AN APPROACH FOR IMPROVING PERFORMANCE

OF AGGREGATE VOICE-OVER-IP TRAFFIC

A Thesis

by

CAMELIA AL-NAJJAR

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

August 2005

Major Subject: Computer Engineering

AN APPROACH FOR IMPROVING PERFORMANCE

OF AGGREGATE VOICE-OVER-IP TRAFFIC

A Thesis

by

CAMELIA AL-NAJJAR

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Approved by:

Chair of Committee,	A. L. Narasimha Reddy
Committee Members,	Andreas Klappenecker
	Weiping Shi
Head of Department,	Chanan Singh

August 2005

Major Subject: Computer Engineering

ABSTRACT

An Approach for Improving Performance
of Aggregate Voice-over-IP Traffic. (August 2005)
Camelia Al-Najjar, B.S., University of Jordan
Chair of Advisory Committee: Dr. A. L. Narasimha Reddy

The emerging popularity and interest in Voice-over-IP (VoIP) has been accompanied by customer concerns about voice quality over these networks. The lack of an appropriate real-time capable infrastructure in packet networks along with the threats of denial-of service (DoS) attacks can deteriorate the service that these voice calls receive. And these conditions contribute to the decline in call quality in VoIP applications; therefore, error-correcting/concealing techniques remain the only alternative to provide a reasonable protection for VoIP calls against packet losses. Traditionally, each voice call employs its own end-to-end forward-error-correction (FEC) mechanisms. In this paper, we show that when VoIP calls are aggregated over a provider's link, with a suitable linear-time encoding for the aggregated voice traffic, considerable quality improvement can be achieved with little redundancy. We show that it is possible to achieve rates closer to channel capacity as more calls are combined with very small output loss rates even in the presence of significant packet loss rates in the network. The advantages of the proposed scheme far exceed similar or other coding techniques applied to individual voice calls.

To my parents,
brothers, and sister

ACKNOWLEDGMENTS

I would like to express my sincere appreciation to Dr. A. L. Narasimha Reddy whose continuous guidance, numerous suggestions, and valuable insights have guided me throughout the course of this research.

I would like to thank my committee members Dr. Andreas Klappenecker and Dr. Weiping Shi for their time and cooperation. I also thank Dr. Krishna Narayanan and Abhiram Prabhakar for their assistance with my work with LDPC codes.

And finally, I would also like to express my gratitude to my family for their unwavering support and constant encouragement during the course of my graduate program.

TABLE OF CONTENTS

Page

ABSTRACT iii

DEDICATION iv

ACKNOWLEDGMENTS v

TABLE OF CONTENTS vi

LIST OF FIGURES viii

LIST OF TABLES ix

1 INTRODUCTION 1

2 PACKET AUDIO: CALL QUALITY AND SPEECH CODING 6

2.1 Evaluation of Call Quality 7
2.2 Speech Processing and Compression 8

3 ERROR CONCEALMENT AND CORRECTION FOR PACKET-AUDIO 10

3.1 Signal-Processing Based FEC (SFEC) 11
3.2 Forward-Error Correction 12

4 AGGREGATE CODING OF VoIP TRAFFIC 14

5 RESULTS AND EVALUATION 15

5.1 Recovery Performance of LDPC 15
5.2 Aggregation vs. Single Call 17
5.3 Evaluation of the Performance by Network Simulations 22
5.3.1 The Opposing Traffic 25
5.3.2 Comparison of the Recovery Level of SFEC and LDPC 30
5.3.3 Effect of Changing Parameters 33
5.4 Network Simulations with Multiple Links 38

