Smart Multipath Capacity Aggregation for High Quality Video Streaming

Omesh Tickoo* *Member, IEEE*, and Shivkumar Kalyanaraman *Member, IEEE*

***Abstract*— We present an end-to-end scheme to aggregate a diversity of wired or wireless network paths in an application content-aware manner. The objective is to realize a virtual pipe abstraction that offers high capacity, low perceived jitter and low perceived loss for video streaming. The scheme called "Smart Multipath Capacity Aggregation (SMCA)" uses elementary information about the video stream (eg: I, P, B frame type, packetization and sequencing information) to intelligently map packets to flows in different groups. Flows are classified into delay- and loss-based classes using end-to-end estimation of delay, loss and rate information. The size of these groups is adaptively determined by the current set of application packets and flow characteristics. The gains are realized by efficient matching of application content diversity to the network performance diversity at any instant of time. SMCA is designed to be scalable with increasing number of available network paths and with increasing content and network diversity. Our experiments demonstrate marked improvement in video playback quality measures both on an absolute and relative basis (compared with other path-diversity based schemes). Interestingly, our relative gains are even better with an increase in performance diversity of network paths.**

***Index Terms*— Multimedia streaming, multipath, path diversity.**

I. Introduction

The notion of best-effort service in the Internet has traditionally implied an unpredictable packet-by-packet service delivered over on a single path. Today, demanding applications like video streaming are dependent on the performance vagaries and bottlenecks on a single path. The performance bottlenecks on such single paths are also moving away from the access link (eg: last mile) due to the deployment of broadband access and the general availability of alternate access options (eg: why not use cable modem, 3G wireless and DSL together?).

Once the access bottleneck is removed, we realize that the Internet intrinsically has a multiplicity of end-to-end paths because hosts, networks and autonomous systems (of enterprises and ISPs) are increasingly multi-homed. The spatio-temporal statistical multiplexing gains from these paths can be harnessed to deliver a superior form of end-to-end best-effort service to applications (beyond the temporal multiplexing gain offered by packet switching on single paths). If both end-systems have broadband or high-speed access, the goal

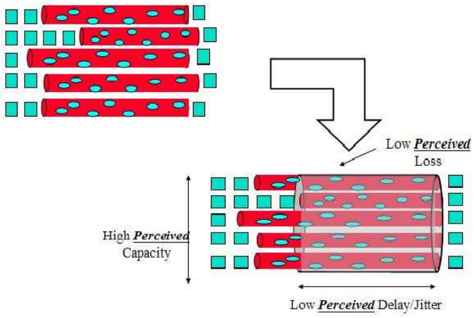

Fig. 1: SMCA logically rearranges the available paths for improved streaming experience for the end-user.

is to create the abstract of an end-to-end broadband pipe built out of purely best-effort underlying components (Fig. I). End-to-end path multiplicity gain can be realized *even if* routing protocols only offer single (shortest) paths, as long as key edge-nodes support the mapping of end-to-end flows to different exit choices. Connectionless routing frameworks have been proposed for incrementally upgrading the Internet to support multi-paths (e.g. Bananas [23]). Multi-paths may also be provisioned through overlay networks or peer-to-peer networks ([1], [2], [3], [4]). The responsibility of end-systems is then to instantiate multiple flows, locally map them to multiple interfaces, perform congestion control on each flow and manage the mapping of application packets to flows [23], [16].

This paper shows how end-systems can effectively harness multiple paths *even if* these paths are very diverse in performance characteristics and *even if* applications are demanding in terms of deadline, reliability and sequencing expectations. In fact, we show that increased diversity of path performance and application expectations can be leveraged to provide even better perceived performance! In particular, we propose a scheme called "Smart Multipath Capacity Aggregation (SMCA)" that matches application content-diversity (in terms of per-packet performance expectations) to the per-flow performance diversity to realize a virtual end-to-end pipe abstraction that offers high capacity, low perceived delay jitter and low perceived loss for video streaming (Fig. I).

SMCA uses elementary information about the video stream (eg: I, P, B frame type, packetization and sequencing information) to intelligently map packets to flows in different groups. Flows are *dynamically* classified into delay- and loss-based groups using end-to-end estimation of delay, loss and rate information. The sizes and boundaries of these groups are *adaptively* determined by the current set of application packets and flow characteristics. SMCA maps more important

Manuscript received 08-01-2006. This work was supported by the ARO Contracts (W911NF-04-1-0300, DAAD19-00-1-0559), NSF-ITR 0313095 and a grant from Intel Corporation

O. Tickoo is with the Intel Corporation, 2111 NE 25th Ave., Hillsboro OR 97124 USA (phone: 503-712-5019; fax:503-264-1544; email:omesh.tickoo@intel.com)

S. Kalyanaraman is with ECSE Dept., Rensselaer Polytechnic Institute, 110 8th St., Troy NY 12180 USA (email:shivkuma@ecse.rpi.edu).

video packets to superior flow groups to maximize probability of delivery and timeliness. Poor quality flow groups (low available bandwidth, high latency, high loss rates) are used to send lower priority data or packets with later decoding deadlines. This out-of-sequence and importance-based mapping scheme helps us utilize the flow resources that are otherwise unusable due to playback deadlines at the receiver. Our experiments demonstrate marked improvement in video playback quality measures both on an absolute and relative basis (compared with other path-diversity based schemes). At the same time we show that naive mapping of packets to available paths may not deliver these diversity gains to the application. Our approach fundamentally differs from many of the earlier proposed packet-mapping schemes ([5], [6], [10], [11], [12], [16]) because we are able to use paths that would potentially be deemed unusable by other schemes. Therefore our SMCA scheme concepts are applicable to both wired and wireless networks, overlay and peer-to-peer networks, and to a broader range of applications beyond video streaming.

