Path Diversity Based Techniques for Resilient Overlay Multimedia Multicast

Matulya Bansal and Avideh Zakhor
Department of Electrical Engineering and Computer Science
University of California, Berkeley
Berkeley, CA 94720
{matulya, avz}@eecs.berkeley.edu

***Abstract*— Network Losses due to congestion and node failures severely impair the performance of multimedia streaming on the Internet. In previous work [14], we proposed a content distribution mechanism based on k-DAGs, a special Directed Acyclic Graph (DAG) with k parents per receiver, exploiting path-diversity and rate-adaptation to improve the performance of multicast multimedia streaming. In this paper, we propose and evaluate two new techniques. In the first scheme, each receiver streams an independent multimedia stream, such as the Forward Error Corrected (FEC) bit-stream or a Multiple Description Coded (MDC) description, from a parent, depending upon the existing error conditions in the network. This simplifies the receiver implementation as no packet-partitioning [5], [14] is required. In the second scheme, we construct multiple interior-node disjoint k-DAGs, and stream mutually exclusive stripes of content on each k-DAG. This improves tolerance to node-failures [1], [4], as any node is an interior node in at most one of the k-DAGs. Our results show that the proposed techniques are very effective in dealing with packet losses in the network, improving FEC goodput by 20-35% and MDC connectivity by 15-20%.**

I. INTRODUCTION

With the increasing user demand for high-bandwidth multimedia content, multimedia multicast is becoming more important for both content as well as network service providers. However, IP multicast has been unable to elegantly and adequately support the requirements of emerging applications such as streaming video multicast. This has led to the evolution of several application layer multicast solutions. In this paper, we present three overlay multicast techniques for improving the performance of Forward Error Corrected (FEC) and Multiple Description Coded (MDC) media under conditions of congestion and node failures in the network. The proposed techniques are based on our earlier work on k-DAGs, special Directed Acyclic Graphs (DAGs) characterized by the property that every receiver in a k-DAG has k parents [14].

In the first scheme, called Robust Overlay Multicast (ROM), a receiver streams simultaneously from multiple parents, thereby decorrelating losses and mitigating the effect of node failures. Further, for FEC media, rates are adapted to stream as much data as possible from parents experiencing the minimum loss. The second technique, Simple Parent Selection (SPS), trades-off some of the advantages of path-diversity for protocol simplicity. In case of SPS, a receiver streams an independent multimedia stream, such as the FEC coded bit-stream or a MDC description, from a parent. Parents to stream content from are selected on the basis of loss rates experienced by the receiver, and hence the scheme is referred to as Simple Parent Selection. The third technique, called Multiple Interior Node disjoint k-DAGs (MINK), constructs multiple interior node disjoint k-DAGs, and streams mutually exclusive stripes of content on each, thereby not only employing rate-adaptation and path-diversity, but also minimizing the effect of node failures by ensuring that every node is an interior node in at most one DAG.

The remainder of the paper is organized in the following manner. In Section II, we briefly present and characterize k-DAGs. In Section III, we present the three k-DAG based multimedia multicast schemes. In Section IV, we evaluate the performance of our schemes using NS-2 [15]. Finally, after presenting a brief overview of related work in Section V, we conclude in Section VI.

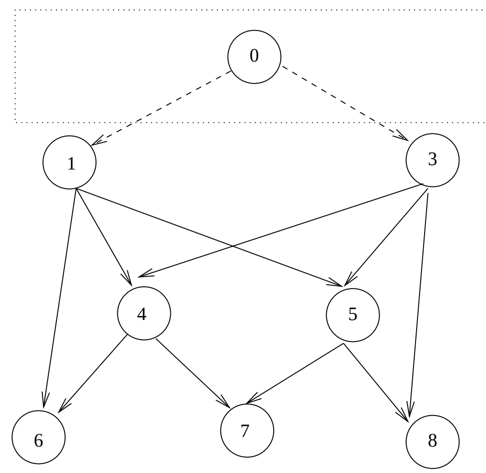

Fig. 1. A 2-DAG. Unlike a traditional multicast tree, every receiver in a k-DAG is connected to k parents, and therefore a k-DAG has k points of failures. The dotted line shows how to handle a single source situation. In this case, the first k nodes that join this source serve as the source nodes for building the k-DAG downstream from then on.

II. K-DAGS

In this section, we present the data-structure for constructing content distribution networks such that every receiver has k senders. Since there are multiple parents per node, our proposed distribution network is no longer a tree, but a DAG. Figure 1 shows an example of a 2-DAG.