6 DISCUSSION 43

Page

7 CONCLUSION 45

REFERENCES 46

VITA 50

LIST OF FIGURES

FIGURE Page

1 Configuration of a VoIP network. 3

2 Process of aggregation, encoding, transmission, and decoding for VoIP calls... 4

3 Sparse graph representation of LDPC codes. 13

4 Overhead vs. maximum loss per single block. 16

5 Input loss rate vs. redundancy (32KB block, 128B packet). 17

6 Output loss vs. redundancy for different input loss rates. 19

7 Output loss vs. redundancy for different input loss rates (closer view). 19

8 Output loss vs. redundancy for different link speeds; assuming loss rate = 0.1, delay = 50 ms. 20

9 Output loss vs. redundancy for different link speeds; assuming loss rate = 0.1, delay = 10 ms. 22

10 Resilience to opposing traffic with large packet sizes. 26

11 Recovery of SFEC and LDPC with a single CBR flow at different loads. 27

12 The effect of a drift in the sending rate of the flows on loss rates. 34

13 The effect of drift on recovered SFEC loss rate. 35

14 Effect of varying inter-flow spacing for two CBR setups. 36

15 Effect of jitter in packet inter-arrival times on loss rate (scenario 3). 37

LIST OF TABLES

TABLE Page

1 Properties of different coding standards. 9

2 Characteristics of the evaluated SFEC streams. 24

3 Resulting quality for a load of 20% (11% loss) for a single CBR flow. 28

4 Various scenarios where coding techniques are evaluated. 29

5 Distribution of quality for the 82 calls on the link. 30

6 Scenario 1: losses focused on subset of flows (loss 9.9%). 31

7 Scenario 4: losses focused on subset of flows (loss 18.9%). 32

8 Scenario 6: losses spread out among the flows (loss 10.6%). 33

9 Performance of SFEC vs. LDPC with -5% variation (loss = 9.8%). 35

10 Effect of a jitter of 5% on loss recovery in scenario 3. 37

11 Resulting quality with 16 msec-spaced attack traffic. 39

12 Quality distributions for the three cases of table 11. 39

13 Comparison of the different burst spacing (equal to voice packet spacing or not). 39

14 Recovery of LDPC and SFEC for the 16 msec case. 40

15 Recovery of LDPC and SFEC for the 7 msec case. 40

16 The effect of adding jitter to the 16 msec burst-spacing configuration. 41

17 The effect of adding jitter to the 7 msec burst-spacing configuration. 41

18 Loss run measurement comparison for both SFEC and LDPC. 41

1. INTRODUCTION

Voice-over-IP (VoIP) services have witnessed increased popularity and acceptance in recent years, but this trend comes along with customer concerns about voice quality over these networks. At the same time, most VoIP users may not realize the susceptibility of these applications to the same security threats which data traffic suffers from on packet networks. Achieving acceptable real-time interactivity of voice calls directly depends on loss, jitter, and delay along with some other factors. The unpredictable behavior of packet networks coupled with threats from denial-of-service attacks, which are prevalent in these networks, makes it difficult to achieve toll quality comparable to that of PSTN systems.

When VoIP services emerged, it was strictly PC-to-PC calls which required an Internet connection and compatible software to succeed. As it developed, gateways to the PSTN were installed and PCs were no longer a requirement [21]. But Internet backbones are still used for the long haul for long-distance VoIP calls [20]. The major VoIP providers have moved to the use of private links instead where they can guarantee some behavior of the network [6]. Businesses are incorporating voice into the existing network infrastructure to eliminate the need for switching equipment ([28], [30]). With this, businesses gain the advantages of reducing long-distance call charges, allowing remote access to home or branch offices, establishing remote presence by owning local numbers, and enhancement of customer call centers ([28], [30], [21]). But still, the routing of VoIP calls over the Internet is common in many cases. Small businesses may not be able to afford private circuits between their offices [21]. Smaller providers can compete in the international market through the use of the Internet. Moreover, residential users may be willing to accept the lower quality of calls routed over the Internet for lower prices [21]. Even major VoIP providers use Internet paths as backup to their primary lines [6]. In addition, the development of global commerce leads many companies to enable customers to call through the Internet [28].

This thesis follows the style of *IEEE/ACM Transactions on Networking*.