We present the related work in the next section. High level overview and analytical details of SMCA are presented in Section III. Section IV presents the simulation results and the work is summarized in Section V.

II. Related Work

Providing sufficient bandwidth for delay sensitive applications has been an active area of research during the recent years. The total capacity of a network path puts an upper limit on the bandwidth a user can get from the network. To solve these problems, solutions have been proposed that employ multiple paths to reduce the packet loss and increase the effective bandwidth obtained. Savage etal [13] report that in majority of cases one can find a more optimal path to the destination as compared to the default paths provided by any routing protocol. Thus, there is a good possibility of finding multiple paths that can satisfy a real-time application's transmission requirements. Vutukury etal [14] give methods for near optimal Multipath routing. Nguyen etal [15] propose a Multipath scheme that utilizes source routing and installed network relays for video transmission. Source routing requires the media source to specify the exact path of transmission in terms of the intermediate hops.

In [5] Apostolopoulos etal use two different paths to send even and odd frames encoded using Multiple Description Coding (MDC) but reference [5] does not use any kind of network feedback. The paper suggests that it can be beneficial to send different *amounts* of traffic on different paths. In [12] Lian etal use transmission of multiple redundant descriptions of the voice streams over independent network paths. Receivers use multi-stream adaptive playout scheduling to improve the trade-off among delay, loss-rate and speech quality. The paper reports better quality in Multipath transmitted voice as compared to the FEC protected voice streams (in terms of mean end-to-end latency and loss-rate). The path diversity is achieved by sending media through the default and a source-based route. The two different flows are constructed using the even and odd samples of the voice stream. More recent studies involve schedulers that tightly couple the loss-rate experienced with the transmission capacity of individual paths in a Multipath scenario [20]. Such schemes do not take per-path latency characteristics into consideration.

In [6] Apostolopoulos etal provide models to compare multiple description coding (MDC) plus path diversity against single description coding plus single path. The paper also presents a model for the loss process of a two-path diversity system. In [18] Zhou etal present a transmission scheme to improve MPEG-4 streaming using Multipath. The MPEG stream is divided into a base and enhancement layer. The base layer is duplicated over the paths available to provide robustness for transmission of important frames. The enhancement layer content is then separated into multiple descriptions and sent over the multiple paths leading to incremental increase in the video quality with reception of more and more enhancement packets. The drawbacks of these schemes are the additional complexity introduced at the source (separating the enhancement layer into multiple descriptions is almost as complex as coding the video afresh) and the wasted bandwidth due to duplication of the base layer. Another drawback of using the MDC based schemes is the extra traffic added by introducing the multiple description splitting of the traffic. These schemes also suffer from degraded performance in case the bandwidth of each path is smaller than the overhead introduced. In addition, most of the schemes discussed above are verified for two paths. With two paths all the above algorithms provide performance improvements over the single path but one would expect the complexity of these schemes to increase with the number of paths used.

The closest attempt at analyzing an efficient partitioning scheme has been made in [16] by Xu etal. Reference [16] gives a good overview of the issues involved in partitioning a differentially encoded bit stream over multiple routes. A pruned tree approach with complexity of the order of $O(N^Q)$ is presented where N is the number of paths available and Q is the total number of frames in a group of pictures (GOP). The performance comparisons are made with a greedy multiplexing technique. While the pruned tree approach in [16] provides an optimal solution to the multiplexing problem, it does so at the cost of high complexity. SMCA, on the other hand, is designed to be practical while being efficient and low in complexity. In [17], Xu etal present a new channel coding scheme (product codes) to unequally protect the video for efficient transmission over multiple paths.

III. Scheme Description

This section builds the analytical basis for SMCA. For the sake of simplicity we assume that each frame is transmitted as a single packet. This assumption does not have any effect on our final results since same results are obtained when all packets containing a frame's data are treated in similar fashion. We reinforce the fact that for a video transmission system to take maximum advantage of path diversity, the transport layers must have knowledge of the source coding process and both layers need to work in conjunction with each other. Our choice of layered video is based on previous research that

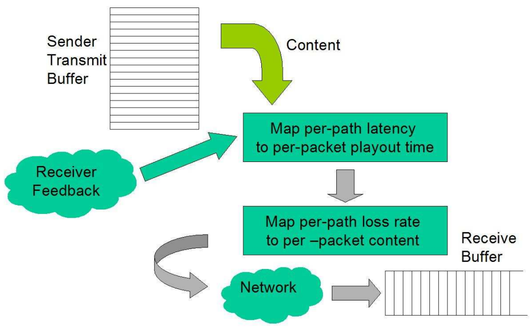

Fig. 2: SMCA scheme overview. The degrading effects of network latency and loss on perceived video quality are reduced in two separate steps.

