Reliable Interactive Video Streaming in Peer-to-Peer Networks

Xiangyang Zhang[1], Glen Takahara[2], and Hossam Hassanein[3]
[1,3]School of Computing, [2]Dept. of Math. & Stat., Queen's University, Kingston, Canada
{[1]xiang, [3]hossam}@cs.queensu.ca, [2]takahara@mast.queensu.ca

***Abstract*—Forward error correction (FEC) coding is the preferred error correction technique for interactive video streaming applications on the Internet. Because its performance is impaired by the burstiness of packet loss of Internet links, peer-to-peer (P2P) networks are often proposed to provide multiple paths between a sender and a receiver. However, peers may leave abruptly and the number of disjoint paths may be limited; it is unclear whether or when the use of P2P networks for path diversity can be justified. In this paper, we study the packet loss ratio after FEC correction when using P2P networks to provide multiple paths. We examine two situations: a sender can find enough disjoint paths, or uses a limited number of disjoint paths. We model Internet links using Markov chains, provide numerical analysis of the performance of systematic FEC codes, and verify the results by simulation. We find that although using P2P networks for path diversity often results in a lower post-FEC loss ratio, conditions apply. There exist guidelines but no simple formula to determine when to use P2P networks for path diversity and coding parameters. An application should carefully evaluate the performance gain before taking actions.**

I. Introduction

The past few years have witnessed significant growth of interactive video streaming applications, such as video telephony and video conferencing, on the Internet. These applications can only tolerate a few hundred milliseconds of delay at best and are sensitive to packet loss. Forward error correction (FEC) coding is a widely used error correction technique on noisy communication channels for interactive streaming applications (as well as other real-time communications). However, FEC performance is impaired when errors are bursty, while studies [1], [2] show that on the Internet, packet loss often occurs in bursts of consecutive loss. The general idea of using multiple paths to combat bursty packet loss has been explored in a number of research efforts [3], [4]. On the Internet, the most feasible way to maintain multiple paths is to use peer-to-peer (P2P) networks. Indeed, many interactive streaming applications, most notably Skype [5], use P2P networks. However, in practical P2P networks, peers are dynamic. Using peers as intermediate hops to achieve path diversity may cause additional packet loss when they leave abruptly. Besides, there may not exist enough disjoint paths. It is unknown whether using P2P networks for path diversity can really reduce the post-FEC loss ratio or under what conditions the ratio can be reduced, and how much the reduction can be.

In this paper, we study the packet loss ratio after FEC correction when using P2P networks to provide multiple paths, and compare with the loss ratio when a single direct path is used. We study two situations of path diversity: a sender uses a large number of disjoint paths if they exist, or uses a limited number of disjoint paths. We model Internet links, with or without peers as intermediate hops, using discrete time Markov chains. Unlike other studies that either use non-systematic FEC codes[1] or compute with heuristic methods, we use systematic FEC codes and compute exact numerical results, and verify these results by simulation. We find that although using P2P networks for path diversity often results in a lower post-FEC loss ratio, conditions apply. There exist guidelines but no simple formula to determine whether the use of P2P networks for path diversity can be justified. For example, using more disjoint paths, stabler peers, and more coding redundancy can reduce the post-FEC loss ratio. However, we also find that if Internet links have a low loss ratio or certain coding parameters are used, using multiple paths may not reduce the post-FEC loss ratio compared with using only the direct path, and certain coding parameters are "better" than others from a performance/cost perspective. An application should carefully evaluate the performance gain before using P2P networks for path diversity or deciding coding parameters.

The remainder of this paper is organized as follows. Section II introduces related work. Section III describes how to use P2P networks to achieve path diversity. Section IV analyzes the post-FEC loss ratio when the direct path is used; the results serve as the benchmark. Sections V analyzes the post-FEC loss ratio when multiple paths are used. Sections VI presents the numerical and simulation results. Section VII concludes this paper.

II. Related Work

Elliott [6] pioneered the study of FEC performance in the context of telephone networks. He used the Gilbert channel model [7] and introduced two general approaches—a recursive approach and a generating function approach—to compute the probability that m errors occur in a block of n bits after FEC correction on a single channel. However, his approaches do not distinguish the loss of original bits and FEC bits and hence only apply to non-systematic FEC codes. Frossard [8] uses the recursive approach and a 2-state Markov chain channel model to compute the post-FEC packet loss ratio and burst

[1]Codes that include the original packets in the output are systematic and codes that do not are non-systematic. In this paper, we focus on systematic codes because they have a significantly lower post-FEC loss ratio and burst length than non-systematic codes.