Whether calls traverse the Internet or private networks, the difficulties are the same. Security threats to VoIP should be expected in either case as it has been shown that a majority of hacking incidents come from the inside [30]. VoIP, similar to data traffic, is also prone to denial-of service (DoS) attacks [26], worms, viruses, and Trojans [9] but here the stakes are higher; an attack can bring down the phone network [30], consume large bandwidth to block many calls [16], or even prevent emergency calls to police or fire departments [9]. Some vulnerabilities of VoIP equipment have been shown to exist. As a result of the large number of open ports and improper access control [23], it is possible to reconfigure the device remotely and plant malicious software [5]. And PC-end users are at even greater risk to attacks.

The challenges for achieving good quality may be magnified under the best-effort service of the Internet, but unpredictability exists on any packet network. As a result, error-correcting and/or concealing schemes are employed as protection against the corruption of the transmitted audio such that reliable (or near-reliable) communication is possible without the need for retransmission.

Forward-error-correction (FEC) protection is typically applied on individual calls on an end-to-end basis; therefore, many of the existing studies address the improvement of quality for a single call. Since these schemes generally require the addition of redundancy over the existing stream (except for basic error-concealment algorithms), it is necessary to consider the global picture. In [25], it is shown that a blind application of FEC does not always give a beneficial result; the optimal level of redundancy is heavily dependent on the loss conditions as well as the behavior and load of traffic in the network. Moreover, the FEC mechanisms which are usually employed for individual voice calls add a considerable amount of redundancy that cannot be ignored when the network carries a large number of FEC-protected voice streams [25].

Low-density-parity-check (LDPC) codes are an FEC technique with linear-time encoding/decoding times which can scale up to large block sizes; as opposed to the impractical processing times for earlier block codes ([22], [20]). Tornado Codes are a form of LDPC codes with very efficient software implementations which can sustain speeds up to 100Mbps [19]. For clarity, in this paper, these LDPC forms of FEC will be referred to as LDPC codes whereas all other FEC-based coding will simply be called FEC. Recovery rate and the overhead for LDPC codes is almost optimal as the block size is increased. However, interactive audio is limited by the end-to-end delay in order to achieve good quality [4]. Because of the lower bit rates and the small delay requirements of VoIP, LDPC codes may not be efficient for an individual call.

On the other hand, when we consider links in a VoIP provider's overlay networks or inter-office tunneling of voice calls, it is very likely that these links carry a large number of voice calls simultaneously. It may be possible to multiplex frames from different voice calls into a stream of blocks each of which is coded with LDPC. Figures 1 and 2 illustrate the proposed scheme for a VoIP provider's network. The VoIP service provider's access routers or gateways can employ LDPC coding on aggregated VoIP traffic with a common egress router of the provider's network. Figure 2 shows how frames from individual VoIP calls are scrambled into a large block, which is LDPC coded, packetized and sent over the network. The process is reversed at the egress point of the network to decode the original transmitted information.

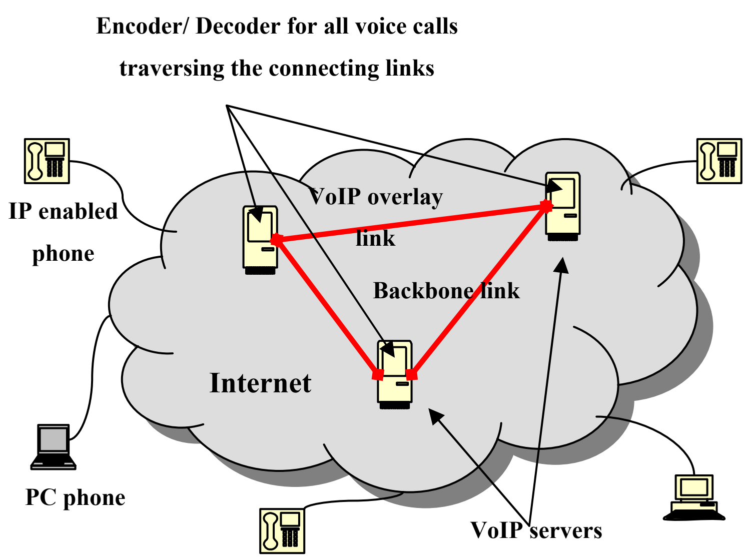

Figure 1: Configuration of a VoIP network.