Parameter	Definition
N	Total available paths
l_0	path with the lowest end-to-end latency
l_i	Path i in the latency based ranked list of paths
B_i	Bandwidth (frames/sec) of path i
D_i	End-to-end delay/latency of path i
f_1	Frame/packet at the head of the transmit buffer
f_i	Frame/packet at the position i in the transmit buffer
τ_i	Transmission time for f_i
k_i	Effective carrying capacity of l_i
$t(i)$	Expected playout time for f_i
S	Source Buffer Size in packets
Δ	Avg. time (sec) between successive packets playout

TABLE I: Parameter definition for out-of-sequence analysis

shows layered coding as the method of choice when content aware packet schedulers are used [19].

The SMCA architecture is represented in Fig. 2. The sender's transmit buffer is filled with video packets by the application in a serial fashion. The packets corresponding to the frame to be transmitted/decoded earliest occupy the head of the buffer. The SMCA scheme is used to *choose* frames for transmission from the transmit buffer as described below. There are two main stages in mapping the video frames to the appropriate paths. The first stage assigns the frames to a set of paths under delay constraints. The second stage optionally protects video frames with FEC and maps each video and redundancy frame to a path using the content information. The delay-based mapping stage helps minimize the effective delay by using out-of-sequence transmissions on high latency paths. Similarly, content-based mapping of frames and error correction reduces the impact due to the network losses.

SMCA estimates the path characteristics at the sender. This information comprises of loss-rate, bandwidth and latency values at different time instants. This estimation is done using the congestion window behavior and acknowledgment information from the transport scheme.

A. Delay based packet-to-path mapping

The delay-based mapping stage of SMCA is concerned with reducing the effective *perceptible* delay of the deadline driven real-time traffic. Assume that the receiver starts the playout after it receives the first P packets completely. Then, in steady state, maintaining the receiver-buffer occupancy of P packets on an average requires the average packet interarrival time to be equal to the playout time of a packet. Let P packets take T seconds to playout. After the first P packets have arrived, the next P packets should reach the destination within T seconds. The third set of P packets must reach within $2T$ seconds and so on. This provides us with an expected arrival time for each packet.

Table I defines the parameters used in the analysis in this section:

At the sender, the packets are assigned to the paths using the following procedure:

Find the largest $n1$ such that

$$\tau_n + D_{n1} \leq t_1 \tag{1}$$

Then, considering just the delay requirements, the paths l_1 to l_{n1} are suitable for any of the packets in the transmit buffer. The $q_1 = \sum_{i=1}^{n1} k_i$ packets are mapped to the paths l_1 to l_{n1} forming the first delay-based subgroup. The packets from $q_1 + 1$ onward are again grouped separately and mapped on paths l_{n1+1} onward creating the second delay-based subgroup. The number of paths, in this case will be given by the largest integer $n2$ such that

$$\tau_{n1+n2} + D_{n1+n2} \leq t(f_{q_1+1}) \tag{2}$$

In case $\tau_{n+1} + D_{n+1} \geq t(f_{q_1+1})$, we skip the packets in the sender's transmit buffer until we reach the packet f_k, $k \geq q_1+1$, that satisfies the delay condition $\tau_{n+1}+D_{n+1} \leq t(f_k)$. The skipped packets between the packets f_{q_1+1} and f_k can wait for transmission and will be transmitted in the subsequent refresh periods.

The second group of paths l_{n1+1} to l_{n1+n2} is assigned q_2 packets for transmission in a similar fashion as q_1 were assigned in the first step. We continue this grouping of paths and assignment of packets for transmission until either all the paths are categorized or we run out of packets. Each group-of-paths and associated set-of-packets is referred to as a Delay Based Subgroup (DBS). The delay reduction unit utilizes the out-of-sequence transmission that reduces the overall transmission delay by mapping the packets positioned higher up in the transmit buffer to the paths with higher relative latencies.

We now derive the effective carrying capacity of each delay-based subgroup. As an example. we consider the first delay based subgroup. For a path l_i within this subgroup, the total time taken by k_i packets to reach their destination is given by;

$$T_i = D_i + \frac{k_i}{B_i} \tag{3}$$

We call the average latency of path 0, denoted by D_0 the *Base Latency* for the this subgroup. Realizing that within this subgroup any packet may be transmitted on any of the $n1$ paths, on an average the condition for timely arrival of packets at the receiver can be summarized as:

$$D_i + \frac{k_i}{Bi} \leq D_0 + \Delta \tag{4}$$

The condition (4) can be derived as following. If packet F is transmitted on path 0, it suffers an end-to-end latency of D_0 (Base Latency). Packet $F+1$ must reach the receiver within $\triangle$ of packet F to avoid buffer underrun. Since packet $F+1$ may be transmitted on any of the $n1$ paths in the tagged subgroup, all the path latencies in the subgroup must be less that $D_0 + \triangle$.

Rearranging (4) and for each path l_i, denoting the difference in the individual latency and the Base Latency by δ_i, we have:

$$k_i \leq (\Delta - \delta_i) \times B_i, \quad 0 \leq \delta_i \leq \Delta \tag{5}$$

Equation (5) is the defining equation for the out-of-sequence mapping scheme and it shows that for out-of-sequence transmission the relative latency rather than the absolute latency of the paths within the delay-based subgroups determines the effective carrying capacity. Thus, for the first subgroup the effective carrying capacity s given by:

$$K_1 = \sum_{i=0}^{n1-1} k_i \;\leq\; \Delta \times B - \sum_{i=0}^{n1-1} (\delta_i \times B_i) \tag{6}$$

where $B = \sum_{i=0}^{n1-1} B_i$ is the aggregated bandwidth (in packets/sec) of the subgroup.

