover the remaining c packets and send them to the child nodes in the corresponding distribution trees. A packet received from the parent node after it has been decoded based on other packets in the block will be discarded.

Fig. 1. Multicast tree structure for $t = 3$, $m = 3$ and $N = 6$

3 Mathematical Model

In this section we present a mathematical model that describes the behavior of the system in the presence of packet losses due to congestion in the network. Our goal is to calculate the probability $\pi(i)$ that a node in layer i of the distribution tree receives or can reconstruct an arbitrary packet, where i can be arbitrarily high. We model the correlated losses at the input-link of the nodes by a two-state

time-discrete Markovian model, often referred to as the Gilbert model [10]. We denote the probability that a packet is lost on the path between two adjacent peer nodes by p_ω ($0 < p_\omega < 1$). The probability that a packet is lost on the input link of a node given that the previous packet from the same block of packets was lost is denoted by $p_{\omega|\omega}$. The parameters p and q of the Gilbert model are calculated as $q = 1 - p_{\omega|\omega}$ and $p = \frac{p_\omega q}{1 - p_\omega}$. Based on the Gilbert model we can calculate the probability of losing l packets out of j consecutive packets denoted by $P(l, j)$ [10]. Losses seen by different nodes are assumed to be independent. We assume that the probability that a node is in possession of a packet is independent of the probability that another node in the same layer possesses a packet from the same block of packets. We will comment on the validity and possible effects of these assumptions later.

In the following we give a nonlinear recurrence equation [11] to calculate the evolution of $\pi(i)$. As the root node possesses all packets, the initial condition is

$$\pi(0) = 1. \quad (1)$$

Consider the n packets of an FEC block that should arrive from different parents to a node in layer $i + 1$. The average number of packets received or reconstructed at the node can be calculated as the average number of packets reconstructed given that j packets have been transmitted from the parents multiplied by the probability that the parents possess j out of the n packets. The probability that a node in layer $i + 1$ ($i \geq 0$) will possess a packet can then be calculated as

$$\pi(i + 1) = R(\pi(i), p_\omega, p_{\omega|\omega}) = \sum_{j=1}^n \binom{n}{j} \pi(i)^j (1 - \pi(i))^{n-j} \frac{1}{n} \sum_{l=1}^j \tau(l) P(j - l, j), \quad (2)$$

where $\tau(l)$ indicates the number of packets after FEC reconstruction if l packets have been received and is given as

$$\tau(l) = \begin{cases} l & 0 \leq l < k \\ n & k \leq l \leq n. \end{cases}$$

If losses occur independently on the input links of the nodes then $P(l, j) = \binom{j}{l} p_\omega^l (1 - p_\omega)^{j-l}$, and the model becomes the same as the one presented in [12] for independent losses.

We can rewrite (2) by subtracting $\pi(i)$ from both sides and omitting the indices to

$$f(\pi) = -\pi + \frac{1}{n} \sum_{j=1}^n \binom{n}{j} \pi^j (1 - \pi)^{n-j} \sum_{l=1}^j \tau(l) P(j - l, j). \quad (3)$$

Figs. 2 and 3 show examples of $f(\pi)$ for independent and correlated losses respectively. Since $\pi(i + 1) - \pi(i) = f(\pi(i))$ we have that for any layer i if $f(\pi(i)) < 0$ then $\pi(i + 1) < \pi(i)$, if $f(\pi(i)) > 0$ then $\pi(i + 1) > \pi(i)$ and if $f(\pi(i)) = 0$ then

Fig. 2. $f(\pi)$ vs. π for different ratios of redundancy and loss probabilities, independent losses ($p+q=1$). At the root in $(0,1)$ closest to 1 (if it exists) the derivative is negative and hence the corresponding fixed point is stable.

Fig. 3. $f(\pi)$ vs. π for different ratios of redundancy and loss probabilities, correlated losses. At the root in $(0,1)$ closest to 1 (if it exists) the derivative is negative and hence the corresponding fixed point is stable.