As in [14], we define the $level$ of a node as the length in hops of the longest path from the source to this node. The level of the source nodes is defined to be 0. We say that a DAG has dag_degree k if all its nodes have k parents. Clearly, $k = 1$ results in traditional multicast trees with one parent per receiver. We define a distribution network to be $k - independent$ if any node in this network is still connected to the root of the DAG by at least one path in the event of up to $k - 1$ node failures. It can be shown that by simply starting with k source nodes and ensuring that every new node has k parents, no matter which these parents are, the resulting DAG is k-independent. Intuitively, k-independence is desirable because a k-independent distribution network has k points of failure, compared to a tree with a single point of failure. We define the outgoing capacity of a node as the ratio between the upstream bandwidth of this node and the maximum bit-rate of the multimedia content being streamed. This is also simply referred to as capacity c, and is different from the incoming capacity, which is the total bandwidth a node needs to stream multimedia content from its parents including that required for rate adaptation. Also, as in [14], we define $frac\_bw\_per\_parent$ to be the ratio between the maximum bit-rate a receiver can stream data from any of its parents, and the total bit-stream rate of that piece of content. For the receiver to receive the entire multimedia-stream, we require that $k * b \geq 1$. For $k * b > 1$, it is possible to adjust the streaming rate among the parents so as to adapt to network conditions. Again, to ensure fairness to the competing Simple Overlay Multicast (SOM) scheme, we assume that receivers proactively $reserve$ the maximum bandwidth they might need, i.e. b * bit-stream rate, from any parent for the entire duration of a session. For MDC media, since we stream each description from a different parent, $frac\_bw\_per\_parent$ is defined as the ratio between the bit-rate of a description and the aggregate bit-rate of all the descriptions in the stream. All MDC descriptions are assumed to have almost the same bit-rate.

The degree of the DAG is an important parameter of our proposed scheme. Intuitively, larger $dag\_degrees$ result in greater flexibility in adapting to losses. However, assuming that the upstream capacity for all nodes stays the same, distribution networks using k-DAGs are likely to become deeper, and hence experience increased delay as $k * b$ becomes large. Assuming that we build r such interior node disjoint k-DAGs, and stream a fraction f of the total content on each, such that $r * f \geq 1.0$ and each k-DAG has N/r interior nodes, the minimum depth of the DAG needed to support N nodes can be shown to be

$$\frac{log(\frac{Nr}{p} * (\frac{c}{b*k*f} - 1) + 1)}{log(\frac{c}{b*k*f})} - 1. \qquad (1)$$

where p is the number of nodes at the level 0 with a total outgoing capacity of $p*c$ [6]. Similarly, the amortized loss rate across all nodes for a k-DAG of depth l, containing the maximum number of nodes that it can support, is

$$1 - \frac{(1-x) * (\frac{(\frac{c*(1-x)}{b*k*f})^l - 1}{\frac{c*(1-x)}{b*k*f} - 1})}{(\frac{(\frac{c}{b*k*f})^l - 1}{\frac{c}{b*k*f} - 1})}. \qquad (2)$$

where x is the average loss rate per link [6].

Note that a necessary condition for feasibility of construction of such r interior node disjoint k-DAG is that $c > b * k * f * r$, or intuitively, the outgoing capacity of each node should be greater than its incoming capacity.

Based on Equations 1 and 2, Figure 2 shows how amortized loss experienced over nodes varies as a function of number of nodes for different values of k, b and r with $c = 2$, $f = 1/r$, $p = 1$ and $x = 0.05$. We observe that for a given number of nodes, the depth and the amortized loss experienced by the receiver nodes in a k-DAG increases with $k * b$. However, depth and amortized loss decrease with r. Intuitively, this happens because as observed in [1], [4], by building interior node disjoint distribution networks, we are utilizing the capacity of what would otherwise be leaf nodes with unused capacity in a simple multicast tree. This reduces depth and hence delay and loss.

In case of a k-DAG where the upstream bandwidth reserved per

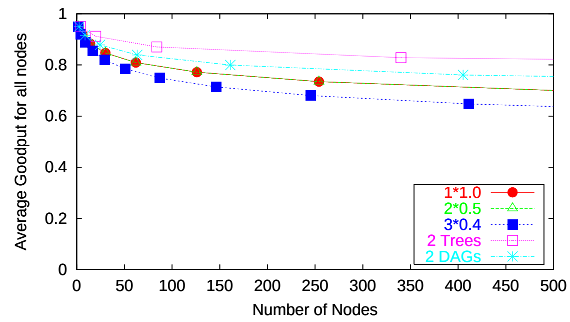

Fig. 2. This figure shows how average goodput of k-DAGs vary with k, b and r. The key uses a k*b notation to denote schemes with different k and b values. 2 trees and 2 DAGs refer to two interior node disjoint trees and DAGs respectively. For 2 DAGs, each of the DAGs uses k=4 and b=0.4. r*f=1.0 for both 2 trees and 2 DAGs.