Equation (6) provides the upper bound on the carrying capacity of the first delay-based subgroups. Similar result holds true for an arbitrary delay based subgroup j with K_1 replaced by K_j and summation executed over the paths in subgroup j.

Example: Consider a set of paths between a source and a destination. We can choose any arbitrary D_0, since the analysis above shows that the carrying capacity of the subgroup is dependent on the *delay spacing* between the paths rather than the absolute value of the base delay. The available paths are ranked in increasing order of latency with a uniform separation of 10 ms between any two successive paths. Assume that each path has an average bandwidth of 1 $Mbps$. Assume that we want to transfer video encoded at 30 fps with an average frame length of 500 Bytes. Armed with this knowledge we can now calculate $\Delta = 1/30 = 33.33$ms, $B = 3$Mbps $= \frac{3\times10^6}{500\times8}$ fps, $\delta_i = 0, 10\times10^6, 20\times10^6$ sec for $i = 0,1,2$ and $B_i = \frac{1\times10^6}{500\times8}$ fps for $i = 0,1,2$.

From the value of Δ we can select upto 3 paths, starting from the path with least delay, for inclusion in the first delay-based subgroup. The maximum number of frames that can be assigned to this subgroup, F_{max} is given by Equation (6).

$$F_{max} = \frac{3\times10^6}{500\times8} \times \frac{1}{30} - \frac{1\times10^6}{500\times8} \times (10\times10^6 + 20\times10^6) \approx 17\text{frames} \tag{7}$$

By symmetry of the example, the next subgroup will be assigned a maximum of 17 frames. In this case we just need to send 13 more frames so two subgroups i.e. 6 paths will suffice. In contrast we consider an opportunistic packet mapping scheme that transmits the packets from the head of the application's transmission buffer onto the paths that are ready to accept a packet at any instant of time. We note that if such an opportunistic packet mapping scheme was used in place of the out-of-sequence transmission, we would have only 3 usable paths (the first subgroup) for our effective use. The frames transmitted on the other paths can lead to late deliveries and consequent receiver-buffer underflow. Thus, out-of-sequence transmission achieves a much higher guaranteed frame rate, 30 fps versus 17 fps for an opportunistic packet mapping scheme. This better performance of out-of-sequence transmission is attributed to the ability of the scheme to use higher delay paths in an efficient manner, like the second delay-based subgroup in the example above.

1) Effect of source buffer size: Source buffer size limits the throughput of the out-of-sequence transmission scheme. The number of packets k_{oos} that can be transmitted on all the available paths in time t_{oos} is bound above by the source buffer size i.e. given the source buffer of length S frames

$$T_{oos} = \frac{k_{oos}}{t_{oos}} \leq \frac{S}{t_{oos}} \tag{8}$$

2) Effect of receiver buffer: The receiver-buffer size limits the number of packets associated with any subgroup. A playout buffer of size R packets drains at the rate of 1 packet per Δ seconds. We need at least 1 packet in the buffer and a maximum of R packets. This in turn implies that in steady state the incoming packet rate into the receiver-buffer must range between [1 to R] frames per Δ seconds. This has the following performance consequences:

1) Each delay-based group on the sender side is upper-bound by R packets i.e. in (6), $k_i \leq R$ for all paths in set O.
2) *Jitter Bounds:* The receiver-buffer needs at least one packet in the queue to avoid underflow. To prevent buffer overflow, receiver-buffer occupancy must stay below R packets. Since the playout occurs at the rate of 1 packet per Δ seconds, on appropriately large time scales $\frac{1}{\Delta}\ to\ \frac{R}{\Delta}$ defines the acceptable range of incoming data rate at the receiver. Any jitter within this range will be absorbed by the receiver-buffer.

B. Loss-based mapping

The second step involves the exact mapping of frames to paths within a DBS. For hybrid video streaming significant quality improvements can be obtained by exploiting the dependencies between different frame types and ***without worrying about the interdependencies between frames of same type.***. We take as an example the case involving transmission of video encoded using MPEG or H.26x. These coders encode each group of pictures (GOP) into three different frame types: I, P and B. Each frame type is of varying importance within a GOP. Typically, each GOP starts with an I or an intra-coded frame. An I frame can be decoded independently without any reference to the frames preceding or succeeding it. On the other hand, P frames need the information from the latest reference frame (I or P) for correct decoding while the B frames need information from both the preceding as well as the succeeding reference frames for correct decoding. Extending the analysis from the performance model presented in [22] to multiflow transmission, we devise a greedy yet simple strategy for prioritized frame transmission. We use the following algorithm for the transmission of I, B and P frames for a given set of GOPs within a DBS. The available paths are ranked in the increasing order of loss rates. The loss rates are measured using the information from acknowledgments received for a given number of packets sent over an interval [8]. The multiplexing scheme is a simple frame type based prioritized mapping wherein the I frames are mapped to the available paths with low loss rates followed by the P frames. The B frames are then mapped to the remaining paths in increasing order of the loss rates. The combination of out-of-sequence delay-based mapping and smart multiplexing uses the network diversity in latencies and loss rates to the streaming application's advantage.