$\pi(i+1) = \pi(i)$ and $\pi(i)$ is a fixed point of (2). Starting with $\pi(0) = 1$ as in eq. (1) the value of $\pi(i)$ will decrease as long as $f(\pi(i)) < 0$. The roots of $f(\pi)$ correspond to the fixed points of eq. (2). If $f(\pi)$ has a real root r in the interval $(0,1)$ and the derivative $f^{(1)}(r) = \frac{d}{d\pi}f(\pi)|_{\pi=r} < 0$ then $\pi(\infty) = \lim_{i \rightarrow \infty} \pi(i) = r$ (e.g. r and r_2 in Figs. 2 and 3), since the fixed point corresponding to this root is asymptotically stable. A fixed point corresponding to a root r with $f^{(1)}(r) > 0$ is unstable on the other hand (e.g. r_1 in Figs. 2 and 3). If $f(\pi)$ does not have a real root in $(0,1)$ then $\pi(\infty) = 0$, since $f(\pi)$ is always negative on $(0,1)$ as we show it later (e.g. the dashed line in Figs. 2 and 3). We will call the system stable if $\pi(\infty) > 0$ and unstable otherwise. To check the existence and the number of real roots of $f(\pi)$ in $(0,1)$ we investigate the signs of $f(\pi)$ at the endpoints of the interval.

For any $p_\omega > 0$ the ratio of successfully received or recovered packets has to be less than 1, so that $f(1) < 0$. Since $\pi = 0$ is a zero of $f(\pi)$ we have to calculate $f^{(1)}(0)$ to see the sign of $f(0+) = \lim_{\pi \rightarrow 0+} f(\pi)$.

If $k = 1$ then we get that $f^{(1)}(0) = n(1 - P(1,1)) - k$, and thus if $P(1,1) = p_\omega < (n - k)/n$ then $f^{(1)}(0) > 0$ and consequently $f(0+) > 0$. Hence there has to be at least one root r in $(0,1)$ for which $f^{(1)}(r) < 0$ resulting in an asymptotically stable fixed point (e.g. the solid line in Figs. 2 and 3). It follows that for any loss probability p_ω there is a ratio of redundancy c/k above which $\lim_{i \rightarrow \infty} \pi(i) > 0$. Otherwise, if $p_\omega \geq (n - k)/n$, then the number of real roots in $(0,1)$ is either zero or an even number, and using Sturm's theorem [13] we find that the number of roots in $(0,1)$ is 0 for any such p_ω (e.g. the dashed line in Figs. 2 and 3).

If $k > 1$ then we have $f^{(1)}(0) = -P(1,1) = -p_\omega$, which is always negative, and thus the number of real roots in $(0,1)$ is either zero or an even number. By using Sturm's theorem we find for any p_ω and $p_{\omega|w}$ that the number of real

roots in $(0, 1)$ is no more than two (counting their multiplicity). If they exist, denoted by r_1 and r_2 ($r_1 \leq r_2$), then $f^{(1)}(r_2) < 0$, r_2 is asymptotically stable and $\pi(\infty) = r_2$ (e.g. the dotted line in Figs. 2 and 3). Since $f(\pi) > 0$ for $r_1 < \pi < r_2$, the above result holds for any $r_1 < \pi(0) \leq 1$ as initial condition. Similarly, even if $r_1 < \pi(i) < r_2$ for some i , we have $\pi(\infty) = r_2$. With other words, the system can recover from disturbances, as long as $\pi(i) > r_1$. Let us denote the bifurcation point in p_ω at which the asymptotically stable fixed point (r_2) annihilates with the unstable fixed point (r_1) and both disappear by $p_{max}(p_{\omega|\omega})$. For $0 < p_\omega < p_{max}(p_{\omega|\omega})$ we have that $0 < r_1 < r_2 < 1$. However, $f(\pi)$ has no roots in $(0, 1)$ for $p_\omega > p_{max}(p_{\omega|\omega})$ (e.g. the dash-dotted line in Figs. 2 and 3). In the special case when losses are uncorrelated, i.e. $p_{\omega|\omega} = p_\omega$, we will denote the bifurcation point in p_ω by p_{max} .