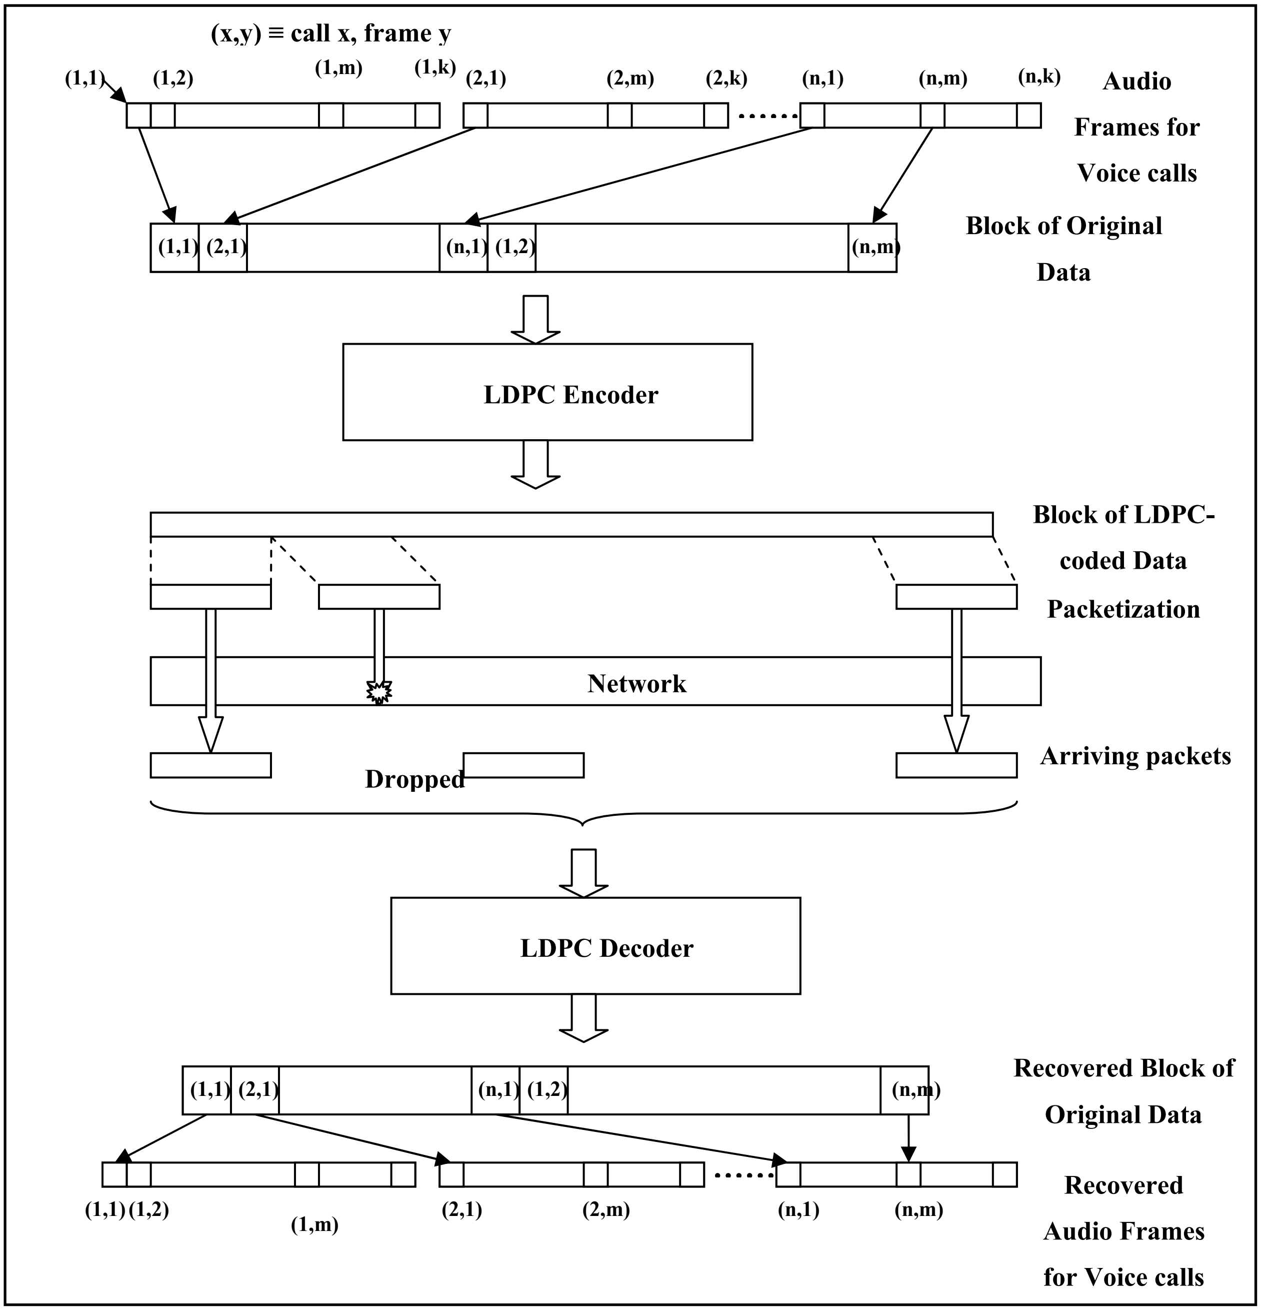

Figure 2: Process of aggregation, encoding, transmission, and decoding for VoIP calls.

With larger achievable block sizes through aggregation over multiple calls, the redundancy can be reduced compared to previous FEC schemes. With an appropriate choice of coding, the loss rates can be reduced to a tolerable level. As we will show later, aggregation enables reaching significantly lower output loss rates than possible with other FEC schemes.

Using high-bandwidth links can shorten the accumulation time before coding, which makes the delay comparable, if not better, than that needed for SFEC techniques.

Furthermore, with the existence of opposing traffic on the network, considerable damage can be done to the voice streams even with little overloading of the link. We will show that the recovery rate and the resulting quality for LDPC encoding are much better than that for SFEC with a comparable level of added redundancy.

2. PACKET AUDIO: CALL QUALITY AND SPEECH CODING

For audio to be digitally transmitted, various coding techniques are being used. Pulse Code Modulation (PCM), the basic type of waveform coding, retains almost original quality of the speech while requiring minimal processing. Its high bit-rate (64kbps) makes it rather restrictive for use on IP networks and particularly so for dial-up users on the Internet [17]. Therefore, alternative coding techniques must be used to compress the stream to lower rates. Several compression techniques have appeared which can reduce the rate to 8 kbps while still maintaining the excellent quality ([17], [18]). This is achieved at the expense of longer encoding times; however, real-time audio cannot tolerate large