1) Complexity analysis: Assume a GOP containing Q frames is to be multiplexed over N paths. If transmission across GOP boundaries was allowed, transmission of a GOP under SMCA involves the following steps:

1) Compare the first frame with the available paths under delay constraints (N comparison operations) to select a *set-of-paths*.
2) Find the suitable number of frames, under bandwidth constraints, that can be multiplexed over the *set-of-paths* found above (Q comparison operations).
3) For the *set-of paths* to frames mapping determined above, find the frame-to-path mapping under loss constraints (N comparisons per frame selected).

We note that including the restriction of transmitting only integral GOPs over a *delay-based subgroup* does not change the complexity of the scheme. Therefore, the complexity of SMCA is $O(N^2+QN)$. We compare the complexity of SMCA smart multiplexing with a completely optimized pruned tree(PT) based approach [16]. The worst case complexity of a pruned tree based approach is $O(N^Q)$. We note that the complexity of SMCA is much lower than that of the pruned tree approach and as we show later, the performance of SMCA is comparable to the performance of a scheme based on the pruned tree algorithm.

C. Error Control to Combat Transmission Errors

If the underlying congestion control scheme used for SMCA does not discriminate between congestion losses and transmission errors, SMCA can suffer performance degradation in lossy environments like wireless networks. To protect the real-time data against transmission losses we propose to integrate Forward Error Correction with SMCA. In this section we introduce two possible FEC strategies. FEC schemes are designed based on the RS(n,k) [24] codes.

Given an average path loss rate of P_a over the n paths in a subgroup, the total video packets that can be transmitted is reduced to:

$$k \leq (1 - P_a) \sum_{n} k_i \tag{9}$$

The remaining capacity is used for transmission of FEC information. Equation 9 provides us with the upper limit on the amount of the data (video) that may be transmitted on the group of n identified paths. The rest of the capacity is used by the FEC packets. We present two different techniques to use this bandwidth.

The first technique uses uniform error control to protect all the application data. In case of uniform control all the video packets are treated alike and the content is protected with RS(m,k) FEC as given by Equation 9. The FEC packets are treated with the same priority as the video data they protect. A point worth noting is that since the data and FEC are *clustered* together, they will be exposed to similar and probabilistically correlated path conditions. Hereafter, we label this scheme as SMCA_UFEC.

The second technique exploits the content priorities of hybrid video to unequally protect the most important video frames. Unequal or prioritized error control involves unequally protecting the video packets according to their relative importance within the GOP. In our case of MPEG and H.26x video, the relative importance of I, P and B frames dictates the amount of FEC allocated to the associated FEC. Of course, the total FEC allocated cannot exceed $m-k$ packets. To make the video transmission robust to path loss correlations, the FEC for a frame is sent on a path that is *far* from the path (within a *set-of-paths*) on which the frame is sent. The scheme uses unequal protection of I and P frames and no protection for B frames. The FEC is *decoupled* from the transmission of the data by reserving the last few paths in the set-of-paths for transmission of FEC. Hereafter, we label this scheme as SMCA_UEP.

IV. RESULTS

We now present the performance evaluation of the complete SMCA scheme with a video-streaming application. The SMCA framework consists of the out-of-sequence delay-based scheme and the content-based smart multiplexing scheme. We also present performance improvements obtained by adding FEC as an additional protection to erasures. Performance comparisons are presented for a uniform FEC protection scheme and a content-based non-uniform FEC scheme.

Figures 3 and 4 show the simulation set-up used for measuring the performance of SMCA. Source S multiplexes video traffic destined for destination D over multiple links constituted by hosts/routers $B1$ to $B16$. Fig. 3 shows the set-up for the case where the source multiplexes traffic over 5 uncorrelated paths while the set-up of Fig. 4 corresponds to the correlated links case. In the case of shared links like those in Fig. 4, the loss correlation among different paths is determined by the link characteristics and the amount of background traffic through each of the shared links. The background traffic generators consist of sources transmitting variable bit rate (VBR) and constant bit rate (CBR) traffic. The bandwidth of each link varies between 300 Kbps to 1 Mbps. A 20 packets buffer is provided at each transmit interface of the source and the nodes $B1$ to $B16$.

We compare the performance of our scheme with OPMS and the pruned tree(PT) algorithm of [16]. We used the *Flower*