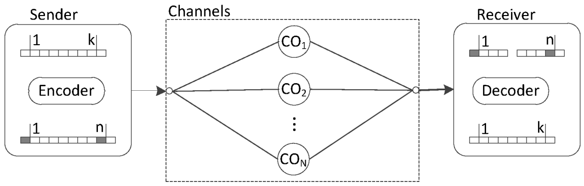

Fig. 1. Network model. The sender encodes $k = 7$ original packets into an FEC block of size $n = 8$ by appending an FEC packet and sends encoded packets via N paths in a round-robin fashion. The receiver can reproduce all the original packets if no more than $n - k = 1$ packet is lost in transmission.

TABLE I
PARAMETERS

Parameter	Comment
n, k	The number of total packets and original packets in an FEC block with $RS(n, k)$ coding
N	The number of disjoint paths
d	The probability that a peer leaves in any time unit
s	The time the sender takes to detect a CO peer's departure
p, q	The parameters of the 2-state and $(s+2)$-state Markov chains
$e, (e_c)$	The packet loss ratio of unicast (or CO) channels
$l, (l_c)$	The average burst length of unicast (or CO) channels

length for systematic FEC codes on a single channel. However, we find his method to be incorrect[2]. Yu et al. [9] study the situation where a number of senders send packets to their respective receivers via a bottleneck node, which is modeled as a $G/M/1/K$ queue. They use non-systematic FEC codes and compute the post-FEC loss ratio using a recursive method.

The general idea of using multiple paths to combat bursty loss has been explored in a number of research efforts [3], [4]. Begen et al. [4] uses a 2-state Markov chain channel model and multiple description coding (MDC), and formulate a path selection problem to maximize video quality. Li et al. [3]'s work is closest to ours. They use systematic FEC codes and N disjoint Internet paths, and extend the 2-state Markov chain to s-state to model packet loss on each path. In each state i, the chain has probability $p_{i,i+1}$ to go to state $i + 1$ (except at the state $s - 1$ it has probability $p_{s-1,s-1}$ to go to itself) and probability $p_{i,0}$ to go to state 0. They use a heuristic method to estimate the post-FEC loss ratio, which is based on two assumptions: (1) both the size n of FEC blocks and the number N of paths are large, and (2) the loss of original packets and FEC packets of an FEC block are independent and have a normal distribution. Both assumptions do not apply to interactive streaming applications.

III. NETWORK MODEL

In a P2P network, a peer acquires its initial membership table, which consists of peers in the P2P network, by contacting a well-known *tracker* or peers already in the system, and maintains the table size to be in a certain range (e.g., between 50 and 100). To reduce communication overhead, a peer does not exchange keep-alive messages with peers in its membership table or uses a large interval (e.g., one keep-alive message every 30 minutes). Therefore, a peer may not know that some peers in its membership table have already left. Two peers can communicate directly or via a number of care-of (CO) peers to achieve path diversity (see Fig. 1). We call paths without CO peers *unicast channel*s and paths with a CO peer *CO channel*s. When a sender and a receiver initiate a communication session, they try to find CO peers that are located between them on the Internet (i.e., the sum of propagation delays from a CO peer to the sender and the receiver is close to the propagation delay between the sender and the receiver). There are many schemes for the sender and receiver to find these peers in a distributed P2P network. For example, they can use virtual network position schemes such as [10] and a center server that maintains all the peers' network positions or distributed position-based routing schemes [11] to do so. The sender specifies the destination peer in a packet and sends the packet to one of the CO peers, which immediately forwards the packet to the destination peer. (We only discuss one direction of the communication; the other direction is the same.) When multiple paths exist, the sender spreads packets across all the paths in a round-robin fashion. Peer departures (and arrivals) are typically modeled as a Poisson process, which is equivalent to a Bernoulli process in discrete time. An existing CO peer has probability $0 \leq d \leq 1$ to leave in any time unit. (Interactive streaming applications typically have a near-constant bit rate, and we denote the interval between two consecutive packets to be one time unit.) The sender and receiver stay in the system for the whole communication session, i.e., they have probability 0 to leave. The sender exchanges keep-alive messages with CO peers periodically; it takes s time units for the sender to detect the departure of a CO peer. Parameters are listed in Table I.