sender is b, up to s node failures can be successfully tolerated, where

$$s = \max_r (k - r) * b \geq 1. \qquad (3)$$

Further, for a (l, m) FEC coding with l data packets protected by m redundant packets, assuming no other losses elsewhere, the tolerable number of node failures, s, becomes

$$s = \max_r (k - r) * b \geq \frac{l}{l + m}. \qquad (4)$$

The degree of k-DAG lends a natural interpretation in context of MDC media. If we construct a k-DAG for MDC media with k descriptions and enforce that each receiver streams a description from a distinct parent, by virtue of the independence property of k-DAGs, a node is guaranteed to receive at least one description for up to $k - 1$ node failures. If k is greater than the number of descriptions, the receiver can choose parents to stream independent descriptions from depending on the network conditions. In Section IV, we show NS-2 [15] simulation results for node failures using both FEC coded and MDC media.

III. Proposed Techniques

In this section we present three techniques, namely ROM, SPS and MINK, based on this content distribution mechanism for resilient multimedia multicast streaming. The techniques vary in both complexity and performance, with SPS being the simplest and MINK the most complex.

A. Robust Overlay Multicast (ROM)

In Robust Overlay Multicast (ROM), we construct a k-DAG and stream simultaneously from multiple parents [14]. The k-DAG construction and maintenance algorithms used by ROM are discussed in [14]. For FEC coded media, ROM adapts rates using a greedy bucket-filling rate-allocation algorithm and a simple receiver-based packet-partitioning and synchronization algorithm, discussed in detail in [14], [5]. For MDC media with k descriptions, a k-DAG is constructed and each receiver streams a distinct description from a distinct parent, thereby reducing its chances of losing multicast session connectivity upon node failures.

B. Simple Parent Selection (SPS)

Simple Parent Selection (SPS) reduces the complexity involved in rate-allocation and packet-partitioning at the expense of path-diversity. In SPS, instead of simultaneously streaming from multiple parents, a receiver chooses the parents from which it experiences

TABLE I
COMPARATIVE DESCRIPTION AND LEGEND

	ROM	SPS	MINK				MINT	
			MINK-ROM		MINK-SPS			
			Style 1	Style 2	Style 1	Style 2	Style 1	Style 2
FEC	ROM-k-b	SPS-k-1	MINK-ROM-k-b-1-r	MINK-ROM-k-b-2-r	MINK-SPS-k-1-r	MINK-SPS-k-2-r	MINT-1-r	MINT-2-r
	stream from multiple parents, adapt rates	stream from one of k parents	construct r k-DAGs, each using ROM-k-b, striping across k-DAGs	construct r k-DAGs, each using ROM-k-b, striping along k-DAGs	construct r k-DAGs, each using SPS-k-1, striping across k-DAGs	construct r k-DAGs, each using SPS-k-1, striping along k-DAGs	construct r trees, striping across trees	construct r trees, striping along trees
MDC	ROM-k	SPS-k-t	MINK-ROM-r		MINK-SPS-k-r		MINT-r	
	t = k, stream each description from a separate parent	k $\geq$ t, choose t out of k parents to receive one description from each	r DAGs, each using ROM-1		r DAGs, each using SPS-k-1		r trees, each using ROM-1	

minimum loss and streams independent stripes of content from each. For rate-allocation using FEC coded media, the receiver chooses one of its k parents in the k-DAG to stream data from. For rate-allocation using MDC media with t descriptions on a k-DAG, such that $t \leq k$, a receiver chooses t of its k parents from which it experiences minimum losses, to stream an independent description. Since an independent stream is streamed from each parent, no packet-partitioning is required for SPS. SPS uses the same k-DAG construction and maintenance algorithms as ROM. Since a receiver might not be streaming from some parents at any given time, loss rates to these parents are estimated out of band, i.e. by periodically sending a train of probe packets and measuring the observed loss rate.

C. Multiple Interior Node disjoint k-DAGs (MINK)

Multiple Interior Node Disjoint k-DAGs (MINK) is motivated by the observation first made in [1] and subsequently adopted in [4]. The underlying idea is to stream on multiple interior node disjoint trees (MINT), and thereby improve the tolerance of the scheme to node failures. We apply the same idea to streaming on multiple interior node disjoint k-DAGs. Since we construct k-DAGs rather than trees, both ROM and SPS can be used for streaming multimedia over individual k-DAGs. As such, there are two classes of MINK, namely, MINK-ROM and MINK-SPS.

Our multiple interior node disjoint k-DAG construction algorithm is similar to the interior node disjoint tree construction algorithm proposed in [4] with two important differences. First, instead of constructing trees, we construct k-DAGs and second, as during k-DAG construction and maintenance phases, instead of using a purely centralized or distributed mechanism, we utilize the indirection provided to us through DHT nodes [14].