delays. Furthermore, encoding audio for unreliable transmission necessitates that no dependencies exist among the frames, otherwise deterioration of the quality from packet losses will be magnified [18]. The various speech codecs used for packet audio and their characteristics will be discussed further in section 2.2.

Because of the various factors that can affect the quality of a call, many measures of quality have been developed. Mean-opinion score (MOS) is a subjective measure based on human evaluation through listening tests [9], which makes it both impractical and time-consuming [12]. Because quality of real-time audio is affected by loss, delay, and throughput on the channel, MOS scores are often related to some range of these metrics as general rule of thumb [4]. However, these cannot accurately predict human assessment of quality ([12], [20]). Perceptual Evaluation of Speech Quality (PESQ) is one psycho-acoustic quality model which predicts MOS scores for speech segments ([31], [14]). The E-model is another quality model developed for the evaluation of telephone networks ([7], [14], [20]).

2.1 EVALUATION OF CALL QUALITY

Several objective and subjective quality measures have emerged in order to asses the quality of voice calls. Many factors contribute to quality degradation in VoIP; such as delay, loss, and distortion. Voice packets are delayed either at the sender (encoding and packetization), while in transit (propagation, transmission, and queuing delays), or at the receiver (buffering and decoding) [20]. And distortion is caused by encoding, loss of frames, and overflow of the playback buffer. Therefore, the receiver buffering which eliminates jitter results in a trade-off between loss and delay [20].