Internet users typically connect to a point-of-presence (PoP) of an ISP, or several PoPs for higher reliability, to connect to the Internet. For example, residential users usually connect to a single PoP, possibly via a home network; campus users may connect to several PoPs via a local area network (LAN). Beyond PoPs, the Internet core is a rich mesh network. Since there may not be enough disjoint channels between a sender and a receiver, we study two situations. (1) A sender uses a limited number of disjoint paths. (2) Because the path segment from a user to its PoP usually has large bandwidth and low loss ratio compared with the Internet core, when this loss ratio is small and can be ignored, we can considere that there exist sufficient disjoint paths between a sender and a receiver.

Systematic Reed Solomon codes $RS(n, k)$ are most widely used FEC codes on the Internet. As shown in Fig. 1, a sender appends $n - k$ FEC packets to k original packets to form an FEC block. The receiver can reproduce the original k packets if it receives any subset of k packets out of the n packets in an FEC block. If the receiver receives $k' < k$ of the k original packets and less than $k - k'$ FEC packets, then $k - k'$

[2]We use a Reed Solomon code $RS(3, 2)$ as an example. The post-FEC loss ratio should be $\frac{p(2-q+pq)}{2(p+q)}$ but is $\frac{p(2-q^2+pq)}{2(p+q)}$ according to Frossard [8], where p and q are parameters of the 2-state Markov chain.

original packets are lost. Note that the loss of original and FEC packets has different impacts on the post-FEC loss ratio when more than $n-k$ packets of an FEC block are lost. For example, for $RS(10,9)$, if 2 original packets of an FEC block are lost, the post-FEC loss ratio is $\frac{2}{9}$; if 1 original packet and 1 FEC packet are lost, the post-FEC loss ratio is $\frac{1}{9}$. The main disadvantages of FEC are the coding overhead and coding delay. For interactive streaming applications, FEC should not introduce a delay of more than several hundred milliseconds, i.e., the FEC block size n should be small.

IV. FEC Performance with the Direct Path

In this section we compute the post-FEC loss ratio when a sender uses only the direct path (which is a unicast channel). The results are used as the benchmark.

A. Loss Model for Unicast Channels

We choose to use a 2-state Markov chain to model a unicast channel. This model is a simplification of the Gilbert model [7] and is widely used in the literature [4], [8]. Like the Gilbert model, the channel is assumed to have constant bit rate (or packet rate). As long as an Internet path between a sender and a receiver has higher bandwidth than the video's streaming rate, it can be considered that the Internet path provides a constant bit rate channel. In each time unit, the channel is either in the good state (state 0) or in the bad state (state 1). Packets transmitted over the channel are error-free when the channel is in the good state and are erroneous when the channel is in the bad state[3]. Given the channel's current state, the channel's state for the next time unit can be described by the one-step probability transition matrix $P = \begin{pmatrix} 1-p & p \\ q & 1-q \end{pmatrix}$.

The 2-state Markov chain has stationary distribution $\pi = (\frac{q}{p+q}, \frac{p}{p+q})$. The channel's loss ratio e is the proportion of time that the chain stays in state 0 and hence $e = \frac{p}{p+q}$; Denote 0^i and 1^i as a run of i consecutive 0's and 1's respectively. The probability that a burst of i consecutive 1's occurs is $Pr\{1^{i-1}0|01\} = q(1-q)^i$, i.e., the burst length has a geometric distribution, and hence the average burst length is $l = \frac{1}{q}$. Sometimes it is more convenient to describe the channel by parameters e and l rather than by parameters p and q. Given e and l, $p = \frac{e}{(1-e)l}$ and $q = \frac{1}{l}$.

B. Post-FEC Loss Ratio When Using the Direct Path

Let X denote the number of unrecoverable original packets lost in an FEC block. The post-FEC loss ratio e_{fec} is the expectation of X divided by the number of original packets in an FEC block. To compute $E[X]$, we need to count the number of unrecoverable original packets lost in each possible FEC block and compute its probability, i.e., $e_{fec} = \frac{1}{k}\sum_{\bar{x}} Pr\{\bar{x}\}F(\bar{x})$, where $F(\bar{x})$ is the number of unrecoverable original packets lost in FEC block $\bar{x}$. By properties of Markov chains, the probability of that an FEC block $\bar{x} = (x_1 \dots x_n)$, where $x_i \in \{0,1\}$, occurs can be reduced to

$$Pr\{x_1 \dots x_n\} = Pr\{x_n|x_{n-1}\} \dots Pr\{x_2|x_1\}Pr\{x_1\}$$

[3]In the Gilbert model, when the channel is in bad state, a packet transmitted over the channel has a certain probability to be error-free.