As in [4], a node upon joining the multicast session, is assigned a k-DAG in which it is $active$, i.e. the k-DAG in which this node will be an interior node and hence support multimedia streaming. However, in case of MINK, as opposed to a central server, this assignment is done by the DHT nodes, which keep track of the number of active nodes in each k-DAG. In all other k-DAGs, this node will be a $passive$ node, i.e. only receive but not stream multimedia content. To join its active k-DAG, a node, say A, contacts the DHT to find nodes with available capacity in this k-DAG. In case there are no such nodes, a parent of k of the passive nodes in this k-DAG is aborted, node A joins the parents of these aborted nodes, and accepts the aborted nodes as its children. Assuming that the outgoing capacity of each node is greater than or equal to its incoming capacity, node A should be able to support these k nodes. For minimizing delay though, the orphaned nodes may choose to request the DHT nodes for new potential parents in this k-DAG. To join a passive k-DAG, node A again requests the DHT for k nodes with available capacity, and joins these as its parents. The DHT nodes attempt to maintain an approximately equal number of $active$ nodes in each of the k-DAGs by asking an incoming node to be active in the k-DAG with the minimum number of active nodes. In case of imbalances due to node departures or failures or capacity constraints, active nodes are reactively moved around to restore an approximately equal number of active nodes in each k-DAG.

For streaming FEC coded media with MINK, we evaluate two different techniques, referred to as style 1 and style 2 of striping. In the first technique, FEC blocks are split $across$ k-DAGs. Assuming a (n, d) FEC coding and r k-DAGs, $\lfloor (n+d)/r \rfloor$ packets are assigned to be sent on each k-DAG. If $n+d-(\lfloor (n+d)/r \rfloor * r) > 0$, for every FEC block an additional packet is sent over $n+d-(\lfloor (n+d)/r \rfloor * r)$ k-DAGs. For example, using a (21,7) FEC coding and 2 k-DAGs, 14 packets in each FEC block are transmitted on each k-DAG. If a (21, 8) FEC coding were to be used, 15 packets in a FEC block would be sent over one k-DAG and 14 over the other. In the second technique, FEC blocks are split $along$ the k-DAGs, and the i^{th} FEC block is transmitted on the $(i\ mod\ r)^{th}$ k-DAG, where r denotes the number of constructed k-DAGs and k-DAGs are numbered from $0, 1, ..., r-1$. For example, assuming that 2 k-DAGs are constructed, FEC blocks 1, 3, 5, ... are sent on one k-DAG while FEC blocks 2, 4, 6, ... are sent on the other. While data being transmitted on different k-DAGs is coupled by the FEC coding in the first case, in the second case, data transmitted on each k-DAG is independent at the FEC coding level. For MINK using k-DAGs, where $k \geq 2$, ROM-like rate-allocation and packet-partitioning is used by each receiver to receive the stripe for that k-DAG from its parents. For streaming MDC media with MINK, we stream a distinct description on each k-DAG. If $k \geq 2$, then the parent from which a receiver experiences the minimum loss is chosen for streaming the description being streamed on this k-DAG.

IV. PERFORMANCE EVALUATION

In this section, we compare and evaluate the performance of the proposed techniques using NS-2 [15] simulations. Due to space constraints, we present results only for the node failure experiments. Results for congestion loss experiments follow similar trends. The notation used for various k-DAG constructs for both FEC and MDC media is shown in Table I.

A. Simulation Parameters

We generate Albert-Barabasi model topologies for our simulations using the BRITE[16] topology generator. We use 500 network nodes

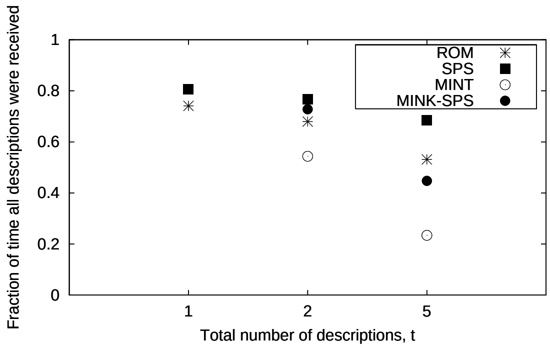

Fig. 3. Fraction of time all descriptions are received versus number of descriptions received.