Human rating of quality is the ultimate goal of any quality metric, so quality is often judged by performing formal listening tests to obtain MOS scores ([4], [25]). The MOS measure has a range of 1(poor) to 5(excellent), and toll quality should have MOS ≥ 4 ([14], [20]). Quality is also considered acceptable for MOS≥ 3.6, which is the level which the PSTN provides [20].

Relating the perceived quality to network characteristics has been the focus of many studies [20]. For example [4] compares objective and subjective measures of users' perception in order to define tolerable levels of stream deterioration. In [4], it was found that 'good' quality audio is restricted by an end-to end delay of 150ms, a jitter of 20ms, and a loss rate of 0.5%; whereas others ([1], [24]) consider a 5% loss level as the threshold for 'acceptable' voice quality. Although these properties are easy to measure and control, the ranges are obtained by statistical analysis and they cannot be generalized to predict human assessment in any network situation ([12], [20]).

The PESQ measure (ITU-T Recommendation P.862) ([31], [14]) is a model developed to evaluate speech quality in telecommunications networks by calculating a MOS score [12]. Studies show that its values correlate well with human ratings ([13], [14]).

The E-model was developed to evaluate transmission quality. The quality measure, which is called the R-factor, is evaluated with the following equation ([7], [20]).

$R = R_o - I_s - I_d - I_e + A$ ([7], [20])

R_o: The effect of noise and loudness

I_s: Impairment which is concurrent with the speech signal (e.g. quantization)

I_d: Impairment from one-way delay, talker/listener echo, and interactivity

I_e: Impairment from signal distortion (e.g. low rate codecs, lost frames)

A: Advantage; when users sacrifice quality for convenience (e.g. cellular systems)

The terms in the equation cover the various impairments to quality that a call suffers from one end to another. The effect is additive when it is converted to the R psychoacoustic scale [20]. Toll quality corresponds to an R value of 80 or higher, whereas R values above 70 are considered of acceptable quality [20]. The factors (R_o, I_s, and A) are network-independent and we will not be concerned with them here.

Some limitations to these models have been shown to exist. PESQ only accounts for speech quality with no consideration for delays, echoes, and other factors which would affect call quality [14]. Although the E-model accounts for these factors, it also lacks in other areas. The E-model has been extended in [7] and [14] to account for bursts and the recency of losses in order to predict human ratings more accurately.

2.2 SPEECH PROCESSING AND COMPRESSION

Waveform coding is a lossy speech processing method which can provide good quality speech with a rate of 16 kbps. The simplest form of waveform coding is PCM. However, because of its large bandwidth requirement (64 kbps), more compressed codes are often used in IP telephony applications [18]. ADPCM coding can achieve rates in the range of (16-40 kbps) while not sacrificing much quality for the lower bit-rates. When very low bit rates are needed, vocoding is used. LPC is one such code, which gives intelligible speech for a rate of 4.8kbps or below. GSM is also a popular coding scheme which is a hybrid of both waveform coding and vocoding [18]. Table 1 summarizes the characteristics for several codecs taken from ([18], [20], [10]).

The encoding delay consists of three components: the processing delay, the framing delay, and the lookahead delay. The frame delay is the length of the voice compressed in a single packet (or frame) and the lookahead is the length of the samples from the next packet which are needed to code the current packet [17]. GSM and LPC both have a total encoding delay of 20 ms, whereas delays in PCM and ADPCM are negligible in comparison [10].

Table 1: Properties of different coding standards.