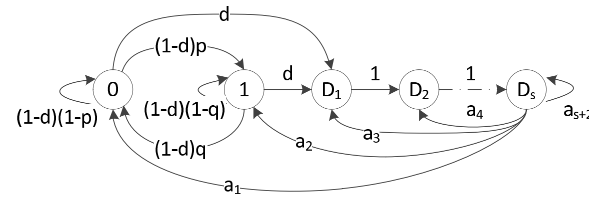

Fig. 2. Loss model for CO channels

where $Pr\{x_i|x_{i-1}\}$ and $Pr\{x_1\}$ are available from the probability transition matrix and stationary distribution respectively. There are $O(2^n)$ possible FEC blocks of length n. Because the FEC block size n is a small number for interactive streaming applications, this algorithm is computationally feasible.

V. FEC Performance With Multiple CO Channels

In this section, we first model CO channels and subchannels and then compute the post-FEC loss ratio when a sender uses a sufficient or limited number of disjoint paths.

A. Loss Model for CO Channels

A CO channel consists of two unicast channels and a CO peer between them. The state of the CO peer and the states of the two unicast channels are independent from one another. Therefore, we can move the CO peer to right before the receiver and consider a CO channel as a concatenation of two unicast channels (which is also a unicast channel) with a CO peer at the end point. In the following, we consider a CO channel as a unicast channel with parameters p and q plus a CO peer that has probability d to leave in any time unit.

When a sender forwards packets to a receiver via N disjoint CO channels, if the CO peer of channel i leaves, the sender will find a new CO peer to replace it. (We still call the channel channel i.) Assume all the CO channels are homogeneous (i.e., the unicast channels have the same parameters p and q and the CO peers have the same probability d to leave the system in any time unit). Then we can model a CO channel, with CO peer replacement, as an $(s+2)$-state Markov chain.

As shown in Fig. 2, state 0 refers to when the CO peer is present and the underlying unicast channel is in the good state, and state 1 refers to when the CO peer is present but the unicast channel is in the bad state. The CO peer has probability d to leave the system, regardless of whether the unicast channel is in state 0 or state 1. A sender takes s time units to detect the departure of a CO peer x. When peer x leaves after forwarding the i-th packet at the t-th time unit, the sender will continue to send s packets to peer x before it switches to a new CO peer y at the $(t+s)$-th time unit. These states are represented by $D_1, D_2, \dots, D_s$. Peer y may be in any of the $s+2$ states, which is reflected by the transition probabilities $a_1, a_2, \dots, a_{s+2}$. Since the sender takes s time units to detect the departure of a CO peer, peer y may have already left between the t-th and the $(t+s)$-th time units. (The sender will know peer y has left and will not use peer y if it has left before the t-th time unit). Peer y has equal probability to have left in any of the s time units, and the probability that peer y has left is $1-(1-d)^s$. By duality, we also have $\frac{a_1}{a_2}=\frac{\pi_1}{\pi_2}=\frac{q}{p}$. Now we can write the chain's one-step probability transition matrix P_c as

$$\begin{pmatrix} (1-p)(1-d) & p(1-d) & d & 0 & \dots & 0 \\ q(1-d) & (1-q)(1-d) & d & 0 & \dots & 0 \\ 0 & 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & 0 & \dots & 1 \\ a_1 & a_2 & a_3 & a_4 & \dots & a_{s+2} \end{pmatrix}$$

where

$$\begin{cases} a_1 = \dfrac{q}{p+q}(1-d)^s \\ a_2 = \dfrac{p}{p+q}(1-d)^s \\ a_i = \dfrac{1-(1-d)^s}{s},\ i \in [3, s+2] \end{cases} \quad (1)$$

The $(s+2)$-state Markov chain has stationary distribution π such that

$$\begin{cases} \pi P_c = \pi \\ \sum_{i=1}^{s+2} \pi_i = 1 \end{cases}$$

Solving the above equations, we have

$$\begin{cases} \pi_1 = \dfrac{q}{p+q} \times \dfrac{(1-d)^s}{d} \times \dfrac{1}{\alpha} \\ \pi_2 = \dfrac{p}{p+q} \times \dfrac{(1-d)^s}{d} \times \dfrac{1}{\alpha} \\ \pi_{i+2} = \dfrac{i+(s-i)(1-d)^s}{s} \times \dfrac{1}{\alpha},\ i \in [1, s] \end{cases} \quad (2)$$

where $\alpha = \frac{s+1}{2} + \frac{s-1}{2}(1-d)^s + \frac{1}{d}(1-d)^s$.