In the following we discuss the validity and effects of certain assumptions made in the model. The model considers loss correlations at the input links of the nodes. If losses occur in bursts at the output links of the nodes, the burstiness influences the results if packets from the same block are distributed over the same distribution tree, i.e. $t < n$, but does not influence them otherwise. We do not consider this case in this analysis. The assumptions $n \leq t$ and $m \geq t$ are made to ensure independence of the losses of packets in the same block and to ensure that each node has different parents in all of the trees respectively. Removing these assumptions will make losses more correlated, and hence worsen the performance of the distribution tree. On the other hand, setting $t > n$ will not improve the performance of the system compared to $t = n$. Hence, when choosing n and t there are two factors to be considered: the delay introduced by an FEC block of length n and the administrative overhead of maintaining t distribution trees.

4 Performance Evaluation

In this section we show results obtained with the model presented in the previous section and simulations. In all scenarios we set $t = n$ and we consider $m = 80$ to $m = 320$ for easy comparison. For the simulations we considered the streaming of a 112.8 kbps stream to nodes organized in 1000 layers, hence the number of nodes in the overlay is between 80000 and 320000. The packet size is 1410 bytes. The peer nodes have 128 kbps connections both uplink and downlink. Nodes choose their parent nodes at random, and avoid having the same parents in different trees whenever possible. During each run of the simulation the root node sends about $10000 \times m$ to $30000 \times m$ packets, so that there are 0.8 to 9.6 million packets sent per layer, and 0.8 to 9.6 billion packets sent in the overlay.

4.1 Independent and Homogeneous Losses

We start the evaluation by considering the simplest scenario, homogeneous, uncorrelated losses. All nodes experience the same packet loss probability, and packets arriving to a particular node are lost independent from each other.

Fig. 4. $\pi(1000)$ vs p_ω for different ratios of redundancy

Fig. 5. $\pi(i)$ vs. i for different ratios of redundancy and loss probabilities

Figure 4 shows $\pi(1000)$ as a function of p_ω for $k = 10$ and $k = 20$ and different values of c . The figure shows that for every (n, k) pair there is a loss probability p_{max} above which the reception probability in nodes far from the root node suddenly becomes 0. Below p_{max} the reception probability is close to 1 and is slowly decreasing. This stepwise, ungraceful decrease of the reception probability is an undesired feature for systems working in a dynamic environment such as the Internet. The figure shows that increasing the number of trees, i.e. the FEC block length, slightly improves the resilience of the distribution tree to losses, which is in accordance with [1, 8].

Figure 5 shows $\pi(i)$ as a function of i for different block lengths n and loss probabilities and a ratio of redundancy of $c/k = 0.2$. We see that $\pi(i)$ is close to one in the cases when $p_\omega < p_{max}$, while it becomes almost 0 after some i otherwise. The value of i at which $\pi(i)$ breaks down depends on how far p_ω is from p_{max} . For $k = 10$ $p_{max} = 0.0799$ and for $k = 20$ $p_{max} = 0.0885$. The positive effects of the increased block length can be seen by comparing results

Fig. 6. $\pi(i)$ vs. i for different ratios of redundancy and loss probabilities, simulation results

Fig. 7. Number of redundant packets vs. p_ω for different performance objectives

at $p_\omega = 0.08$, where for $k = 20$ the system is stable, whereas for $k = 10$ it is unstable. Figure 6 shows simulation results for the same scenarios as Fig. 5. The comparison shows perfect match for $p_\omega = 0.07$ and $p_\omega = 0.10$, while the simulation results show worse behavior than the analytical ones when p_ω is close to p_{max} . The difference is rather big for $k = 10$ and $p_\omega = 0.08$, when $p_\omega - p_{max} = 10^{-4}$, in which case deviations from the mean loss probability in the individual layers make the deterioration faster than that predicted by the model. For higher values of m the simulation gives more accurate results as the probability of deviation is lower due to the central limit theorem.

Figure 7 shows c , the number of redundant packets needed to ensure $\pi(\infty) > 0$, $\pi(\infty) > 0.99$ and $\pi(\infty) > 0.999$ for $k = 10$ and $k = 20$. The figure shows a closely linear relationship between the number of redundant packets needed and the loss probability. For low values of p_ω the number of redundant packets needed to ensure $\pi(\infty) > 0.999$ is close to the number of redundant packets needed for $\pi(\infty) > 0$, and hence in a dynamic environment the ratio of redundancy has to be set higher to prevent a severe decrease of $\pi(i)$ due to a sudden increase of the loss probability.