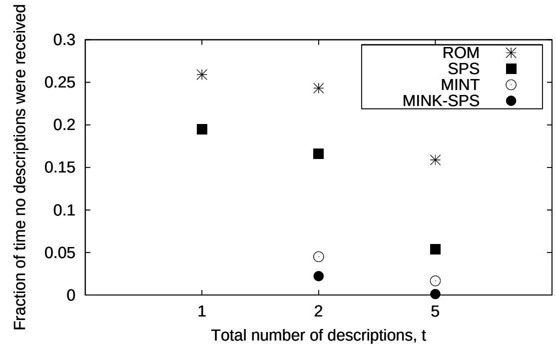

Fig. 4. Fraction of time no descriptions are received versus number of descriptions received.

and 250 overlay nodes and all the overlay nodes are multicast receivers. Only one DHT node is used. The overlay nodes and their order of joining are both randomly chosen for each simulation. We use 7 source nodes for all the simulations. The bit-stream rate is 128Kbps and the packet size is 512 bytes. A (21,7) FEC code is used to protect data. To simulate the effect of increased aggregate bandwidth required to transmit multimedia content at similar Peak Signal to Noise Ratio (PSNR) values when multiple descriptions are used, we assume an overhead of 20% over the bit-stream with one description for every additional description used. The outgoing capacity of all the nodes, unless otherwise mentioned, is fixed to be twice as large as the multimedia bit-stream rate. Loss rates are updated, and the decision to re-allocate rates is evaluated every 5 seconds. Every point in a plot corresponds to an average of 50 repetitions. The source is started at the beginning of the simulation and all the overlay nodes join the multicast session within the first 4-5 minutes. Losses are introduced once all the nodes have joined the session. We randomly fail 20 percent of the nodes within an interval of 60 seconds and then record the performance of the live multicast nodes over this period.

B. Metrics

We evaluate the effectiveness of the proposed schemes using FEC goodput, delay, jitter and control overhead as metrics. FEC goodput is defined as the ratio between the number of received non-redundant FEC blocks to the number of expected non-redundant FEC blocks. Average delay is the average time it takes for a data packet sent by the source to reach a receiver, averaged over all the packets received by a receiver, and all receivers. Jitter is defined as the standard deviation in the delay experienced by a receiver. Average jitter is jitter averaged across all the receivers. Control overhead is defined as the ratio between the number of control packets sent and the data packets received. Normalized node ranks are obtained by dividing the rank of a node by the maximum rank of any receiver. For a set of n receivers, ranks are calculated by sorting the list of metric values for receivers and assigning a number from 1 to n to elements in the sorted list.

C. Node Failures with MDC Media

In these experiments, we evaluate the robustness of our schemes in delivering MDC media with node failures in the network. In order to facilitate DAG construction using 5 descriptions, whence increased bandwidth is required, we fix the outgoing capacity of nodes in these experiments to be thrice the bit-stream rate using a single description.

Figures 3 and 4 show the fraction of time all and no descriptions are received respectively for a variety of schemes. As seen, for all schemes, the fraction of time more descriptions are received goes down as a function of the number of descriptions, due to increased bandwidth requirement of MDC. This is counterbalanced by observing that the fraction of time no description is received decreases with the number of descriptions. For both MINT and MINK, the fraction of time when no description is received is significantly lower than ROM and SPS. Intuitively, this can be explained by considering that any given node is an interior node in only one tree and a leaf node in all the others. So, while the chance of the node losing a description in every tree is low, as any node failure affects only interior of one tree. Finally, as seen, MINK has fewer blackouts than MINT by virtue of rate-adaptation.

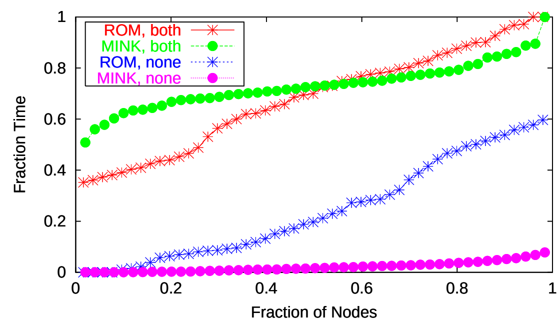

Fig. 5. Percentage of time when both and neither of the two descriptions are received using a 2 description MDC coding using ROM-2 and MINK-SPS-2-2.

Figure 5 shows the distribution of fraction of time both and no descriptions are received across nodes when 2 description MDC media is streamed using MINK-SPS-2-2 and ROM-2. As seen, the MINK curves are flat, indicating that most nodes achieve similar performance, while curves for ROM slope considerably, indicating that nodes near the source perform much better than the nodes towards the bottom of the distribution network. As observed earlier, fraction of time no description is received is substantially lower for MINK, while the fraction of time both descriptions are received are comparable.

Table II summarizes the connectivity, delay and control overhead information for different schemes. We observe that MINK provides the best connectivity, and results in minimum blackouts. Also, while a larger number of descriptions, each being streamed from a distinct parent improves connectivity, it also results in increased delay because of the greater aggregate bandwidth requirements with multiple descriptions. ROM and SPS incur higher delay than SOM.