The CO channel's loss ratio e_c is the fraction of time that the chain stays in states other than state 0, i.e., $e_c = 1-\pi_1$. The probability that a gap of length i occurs between two bursts of non-zeros is $(1-d-p+pd)^i(d+p-pd)$, i.e., the gap length has a geometric distribution, and hence the average gap is $\frac{1}{d+p-dp}$. Since there are the same number of gaps and bursts, the average burst length is $l_c = \frac{1-e}{(d+p-dp)e}$. However, the burst length no longer has a geometric distribution.

B. Subchannels of a CO Channel

When the sender sends packets via N disjoint CO channels in a round-robin fashion, only $\frac{1}{N}$ of each channel is used. We consider the two parts of a CO channel—the unicast channel and the CO peer—separately. Given a unicast channel with one-step probability transition matrix P, its subchannel can also be modeled as a 2-state Markov chain, and its one-step probability transition matrix P' is the N-step probability matrix of the whole channel. By the Chapman-Kolmogorov equations, $P' = \begin{pmatrix} 1-p & p \\ q & 1-q \end{pmatrix}^N$.

From a subchannel's perspective, if the CO peer is present for the current packet at time t, it will leave with probability $1-(1-d)^N$ before time $t+N$ for the next packet. When the CO peer leaves, the sender will continue to send $\lfloor \frac{s}{N} \rfloor$ packets to the CO peer. Therefore, the subchannel of a CO channel can be modeled as a $(\lfloor \frac{s}{N} \rfloor + 2)$-state Markov chain with parameters

$$\begin{cases} p' = \dfrac{p}{p+q}(1-(1-p-q)^N) \\ q' = \dfrac{q}{p+q}(1-(1-p-q)^N) \\ d' = 1-(1-d)^N \\ s' = \lfloor \dfrac{s}{N} \rfloor \end{cases} \quad (3)$$

C. Post-FEC Loss Ratio When Using Sufficient CO Channels

When there exist $N \geq n$ disjoint channels with a loss ratio e_c, since each packet in an FEC block is sent via a channel different from other packets in the FEC block, whether the packet can reach the receiver without error is independent from the states of other packets. Therefore, the probability $P(m,n)$ that m packets are erroneous out of n consecutive packets has a binomial distribution

$$P(m,n) = \frac{n!}{m!(n-m)!} e_c^m (1-e_c)^{n-m}$$

Conditioning on the number of original packets lost in an FEC block, the post-FEC loss ratio is

$$e_{fec} = \frac{1}{k} \sum_{i=1}^{k} iP(i,k) \sum_{j=n-k+1-i}^{n-k} P(j, n-k) \quad (4)$$

D. Post-FEC Loss Ratio When Using Limited CO Channels

As in the case of using the direct path to send packets, when $N < n$, the state of a packet may depends on the states of a previous packet in the same FEC block. To compute the post-FEC loss ratio, we need to count the number of unrecoverable original packets lost in each FEC block and compute its probability.

Assume N divides n and s. Denote $\bar{x}_i = (x_{1i}x_{2i}\dots x_{bi})$, where $b = \frac{n}{N}$, as the subsequence of channel i. Then an FEC block $\bar{x}$ consists of N interleaved subsequences,

$$\bar{x} = (x_{11}x_{12}\dots x_{1N}x_{21}x_{22}\dots x_{2N}\dots x_{b1}x_{b2}\dots x_{bN})$$

Given the CO channel's parameters p, q, d and s, the subchannel's parameters p', q', d' and s' can be computed using Equation 3. The subchannel's stationary distribution π' can be computed using Equation 2 (with parameters p', q', d' and s'). Because these N CO channels are disjoint, the probability of sequence $\bar{x}$ is the product of the probabilities of subsequences, i.e., $Pr\{\bar{x}\} = \prod_{i=1}^{N} Pr\{\bar{x}_i\}$. The probability of each subsequence can be computed using the same method as described in Section IV-B (with the subchannel's probability transition matrix P_c' and stationary distribution π'). The set of all possible subsequences can be generated packet by packet using a recursive method. The state of the first packet ranges from 0 to $s'+1$. Given the state of the current packet, the possible state of the next packet can be obtained from Fig. 2.