4.2 Inhomogeneous Losses

In this subsection we consider the scenario when the packet loss probability experienced by the individual nodes is not homogeneous, i.e. different nodes experience different loss probabilities. We denote by Q the joint distribution function of p_ω and $p_{\omega|\omega}$ experienced by individual nodes. If the multiplicity m of the root node is high, then the evolution of the packet reception probability can be described by the equation

$$\pi(i+1) = \int_Q R(\pi(i), p_\omega, p_{\omega|\omega}) dQ, \quad (4)$$

where $R(\pi(i), p_\omega, p_{\omega|\omega})$ was defined in eq. 2. This approximate model treats layer i as homogeneous, and calculates the mean of the packet reception probability in layer $i+1$ given the distribution Q of the packet loss probability. In the following we consider non-correlated losses ($p+q=1$) to keep the number of parameters low and we show results for two cases of inhomogeneous losses. In the first, bimodal case, γ_l portion of the nodes experiences loss probability p_ω^l and the rest p_ω^h , the mean packet loss probability in the overlay is $p_\omega = \gamma_l p_\omega^l + (1 - \gamma_l) p_\omega^h$. In the second, uniform case, the loss probabilities are uniformly distributed between p_ω^l and p_ω^h , the mean packet loss probability in the overlay is $p_\omega = (p_\omega^l + p_\omega^h)/2$. For easy comparison we consider $n = 12$, FEC(12,10) and $m = 160$. We do not consider $p_\omega^h > 0.1$, as nodes that experience higher loss probabilities will leave the overlay due to poor quality. Figure 8 shows results obtained with the model for various bimodal and uniform distributions. The figure shows that the presence of nodes with high loss probabilities decreases $\pi(i)$ compared to the homogeneous case. This is due to that R is a concave function of p_ω and π for values of interest of p_ω ($p_\omega < 0.1$). Simulation results shown in Figure 9 for

Fig. 8. $\pi(i)$ vs i for inhomogeneous losses

Fig. 9. $\pi(i)$ vs i for inhomogeneous losses. Simulation results.

$\pi(i)$ as a function of i show a good match with the mathematical model. Since the number of nodes in each layer is finite, the mean packet loss probability in individual layers can deviate from the mean loss probability in the overlay, which explains the high variance of the simulation results. The variance of the results decreases as m increases due to the central limit theorem.

4.3 Correlated Losses

Figure 10 shows $\pi(1000)$ as a function of p_ω for various values of $p_{\omega|\omega}$ obtained with the model. The figure shows that the value of both r_2 and $p_{max}(p_{\omega|\omega})$ decreases as losses become more correlated (i.e. $p_{\omega|\omega}$ increases). Figure 11 shows $\pi(i)$ as a function of p_ω for correlated losses for scenarios similar to those in Fig. 5 as obtained with the model. Comparing the two figures shows that correlated losses decrease the packet reception probability significantly whenever the system is stable. Figure 12 shows matching simulation results for the same scenarios.

Fig. 10. $\pi(1000)$ vs. p_ω for $m=80$, FEC (13,10) and FEC(26,20) and different values of $p_{\omega|\omega}$

Fig. 11. $\pi(i)$ vs i for different ratios of redundancy and loss probabilities, correlated losses

4.4 Malicious Layers

In this subsection we investigate how the presence of layers with extreme loss probabilities (e.g. a DDoS attack) influences the packet reception probability. We consider an overlay, where the loss probability is p_ω , except for layers 50 and 100, which experience significantly higher (p_ω^m) packet loss probability. Such layers decrease $\pi(i)$ below $r_2(p_\omega)$ (the stable fixed point corresponding to p_ω) in the layers following them. Based on Figs. 2 and 3 we expect that the overlay can recover as long as $\pi(i)$ remains larger than $r_1(p_\omega)$. Figure 13 shows $\pi(i)$ as a function of i for $p_\omega = 0.05$. The figure shows that for $m = 80$ the malicious layers influence the packet reception probability for all following layers already at $p_\omega^m = 0.25$, while for $m = 320$ the overlay is able to recover. The different behavior is due to that if less than $r_1(0.05) = 0.7499$ portion of the packets of an FEC block is received in the malicious layers, then that block will be entirely lost in the following layers. For $m = 320$ the probability that less than $r_1(0.05)$ portion of the packets in a block will be received in layers 50 or 100 is only significant for $p_\omega^m = 0.4$, while for $m = 80$ this probability is significant already for $p_\omega^m = 0.25$ (around 0.05). Hence, once again, increasing m improves the robustness of the overlay.