TABLE II
MDC CONNECTIVITY, DELAY AND CONTROL OVERHEAD

Scheme	SOM	ROM-5	SPS-6-5	MINT-5	MINK-SPS-2-5
Blackouts	0.2591	0.1586	0.0535	0.0167	0.0012
Delay	0.2476	0.4742	0.6563	0.1941	0.2539
Control	0.0113	0.0306	0.0457	0.0254	0.0541

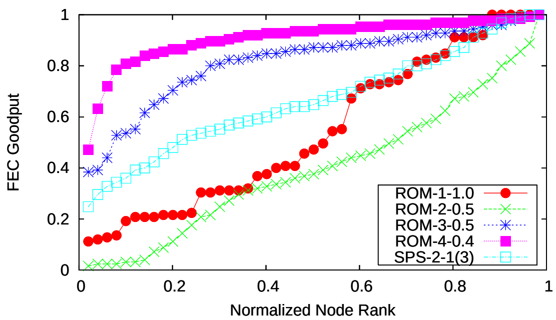

Fig. 6. FEC goodput distribution across nodes for variants of ROM and SPS.

However, MINK and MINT, by utilizing the capacity of all the nodes, outperform SOM. Control overhead is higher for ROM, SPS, and MINK; however, it is reasonable, i.e. less than 6% in all cases, and therefore negligible.

D. Node Failures with FEC Coded Media

In these experiments, we characterize the performance of our schemes in delivering FEC coded media in presence of node failures. To enable k-DAG construction with SPS, we set the outgoing capacity of nodes for SPS simulations to be thrice the bit-stream rate.

Figure 6 shows the FEC goodput distribution across nodes for variants of ROM, and SPS. We observe that ROM-4-0.4, followed by ROM-3-0.5, outperform the rest of the schemes, due to greater path-diversity and ability to adapt. SPS-2-1 performs poorly as compared to ROM-4-0.4 and ROM-3-0.5 because of the coarse granularity of adaptation of SPS for FEC coded media, as it can either receive the entire bit-stream from a parent or none. As such, SPS needs to reserve more upstream bandwidth than ROM for adapting rates, and results in deeper k-DAGs. Contrary to expectation, ROM-2-0.5, performs worse than ROM-1-1.0, because upon failing, every node passes off half of bit-stream loss to twice the number of nodes, and consequently a (21,7) FEC coding results in FEC blocks being lost at twice the number of nodes.

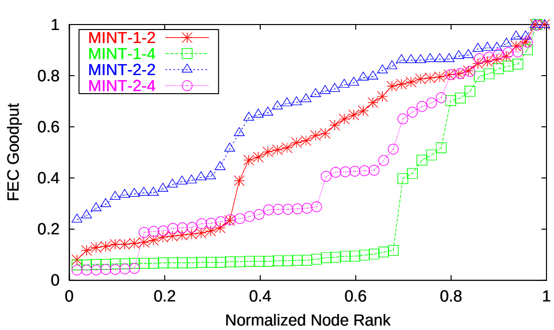

Fig. 7. FEC Goodput distribution across nodes for MINT using striping styles 1 and 2.

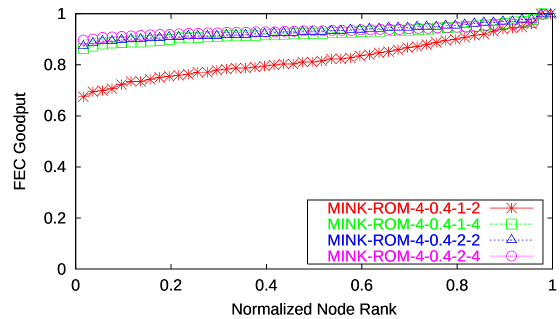

Fig. 8. FEC goodput distribution across nodes for MINK using striping styles 1 and 2.

TABLE III
FEC GOODPUT, DELAY, JITTER AND CONTROL OVERHEAD

Scheme	SOM	ROM-4-0.4	MINT-2-2	MINK-ROM -4-0.4-2-2
FEC goodput	0.5408	0.9040	0.6475	0.9341
Delay	0.3118	0.5807	0.1917	0.2174
Jitter	0.0006	0.0751	0.0286	0.0679
Control	0.0126	0.0317	0.0246	0.0567

Figure 7 shows FEC goodput distribution across nodes using MINT-1-2, MINT-2-2, MINT-1-4 and MINT-2-4. As seen, the overall performance is poor, sometimes even worse than SOM, shown in Figure 6. This happens because every node is a leaf node in all but one of the trees. Therefore in absence of any possibility of rate-adaptation, FEC goodput for style 1 of striping gets adversely affected even if an upstream node fails in any one of the trees. For style 2, since independent blocks are transmitted on different trees, FEC goodput is better, but still poor as compared to ROM and its variants. MINT with 2 trees outperforms 4 trees. This can be explained considering that in case of 4 trees, any node is a leaf node in 3 trees rather than the 2 tree case, where any node is a leaf node in only one tree. Loss compounds with depth, and without any rate-adaptation, leaf nodes suffer from greater loss.