VI. Numerical and Simulation Results

In this section we present the numerical and simulation results. For each numerical result, we run corresponding simulation test for 10 times. The simulation results are very close to the numerical results. Except when the post-FEC loss ratio is less than 0.1%, the 95% confidence interval of the simulation result is with 5% of the corresponding numerical result. Therefore, we choose not to draw simulation results on figures since there will be no discernible difference.

A. Methodology

Both the numerical analysis program and simulation program are written in the Java programming language. The numerical analysis program uses the BigDecimal package for arbitrary-precision computing. In the simulation program, a unicast channel is simulated using the 2-state Markov chain model, a CO channel is simulated as the concatenation of a unicast channel and a CO peer, and a subchannel is simulated as every N-th slot of a channel. The initial state of a unicast channel is set to 0 with probability $\frac{q}{p+q}$ and to 1 with probability $\frac{p}{p+q}$. We set the initial state of a CO peer as follows to simulate a peer that has a random age between 0 and 10 minutes and was present during the sender's last detection. First we select a random age for the CO peer. It has probability d to leave in any time unit. If it has not left until s time unit before the current time, we use its state at the current time. Otherwise, we repeat the above process. The simulation time is 600 minutes for all the tests.

We use the following default parameters unless specified otherwise. The streaming rate is 480 Kbps for TV-quality streaming, the packet size is 1500 bytes, and thus 1 time unit is 25 milliseconds. We use an FEC block size of $n = 16$, which introduces 400 milliseconds of coding delay. According to [1], [2], we use two link loss ratios: 1% and 10%, and two burst lengths: 2 and 4. According to an empirical study [12] on P2P video live streaming, we set a peer's average lifespan to 10 minutes, i.e., $d = \frac{1}{40 \times 60 \times 10}$. A sender takes 1 second to detect a CO peer's departure, i.e., $s = 40$.

B. Results When Using Sufficient CO Channels

Figs. 3 and 4 show the post-FEC loss ratio when a sender uses $N \geq n$ disjoint CO channels. First, compared with using the direct path, when multiple CO channels are used, the post-FEC loss ratio decreases sharply with the increasing number of FEC packets. This observation indicates that path-diversity can greatly improve FEC performance. Also note that the post-FEC loss ratio is significantly lower when l is 2 than when l is 4 if the direct path is used, but is not impacted by the burst length if $N \geq n$ CO channels is used because each packet is transmitted via a separate channel.

Second, the use of CO peers achieves path diversity but also introduces extra packet loss. Compared with a unicast channel with parameters p and q, an CO channel with parameters p, q, d and s has extra loss ratio $\Delta = e_c - e$, which is proportional to the detection time s and inversely proportional to peers' lifespans, and has larger impact when Internet links are less

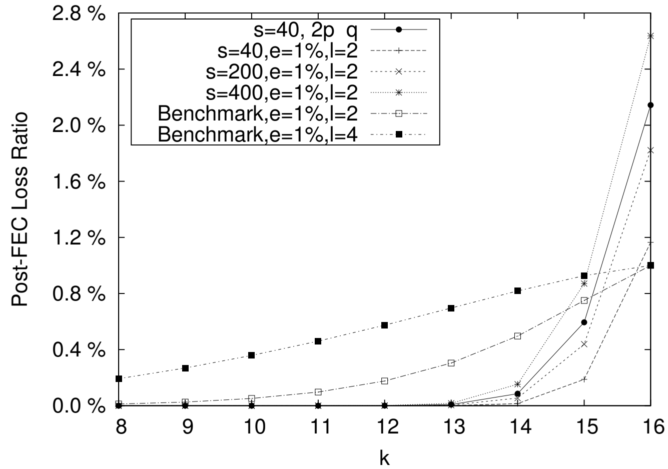

Fig. 3. Post-FEC loss ratio when $N \geq n$ and e=1%.