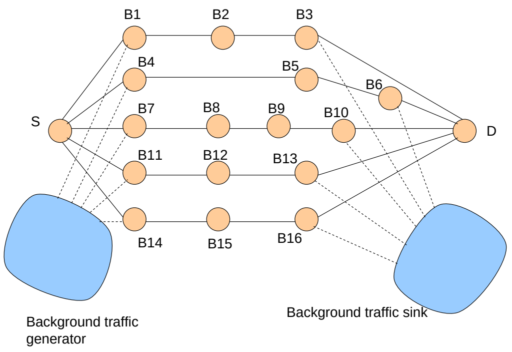

Fig. 3: Uncorrelated link topology

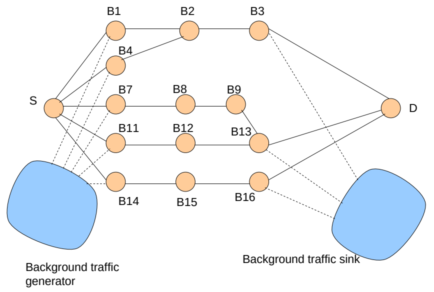

Fig. 4: Correlated link topology

Garden video test sequence with SIF resolution (352 x 240 pixels) at 30 fps. The Flower Garden sequence was encoded using the H.26L encoder. The GOP had 16 frames in the following order: IBBPBBPBBPBBPBBP. The bit-rate was 1.7 Mbps and the average packet size was 700 bytes. In the following subsection we present our performance evaluation results. It is important to note that the average delay values we quote in the following subsection correspond to the values that were administratively configured in the associated topologies. The delays due to intermediate node buffer occupancy are in addition to the values we quote.

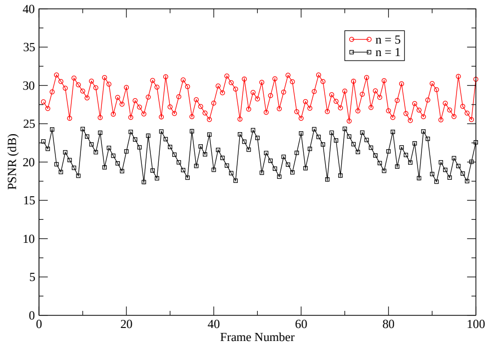

Fig. 5: Gain versus number of paths (more than 3dB gain provides substantial video quality improvement.)

Paths	1	2	3	4	5
PSNR(dB)	20.98	22.48	25.42	26.02	28.34

TABLE II: Variation in the average PSNR with number of paths

1) Gain with increasing number of paths: Table II and Figure 5 show the SMCA performance with changing number of paths. The substantial gain of more than 7 dB when the network resources are diversified among 5 paths shows that SMCA uses path diversity to the user's advantage. The total number of paths was varied from one path ($n = 1$) to five paths ($n = 5$) while keeping the aggregate bandwidth fixed at 1.3 Mbps, average loss probability fixed at 0.1 and average path-delay fixed at 30 ms.

We now compare the performance of SMCA with OPMS and PT over varying values of loss and delay.

2) Uncorrelated topology - path loss variation: Fig. 6 and Table III presents the PSNR gains achieved by using SMCA compared to the OPMS for the topology of Fig. 3. The average path-delay in the simulations was set at 30 ms. Fig. 6 presents the performance results obtained with different average path loss-rates. We observe that the difference in gain between SMCA and the OPMS is positive in all cases. This difference in gain increases with the increase in the average loss-rate and increase in loss-rate diversity among the links. This gain is due to the smart multiplexing scheme that exploits the loss diversity in the paths to transmit high importance frames over low loss paths. However, Table III shows that the average PSNR declines as the average loss-rate of the paths increases. Table III also shows the PSNR comparison with the completely optimized PT scheme. We notice that the average PSNR in case of the SMCA degrades less steeply as compared to the greedy and PT schemes. The average PSNR for the 0.4 average loss-rate case in Table III is worth mentioning here. In this case we observe that the content received using the OPMS has an average PSNR of 11.64 dB which makes it almost impossible to display a fair quality video stream to the user. Under similar conditions the performance of SMCA is within acceptable limits for fair decoded video quality (avg. PSNR = 22.78 dB). We also observe that SMCA achieves comparable performance to the PT approach at a much lower complexity.

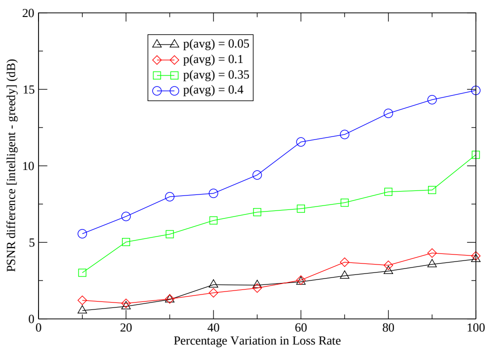

Fig. 6: Gain versus loss diversity (uncorrelated paths) from a greedy (OPMS) scheme by using intelligent path aggregation (SMCA)

Avg. Loss Prob.	SMCA PSNR(dB)	PT PSNR(dB)	OPMS PSNR(dB)
0.05	29.32	31.82	26.06
0.1	29.03	29.02	24.43
0.35	26.32	26.86	18.21
0.4	**22.78**	**20.31**	**11.64**

TABLE III: Gain versus loss diversity - uncorrelated paths. The gains with SMCA over OPMS increase with the increase in average loss-rate and the diversity in loss-rate.

3) Correlated topology - Path loss variation: Fig. 7 and Table IV presents PSNR gains achieved by using our SMCA scheme compared to the OPMS and PT schemes for the topology of Fig. 4. The average path-delay in the simulations was set at 30ms. Fig. 7 presents the performance results obtained with different average path loss rates. The results observed indicate a similar comparative performance as the uncorrelated topology case. This result is important since the users have almost negligible control over the paths provided by the network and a scheme that provides better performance under all conditions of loss correlations is highly desirable for real-time applications.