Figure 8 shows FEC goodput distribution across nodes using MINK-ROM-4-0.4-1-2, MINK-ROM-4-0.4-2-2, MINK-ROM-4-0.4-1-4 and MINK-ROM-4-0.4-2-4. Taking into account Figures 6 and 7, we observe that MINK outperforms MINT, ROM and SPS across nodes. Again, while both variants perform well, as explained above, striping style 2 performs better than striping style 1.

Table III lists the average values for observed average FEC goodput, delay, jitter and control overhead over nodes. As seen, MINK achieves the highest FEC goodput, followed by ROM and SPS in that order. MINT performs poorly as compared to MINK, ROM and SPS, with striping style 1 resulting in FEC goodput values even worse than SOM. Variants of ROM and SPS that reserve more upstream bandwidth result in greater average delay, while MINK and MINT, by utilizing the spare capacity of leaf nodes, achieve delay even less than that of SOM. Both control and jitter increase as data is streamed from multiple parents; however, both are reasonable.

V. RELATED WORK

Multimedia Streaming using Overlay Networks has been an active area of research and as such, several interesting overlay multicast schemes have been proposed. In this section, we present a brief overview of existing related work and compare our work with them.

Narada [3] builds a dynamic DVRMP style overlay multicast tree for video conferencing. Even though it adapts the overlay topology to changing network conditions, every receiver is essentially connected to a single parent. Splitstream [1] addresses the issue of handling node failures by building multiple multicast trees such that any node is an interior node in at most one of the trees. Each tree is used to stream an independent MDC description. Co-Op Net [2], [4] is similar in spirit to Splitstream. It too builds multiple multicast trees each for streaming a single MDC description. However, Co-Op Net is centralized with tree management operations based at the source rather than being peer-to-peer. Our schemes are different because unlike multiple interior node disjoint trees, they rate-adapt to stream FEC coded streaming multimedia effectively as well. Further, for MDC media, use of rate-adaptation, as in MINK, improves the percentage of time more descriptions are received at the nodes.

In terms of motivation, distributed video streaming (DVS) [5], [7], [9], [10], [11], [12], [13] comes the closest to our work. DVS uses both path diversity and rate-adaptation for streaming video content from multiple sources to a receiver. In [7], the effect of path diversity on file transfers involving TCP is examined. By sending acknowledgments, TCP headers and TCP packets on different paths in different proposed variants of TCP, it is shown that by dynamically switching the forward data path based on the observed delay, file transmission times can be shortened. In [8], the authors use path-diversity to minimize the rate-distortion observed at the receiver. By dynamically scheduling packets to be sent over different paths from the sender based on the paths' delay characteristics, end-to-end distortion is minimized. In [5], [11], the authors show how Forward Error Coding coded video content can be transmitted over multiple paths on the Internet to mitigate losses. In [12], [13], [10], the authors propose the use of streaming multiple descriptions on multiple paths to reduce video distortion. In [9], the authors propose the use of multistream video coding over multiple paths for mitigating the effect of errors in ad hoc networks. By using k-DAGs, we have extended the idea of exploiting path-diversity to multicast.

VI. Conclusion

In this paper, we have proposed three techniques for improving the resilience of multimedia multicast to loss in the network. The main observations to be made from the simulation results are

1) Larger rate adaptivity achieved via larger values of $k * b$ could potentially help in combating loss and node failures.
2) At the same time, larger $k * b$ values increase the depth of the resulting DAG.
3) Shallow DAGs have lower delay and better loss and node failure resiliency than deeper ones.
4) Multiple Interior Node Disjoint schemes utilize upstream bandwidth of what would otherwise be leaf nodes and reduce the depth of the resulting DAGs.

Combining (1), (2) and (3), we conclude that k*b values must be chosen large enough to allow for rate-adaptation, and yet small enough to avoid deep DAGs. Furthermore, taking into account the simulation results in Section IV and observation (4) above, we conclude that overall MINK-ROM results in best performance for FEC goodput, and MINK-SPS for MDC connectivity, under loss and node failure conditions. This is because MINK-ROM and MINK-SPS enjoy not only the rate adaptability features of ROM and SPS for FEC and MDC media respectively, but also the shallow constructs of MINK structure.