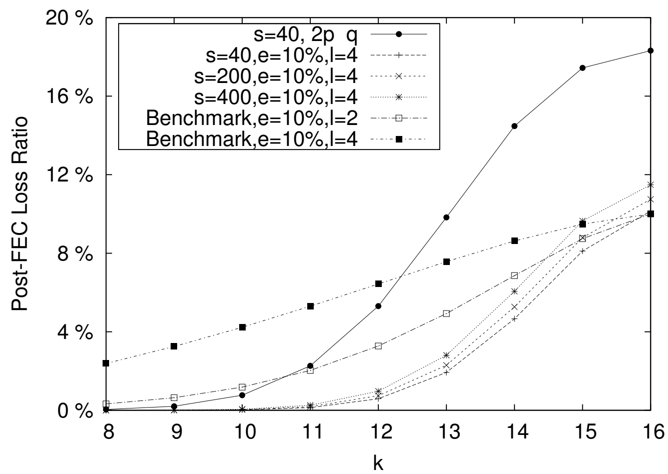

Fig. 4. Post-FEC loss ratio when $N \geq n$ and e=10%.

noisy. (In Figs. 3 and 4, when k is 16, the post-FEC loss ratios when using multiple CO channels and when using the direct path equal the CO channels' loss ratio e_c and the direct path's loss ratio e, respectively.) We vary s to show the impact on the post-FEC loss ratio. The values s=40, 200, and 400 correspond to 1, 5, and 10 seconds. With the increase of the detection time, the post-FEC loss ratio increases for every value of k.

Third, if a sender cannot find CO channels with similar parameters p and q as the direct path, it has to use more FEC packets to achieve a lower post-FEC loss ratio than that of the direct path. For example, when both segments of a CO channel have the same parameters p and q as the direct path (i.e., the CO channel has parameters $2p$ and q), because the loss ratio of CO channels more than doubles that of the direct path, 4 redundant FEC packets are required for 12 original packets to achieve a post-FEC loss ratio similar to that of the direct path when Internet links have an loss ratio of 10% (see Fig.4), resulting in a coding overhead of $\frac{1}{3}$.

C. Results When Using Limited CO Channels

Figs. 5 and 6 show the post-FEC loss ratio when a sender uses $N < n$ disjoint CO channels. For every value of k, e and

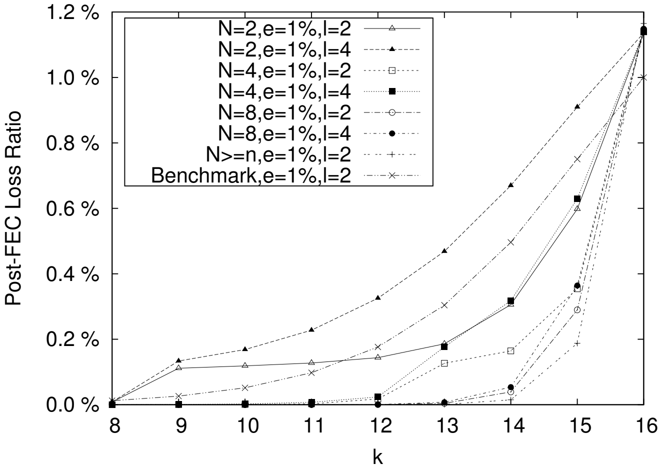

Fig. 5. Post-FEC loss ratio when $N < n$ and e=1%.

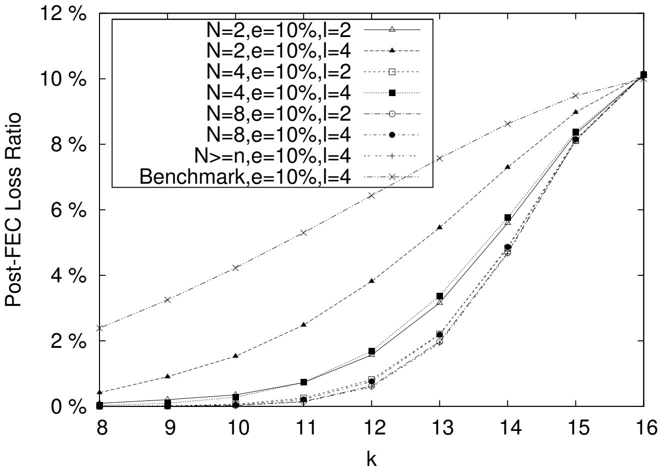

Fig. 6. Post-FEC loss ratio when $N < n$ and e=10%.

l, the post-FEC loss ratio drops when more channels are used, but the decrease becomes marginal when N exceeds a certain value. As in the case of using the direct path, a longer burst length results in a higher post-FEC loss ratio. The increase is more significant when a sender uses fewer channels.