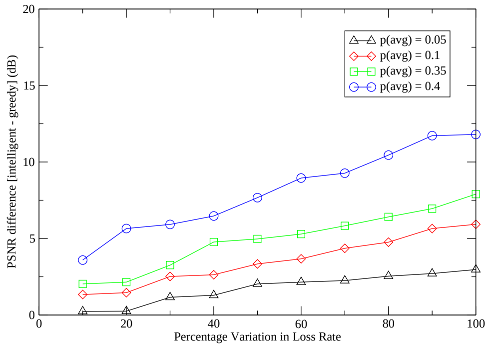

Fig. 7: Gain versus loss diversity - correlated paths. The gain of SMCA over OPMS increases with the increase in average loss-rate and the diversity in loss-rate.

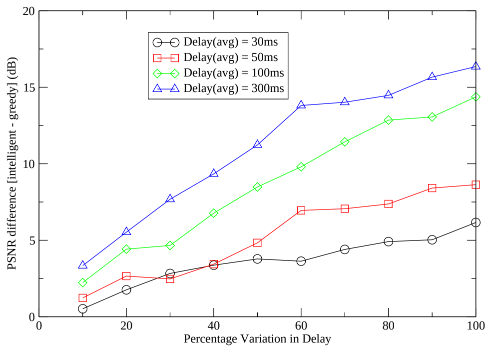

Fig. 8: Gain versus Delay Diversity. The gain of SMCA over OPMS increases with the increase in average delay and the diversity in latency.

4) Uncorrelated topology - Path delay variation: The gain in performance with varying delay characteristics are shown in Fig. 8 and Table V. The different curves in Fig. 8 correspond to different average path-delay settings. The average loss-rate over the paths was fixed at 0.1. The variation in the individual path delays was set at a maximum of 100% from the average. An interesting observation from Fig. 8 is that SMCA performs much better than OPMS in the case of high variations of delay characteristics among paths. This validates the correct operation of the out-of-sequence transmission scheme. From Table V we observe that SMCA achieves comparable performance to the PT approach (at a much lower complexity).

Avg. Loss Prob.	SMCA PSNR(dB)	PT PSNR(dB)	OPMS PSNR(dB)
0.05	28.12	30.42	26.37
0.1	26.33	27.47	23.86
0.35	23.87	20.64	19.75
0.4	**20.26**	**18.36**	**12.54**

TABLE IV: Gain with loss diversity - correlated Paths

Avg. Delay	SMCA PSNR(dB)	PT PSNR(dB)	OPMS PSNR(dB)
30ms	30.12	31.83	27.96
50ms	28.32	29.46	24.33
100ms	25.12	24.21	19.19
300ms	**21.78**	**18.73**	**11.03**

TABLE V: Gain with delay variation

5) Performance with FEC: In this subsection we provide the simulation results of SMCA coupled with two different FEC schemes: the Uniform-FEC (SMCA_UFEC) and FEC using Unequal Error Protection (SMCA_UEP) as described in Section III-C

The topology used for simulation in this subsection is the correlated path topology depicted in Fig. 4. Each path has a mean latency of 30 ms and the average path loss values are controlled for each simulation. Table VI presents the average PSNR values for each of the three schemes. It is observed that while SMCA_UEP outperforms SMCA and SMCA_UFEC the improvements increase with the increase in the average loss-rate of the paths. This is expected since SMCA_FEC provides robustness in the presence of high losses by protecting the important data and making sure that the important data and associated FEC is decoupled in case of correlated losses. Again, we observe that the performance gain increases with the loss-rate and loss-rate variation.

V. Summary

In this paper we presented a route-aggregation scheme, SMCA, that exploits the diversity in network paths to satisfy real-time application's transmission requirements. SMCA uses a novel out-of-sequence transmission strategy to use high latency paths for transferring packets with non-immediate playout times from the transmit buffer. While utilizing this otherwise unusable bandwidth, the out-of-sequence transmission scheme also helps reduce the overall transmission delay. A smart content-based mapping scheme is used by SMCA

Avg. Loss Rate	SMCA Avg. PSNR (dB)	SMCA_UFEC Avg. PSNR (dB)	SMCA_UEP Avg. PSNR (dB)
0.1	26.02	28.34	30.01
0.2	24.75	27.89	29.23
0.3	23.96	27.02	28.96
0.4	**22.32**	**24.53**	**26.59**

TABLE VI: Average PSNR with and without FEC (Avg. delay = 30ms)

Avg. Loss Rate	Variation in Loss Rate 10%	20%	30%	40%	50%
0.1	1.24	2.53	2.95	3.27	3.99
0.2	1.52	2.75	3.16	3.82	4.03
0.3	1.73	2.82	3.37	3.92	4.12
0.4	**2.93**	**3.09**	**3.59**	**4.16**	**4.27**

TABLE VII:
Gain in PSNR (in dB) with SMCA_UEP from SMCA with loss-rate variation (Avg. Delay = 30ms)

to counter effects of network loss. The mapping criteria, though sub-optimal, provides gains at much lower complexity than a fully optimized one. This scheme can be used to map both video and associated FEC in a decoupled manner to avoid performance degradation due to correlated network losses. The mapping is done in an *adaptive* manner to keep up with the network path dynamics. The simulation results show that SMCA performs better than an opportunistic packet mapping scheme and its performance is comparable to a fully optimized multiplexing scheme. Relative performance improvements gained using SMCA increase with the path diversity and higher values of average path loss and latency.