4.5 Effects of Group Dynamics

In this subsection we analyze the effects of node departures on the packet reception probability. We assume that the departure of a node interrupts the flow of data to its child nodes for a random time T . This time T includes the time it takes for the child node to notice that its parent node has departed and the time it takes to find a new parent node. For a description of how the departure of a parent or a child node can be detected see [1]. Several algorithms have been proposed to find a suitable parent node, a comparison of some simulation results is shown in [2].

In this work we consider an ideal parent selection algorithm that maintains the structure of the overlay despite of the node departures. Arriving nodes take the places of the departed nodes, and hence fill the gaps in the distribution tree. We consider the stationary state of the system, when the arrival and departure rates are equal. We assume that the interarrival times are exponentially distributed, this assumption is supported by several measurement studies [14, 15]. The distribution of the session holding times has been shown to fit the log-normal distribution [14].

Based on the model presented in Section 3 we expect that node departures can be included in the model as an increase of the packet loss probability as $p_\omega^d = N_d/N \times \overline{T}$, where N_d is the mean number of nodes departing per time unit and $\overline{T}$ is the mean of the time nodes need to recover from the departure of a parent node. The rationale for this hypothesis is that node departures can be treated as bursty losses on the output link of the departing nodes, and can be modeled as independent if $n \leq t$ and $m \geq t$.

To simulate the ideal tree construction algorithm, instead of removing the departing and inserting the arriving nodes, we switch nodes off after their session

Fig. 12. $\pi(i)$ vs i for different ratios of redundancy and loss probabilities, correlated losses. Simulation results.

Fig. 13. $\pi(i)$ vs i in the case of malicious layers. Simulation results.

Fig. 14. $\pi(i)$ vs i for $1/\mu = 306s$. Simulation results.

Fig. 15. $\pi(i)$ vs i for $1/\mu = 1320s$. Simulation results.

holding time has elapsed, and switch them on after a random time T , which would correspond to the reconstruction of the tree. We can change the value of p_{ω}^d by adjusting T and the session holding time $1/\mu$. We show simulation results for $m = 80$ and $m = 320$, and we consider two mean session holding times, $1/\mu = 306s$ as measured in [14] and $1/\mu = 1320s$ as measured in [2]. The parameters of the corresponding log-normal distributions are $M = 4.93$, $S = 1.26$ and $M = 5.46$, $S = 1.85$ respectively.

Figs. 14 and 15 show results for $1/\mu = 306s$ and $1/\mu = 1320s$ respectively considering FEC(12,10). The two figures show the same characteristics despite of the different mean session holding times, which supports our approximation. The figures show that for $p_{\omega}^d = 0.05$ the overlay is stable for both $m = 80$ and $m = 320$, for $p_{\omega}^d = 0.06$ the overlay is only stable for $m = 320$, while for higher values of p_{ω}^d it is unstable. The overlay would become stable for $p_{\omega}^d = 0.07$ by increasing the value of m similar to the results shown in Fig. 5 obtained with the analytical model. To understand why increasing the number of nodes per layer (m) gives better resilience to node departures, let us consider the evolution

of the number of active nodes per layer (ν). If ν/m in a layer is lower than $1 - p_{max}$ then it is likely that the following layers will not be able to recover the missing data. The evolution of ν can be modeled by an Engset system [16] (due to its insensitivity to the distribution of the service time, i.e. the nodes' lifetime distribution). ν follows a binomial distribution with parameters m and $\beta = \mu/(1 + \mu T)$, its mean is βm and its coefficient of variation (the ratio of the standard deviation and the mean) is $\sqrt{(1 - \beta)/(m\beta)}$. Consequently, the higher the number of nodes per layer, the lower the probability that the ratio of active nodes is lower than $1 - p_{max}$, which explains the improved performance of the overlay as m increases.