Specifically, for FEC media, we observe that MINK-ROM improves FEC goodput by up to 20-35% as compared to traditional overlay multicast, and for MDC media, MINK-SPS reduces the occurrence of blackouts by up to 15-20%. In addition, MINK ensures that performance gains are achieved across nodes, and not at the expense of nodes far away from the source. This performance improvement thereby accrues at the cost of greater jitter values and control overhead only. With increasing end system buffering capabilities, jitter is no longer an issue, except for live applications; meanwhile the extra control overhead is justifiable in context of the performance gains achieved.

References

[1] M. Castro, P. Druschel, A. Kermarrec, A. Nandi, A. Rowstron and A.Singh, "Splitstream: High-Bandwidth multicast in a cooperative environment", *SOSP'03*, Lake Bolton, New York, October 2003.

[2] V. N. Padmanabhan, H. J. Wang, P. A. Chou and K. Sripanidkulchai, "Distributing Streaming Media Content using Co-operative Networking", *ACM NOSSDAV*, Miami Beach, Florida, May 2002.

[3] Y. Chu, S. Rao and H. Zhang, "A case for end system multicast", *In Proc. of ACM Sigmetrics*, pp 1-12, June 2000.

[4] V. N. Padmanabhan, H. J. Wang and P. A. Chou, "Resilient Peer-to-Peer Streaming", *In IEEE ICNP*, Atlanta, GA, November 2003.

[5] T. Nguyen and A. Zakhor, "Distributed Video Streaming Over the Internet", *SPIE*, San Jose, California, January 2002.

[6] M. Bansal, "Resilient Overlay Multimedia Multicast", MS report, Department of EECS, UC Berkeley, 2004.

[7] E. Steinbach, Y. J. Liang and B. Girod, "A Simulation Study of Packet Path Diversity for TCP File Transfer and Media Transport over the Internet", In *Proceedings of International Conference on Digital Communications*, Capri, Italy, September 2002.

[8] J. Chakareski and B. Girod, "Rate-distortion optimized packet scheduling and routing for media streaming with path diversity", In *Proceedings of IEEE Data Compression Conference*, Snowbird, Utah, April 2003.

[9] S. Mao, S. Lin, S. S. Panwar, Y. Wang and E. Celebi, "Video Transport over Ad hoc Networks: Multi-stream coding with multiple transport", In *IEEE Journal on Selected Areas in Communication*, Special Issue on Recent Advances in Wireless Multimedia, vol. 21, no. 10, pp 1721-1737, December 2003.

[10] J. Kim, R. M. Nersereau and Y. Altunbasak, "Network-adaptive video streaming using multiple description coding and path diversity", In *IEEE International Conference on Multimedia and Expo*, July 2003.

[11] T. Nguyen and A. Zakhor, "Path Diversity with Forward Error Correction (PDF) System for Packet Switched Networks", *IEEE INFOCOM 2003*.

[12] J. Apostolopoulos, "Reliable video communication over lossy packet networks using multiple state encoding and path diversity", *In Proceedings of The International Society for Optical Engineering*, January 2001, vol. 4310, pp. 392-409.

[13] J. Apostolopoulos, T. Wong, W. Tan and S. J. Lee, "On Multiple Description Streaming with Content Delivery Networks", *IEEE Infocom 2002*, June 2002, vol. 4310.

[14] M. Bansal and A. Zakhor, "Path Diversity for Overlay Multicast Streaming", submitted to *International Packet Video Workshop 2004*.

[15] K. Fall, "NS Notes and Documentation: The VINT Project", 2000.

[16] A. Medina, A. Lakhina, I. Matta and J. Byers, "BRITE: An Approach to Universal Topology Generation", In Proceedings of *MASCOTS 2001*, Cincinnati, Ohio, August 2001.

[17] J. Byers, J. Considine, M. Mitzenmacher and S. Rost, "Informed Content Delivery Across Adaptive Overlay Networks", in *Proceedings of SIGCOMM 2002*, Pittsburg, Pennysylvania, August 2002, pp. 47-60.

[18] Y, J. Liang, J. G. Apostolopoulos and B. Girod, "Analysis of Packet Loss for Compressed Video: Does Burst Length Matter?", in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, Vol. 5, pp. 684-687, April 2003.

[19] S. Banerjee, C. Kommareddy, K. Kar, B. Bhattacharjee and S. Khuller, "Construction of an Efficient Overlay Multicast Infrastructure for Real-Time Applications", in *Proceedings of IEEE INFOCOM 2003*, vol. 2, pp. 1521-1531, San Francisco, CA, March/April 2003.

[20] X. R. Xu, A. C. Myers, H. Zhang and R. Yavatkar, "Resilient Multicast Support for Continuous-Media Applications", in *Proceedings of IEEE NOSSDAV*, pp. 183-194, 1997.

[21] Y. Wang and Q. Zhu, "Error Control and Concealment for Video Communication", in *Proceedings of the IEEE*, Vol. 86, No. 5, pp 974-997, May 1998.