REFERENCES

[1] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker, “A Scalable Content-Addressable Network,”*Proc. of ACM SIGCOMM*, Aug. 2001.

[2] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, “Chord: A scalable peer-to-peer lookup service for internet applications,” *Proc. of the ACM SIGCOMM*,San Diego, California, August 2001.

[3] A. Rowstron and P. Druschel, “Pastry: Scalable, distributed object

[4] B. Y. Zhao, J. D. Kubiatowicz, and A. D. Joseph, “Tapestry: An infrastructure for fault-resilient wide-area location and routing”, *Technical Report UCB//CSD-01-1141*, U. C. Berkeley, April 2001.

[5] J. G. Apostolopoulos, “Reliable Video Communication over Lossy Packet Networks using Multiple State Encoding and Path Diversity,”*Visual Communications and Image Processing*, January 2001.

[6] J. Apostolopoulos, W. Tan, S. Wee and G. W. Wornell, “Modelling Path Diversity for Multiple Description Video Communication,”*Proc. of ICASSP 2002*, May 13-17, 2002.

[7] I. V. Bajić and J. W. Woods, “Domain-based multiple description coding of images and video,” *IEEE Trans. Image Procesing*, vol. 12, no. 10, pp. 1211-1225, October 2003.

[8] D. Bansal and H. Balakrishnan, “Binomial Congestion Control Scheme,” *Proc. of IEEE INFOCOM*, Tel-Aviv, Israel, March 2000.

[9] I. V. Bajic, O. Tickoo, A. Balan, S. Kalyanaraman, and J. W. Woods, “Integrated end-to-end buffer management and congestion control for scalable video communications,” *Proc. IEEE ICIP*, Barcelona, Spain, September 2003.

[10] A. C. Begen, Y. Altunbasak and O. Ergun, “Fast heuristics for multi-path selection for multiple description encoded video streaming,” *Proc. IEEE Int. Conf. Multimedia and Expo (ICME)*, Baltimore, MD, July 2003.

[11] A. C. Begen, Y. Altunbasak and O. Ergun, “Multi-path selection for multiple description encoded video streaming,” *Proc. of IEEE Int. Conf. Communications (ICC)*, Anchorage, AK, May 2003.

[12] Y. J. Liang, E. G. Steinbach and B. Girod, “Multi-Stream Voice Over IP Using Packet Path Diversity,” *Proc. of IEEE 4th Workshop on Multimedia Signal Processing*, Oct. 2001.

[13] S. Savage, A. Collins, E. Hoffman J. Snell, and T. Anderson, “The End-to-End Effects of Internet Path Selection,” *Proc. of SIGCOM*, 1999.

[14] S. Vutukury and J.J. Garcia-Luna-Aceves, “MDVA: A Distance-Vector Multipath Routing Protocol,” *Proc. of INFOCOM*, 2001.

[15] T. Nguyen and A. Zakhor, “Path Diversity with Forward Error Correction System for Packet Switched Networks,” *Proc. of INFOCOM*, April 2003.

[16] W. Xu and S. S. Hemami, “Efficient Partitioning of Unequal Error Protected MPEG Video Streams for Multiple Channel Transmission,” *Proc. IEEE Int. Conf. on Image Processing*, Rochester, NY, Sept. 2002.

[17] W. Xu and S. S. Hemami, “Delay-Optimized Robust Transmission of Image over Multiple Channels,” *Proc. of IEEE Int. Conf. on Multimedia and Expo*, Baltimore, MD, Jul. 2003.

[18] J. Zhou, H.R. Shao , C. Shen and M. T. Sun, “Multi-path Transport of FGS Video,” *Proc. of Packet Video*, April 2003, Nantes, France.

[19] J. Chakareski, S. Han and B. Girod, “Layered Coding vs. Multiple Descriptions for Video Streaming over Multiple Paths,” *Proc. of the eleventh ACM international conference on Multimedia*, Berkeley, California, USA, 2003.

[20] M. Bansal and A. Zakhor, “Path Diversity for Overlay Multicast Streaming”, *Proc. of Packet Video Workshop*, Irvine, CA, Dec. 2004

[21] J.C. Chen and S.-H. Chan, “Multipath Routing for Video Unicast over Bandwidth-Limited Networks,” *Proc. of IEEE Globecom*, Vol. 3, pp. 1963-1697, San Antonio, Texas, November, 2001.

[22] K. Mayer-Patel, L. Le, and G. Carle, “Transport-level Protocol Coordination in Cluster-to-Cluster Applications,” *Proc. of ACM Multimedia*, 2002.

[23] H. T. Kaur, S. Kalyanaraman, A. Weiss, S. Kanwar and A. Gandhi, “BANANAS: An Evolutionary Framework for Explicit and Multipath Routing in the Internet,” *Proc. of ACM SIGCOMM Workshop on Future Directions on Network Architectures (FDNA)*, Volume 33, Issue 4, pp. 277-288, Karlsruhe, Germany, August 2003.

[24] S. B. Wicker and V. K. Bhargava, “Reed-Solomon Codes and Their Applications,” Wiley-IEEE Press, September 1999.