Two points are worth pointing out. First, each increment of the number of FEC packets results in certain overhead but does not result in equal decrease of the post-FEC loss ratio. Certain combinations of k and n are "better" than other combinations from a performance/cost perspective. For example, in Fig. 5, when N is 4 and l is 2, using 3 FEC packets (i.e., k is 13) only slightly reduces the post-FEC loss ratio compared with using 2 FEC packets, but using 4 FEC packets results in a significantly lower post-FEC loss ratio than using 3 FEC packets. Another example is when N is 2, l is 2, and k is 8, compared with k equal to 12, 11, 10, and 9. Second, using multiple CO paths does not necessarily result in a lower post-FEC loss ratio than using the direct path, even when the extra packet loss caused by peer departures is insignificant. For example, in Fig. 5, when N is 2, l is 2, and k is 9, 10, and 11, using the direct path has a lower post-FEC loss ratio. Both phenomena are the result of the interplay between the distribution of loss ratio and burst length, number of disjoint channels, and the coding parameters n and k. Although there are guidelines, there is no simple formula to determine what are the good combinations of k and n or whether the use of multiple CO channels can be justified.

VII. Conclusion

In this paper we study the performance of systematic FEC codes in interactive streaming applications when using P2P networks to provide multiple paths between a sender and a receiver. We find that although using P2P networks for path diversity often results in a lower post-FEC loss ratio, conditions apply. The interplay of a number of factors—the Internet links' loss ratio and burst length, the number of disjoint CO channels, the lifespans of peers, the time to detect peers' departures, and the coding parameters—determines the post-FEC loss ratio. There exist guidelines but no simple formula to determine whether the use of P2P networks for path diversity can be justified. When a sender uses a large number of disjoint CO channels with similar parameters as the direct path, path diversity typically leads to a lower post-FEC loss ratio unless the direct path has a low loss ratio such that the extra packet loss caused by peer departure is significant. When a sender uses a small number of disjoint CO channels, a longer burst length results in a higher post-FEC loss ratio, and the more disjoint CO channels a sender uses, the lower the post-FEC loss ratio will be. Using multiple CO channels does not necessarily result in a lower post-FEC ratio compared with using only the direct path even when the extra packet loss caused by peer departures is insignificant, and some combinations of k and n of the $RS(n, k)$ codes are better than others from a performance/cost perspective.

References

[1] V. Paxson, "End-to-end Internet packet dynamics," *ACM SIGCOMM Computer Communication Review*, vol. 27, no. 4, pp. 139–152, 1997.

[2] M. Yajnik, J. Kurose, and D. Towsley, "Packet loss correlation in the MBone multicast network," in *GLOBECOM*, 1996, pp. 94–99.

[3] Y. Li, Y. Zhang, L. Qiu, and S. Lam, "Smarttunnel: Achieving reliability in the internet," in *Proc. IEEE INFOCOM*, 2006, pp. 830–838.

[4] A. Begen, Y. Altunbasak, and O. Ergun, "Multi-path selection for multiple description encoded video streaming," in *IEEE Int. Conf. on Comm. (ICC)*, vol. 3, 2003, pp. 1583–1589.

[5] "http://www.skype.com," Accessed Mar. 2011.

[6] E. Elliot, "A model of the switched telephone network for data communications," *Bell Syst. Tech. J*, vol. 44, no. 1, p. 89, 1965.

[7] E. Gilbert, "Capacity of a burst-noise channel," *Bell Syst. Tech. J*, vol. 39, no. 9, pp. 1253–1265, 1960.

[8] P. Frossard, "FEC performance in multimedia streaming," *IEEE Communications Letters*, vol. 5, no. 3, pp. 122–124, 2001.

[9] X. Yu, J. Modestino, R. Kurceren, and Y. Chan, "A model-based approach to evaluation of the efficacy of FEC coding in combating network packet losses," *IEEE/ACM Trans. on Networking*, vol. 16, no. 3, pp. 628–641, 2008.

[10] F. Dabek, R. Cox, F. Kaashoek, and R. Morris, "Vivaldi: A decentralized network coordinate system," in *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 4, 2004, pp. 15–26.

[11] S. Giordano, I. Stojmenovic, and L. Blazevic, "Position based routing algorithms for ad hoc networks: a taxonomy," *Ad hoc wireless networking*, pp. 103–136, 2004.

[12] B. Li, S. Xie, G. Y. Keung, J. Liu, I. Stoica, H. Zhang, and X. Zhang, "An empirical study of the Coolstreaming system," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 9, pp. 1627–1639, 2007.