Node departures can have an aggravating effect towards the end of a broadcast, when departures cause the number of nodes in the overlay to decrease. The increased loss probability due to node churn might exceed the level of stable operation and can lead to $\pi(\infty) = 0$. The root node can prevent this from happening by increasing the ratio of redundancy c/k towards the end of the broadcast.

5 Conclusion

In this paper we presented a mathematical model for the analysis of an end-point overlay for multicast based on multiple distribution trees and forward error correction. We showed that for any loss probability there is a ratio of redundancy which ensures that even nodes far away from the root of the trees receive a non-zero ratio of the information. We showed that this multicast scheme shows a non-graceful performance degradation once the loss probability exceeds a certain threshold. The threshold depends on the number of distribution trees and the ratio of redundancy used. Using the model and simulations we showed that correlated and inhomogeneous losses decrease both the ratio of received packets and the value of the stability threshold of the system. We analyzed how malicious layers can influence the behavior of the system, and concluded that increasing the number of nodes per layer gives improved robustness. The performance evaluation in the presence of dynamic node behavior led to the same conclusion. The results presented here show that the ratio of redundancy has to be adjusted with care to maintain the stability of this overlay, as underestimating the loss probability in the network can lead to the loss of all data.

References

1. V. Padmanabhan, H. Wang, and P. Chou, "Resilient peer-to-peer streaming," in *Proc. of IEEE ICNP*, pp. 16–27, 2003.
2. K. Sripanidkulchai, A. Ganjam, B. Maggs, and H. Zhang, "The feasibility of supporting large-scale live streaming applications with dynamic application end-points," in *Proc. of ACM SIGCOMM*, pp. 107–120, 2004.
3. Y. Chu, S. Rao, S. Seshan, and H. Zhang, "A case for end system multicast," *IEEE J. Select. Areas Commun.*, vol. 20, no. 8, 2002.

4. S. Banerjee, B. Bhattacharjee, and C. Kommareddy, "Scalable application layer multicast," in *Proc. of ACM SIGCOMM*, 2002.
5. S. Banerjee, S. Lee, R. Braud, B. Bhattacharjee, and A. Srinivasan, "Scalable resilient media streaming," in *Proc. of NOSSDAV*, 2004.
6. M. Castro, P. Druschel, A. Kermarrec, A. Nandi, A. Rowstron, and A. Singh, "Splitstream: High bandwidth content distribution in a cooperative environment," in *Proc. of IPTPS*, 2003.
7. B. Lamparter, A. Albanese, M. Kalfane, and M. Luby, "PET - priority encoding transmission: A new, robust and efficient video broadcast technology," in *Proc. of ACM Multimedia*, 1995.
8. K. Kawahara, K. Kumazoe, T. Takine, and Y. Oie, "Forward error correction in ATM networks: An analysis of cell loss distribution in a block," in *Proc. of IEEE INFOCOM*, pp. 1150–1159, June 1994.
9. Y. Chu, A. Ganjam, T. Ng, S. Rao, K. Sripanidkulchai, J. Zhan, and H. Zhang, "Early experience with an Internet broadcast system based on overlay multicast," in *Proc. of USENIX*, 2004.
10. E. Elliott, "Estimates of error rates for codes on burst-noise channels," *Bell Syst. Tech. J.*, vol. 42, pp. 1977–1997, September 1963.
11. I. Gumowski and C. Mira, *Recurrences and Discrete Dynamic Systems, LNM-809*. Springer-Verlag, 1980.
12. G. Dán, V. Fodor, and G. Karlsson, "On the asymptotic behavior of end-point-based multimedia streaming," in *Proc. of International Zürich Seminar on Communication*, 2006.
13. S. Basu, R. Pollack, and M. Roy, *Algorithms in real algebraic geometry*. Springer Verlag, 2003.
14. E. Veloso, V. Almeida, W. Meira, A. Bestavros, and S. Jin, "A hierarchical characterization of a live streaming media workload," in *Proc. of ACM IMC*, pp. 117–130, 2002.
15. K. Sripanidkulchai, B. Maggs, and H. Zhang, "An analysis of live streaming workloads on the Internet," in *Proc. of ACM IMC*, pp. 41–54, 2004.
16. L. Kleinrock, *Queueing Systems*, vol. I. Wiley, New York, 1975.