Codec	PCM	ADPCM	GSM	LPC
Rate (kbps)	64	16 - 40	13	4.8
MOS	4.4	2 - 4.3	3.7	2.6
R	94.3	39 - 89	70	50
Frame Delay (ms)	0	0.125	20	20
MIPS	0.01	2	6	7

The table also shows the R-factor for each codec. This value is the default without consideration for losses. For a PCM stream, $(R_o - I_s)$ has a value of 94.3 ([12], [20]) and R can be calculated as $(R = 94.3 - I_d - I_e)$. Delays up to 175 ms have little effect on R and I_d is negligible [12]. And the loss impairment values of PCM are depicted in [20]. The R-factor approximately declines by 4 for a loss of 1% and drops by 2 for every 1% after that [20]. Since the higher rate ADPCM codec does not cause much loss in quality compared to PCM, the factors for calculating R for ADPCM will be substituted by the PCM estimated values mentioned above. This should not affect the comparison between LDPC and SFEC at all but it will give slightly higher R values for both of them. This difference would not generally be sufficient to change the rating of the audio stream from one category to another.

3. ERROR CONCEALMENT AND CORRECTION FOR PACKET-AUDIO

In the Internet, a majority of loss bursts affect only one or two packets and the distribution of larger bursts is almost geometric ([2], [11], [3]). Nonetheless, a series of short bursts can reduce the quality of an audio signal [25]. Furthermore, backbone connections constitute a portion of the path for long-distance VoIP calls, as well as calls routed over a combination of PSTN and Internet paths. These links can cause poor performance for audio because of their large and variable delay properties even if the averaged loss rate is low ([20], [11]). Due to the many problems involved in admission control and reservation methods, many believe end-to-end schemes are more desirable [25].

Some currently employed receiver-only reconstruction techniques require no redundancy (e.g. silence, white-noise insertion, or repetition). Silence replacement of lost samples works best for small loss rates with small packet sizes, and white noise is just as easy to produce and gives better performance; but both perform poorly for large packet sizes [11]. Repeating the last received packet can only be maintained for a period of 80 ms (length of phoneme) without damaging speech characteristics [11]. Accordingly, there is a need for error-recovery techniques which have the capability to rebuild the stream from the added redundancy [25].

One form of error correction uses block codes (FEC), which are characterized by the values (n, k); where n is the number of encoded packets per block, and k is the original number of packets. The receiver can recover any lost packets after receiving any k of the n packets. The block code used in [31] cannot be performed over a large number of packets because of processing requirements. The redundancy used is very large and may incur up to 50% overhead to result in good quality in the face of losses [31].

Signal-processing based FEC (SFEC) [25] is another technique which uses a lower-quality delayed audio stream which is attached to the original stream. It is also possible to include multiple redundant streams ([2], [25]). This coding can be used for larger packet sizes [11]. When the audio is replayed, the low-quality sample replaces any original sample which is lost. In the next two subsections, we will discuss error-correcting techniques and their use in VoIP applications.

3.1 SIGNAL-PROCESSING BASED FEC (SFEC)

SFEC is used extensively for packet-audio. Many Internet applications such as FreePhone employ it to protect calls from losses [17]. The low quality stream uses codecs such as LPC or GSM which are both CPU-intensive ([2], [25]). LPC has a lower bit-rate than GSM, but its quality is also lower.

The redundant low-quality version of a packet is piggy-backed onto a subsequent packet. At the receiver, the low quality frame replaces the original frame when it is lost ([11], [2]). SFEC also uses error concealment techniques when neither the original nor the redundant packet is recovered [25].

It is possible to vary the number of redundant copies or the delay of the redundant frame from the original one to produce any number of SFEC methods with different redundancy rates and recovery levels [2]. The overhead incurred is lower than that for FEC block codes, but the user-perceived quality – although still acceptable – is worse than that achieved with FEC [31]. An obvious advantage of SFEC over FEC is that the SFEC stream can still recover to some extent when losses are extremely high even though the quality suffers, whereas in FEC methods this situation would result in the loss of the whole encoded block. Consequently, SFEC is very suitable for wireless applications [2].

The larger the separation between the redundant and the original frames, the more resilient the SFEC stream is to burst losses. However, a large separation is not favorable for reducing packet jitter and delay, since the receiver must await the redundant frame whenever a packet is lost ([2], [25]). In addition, a larger number of redundant copies is preferable, but the overhead may also become excessive. In [25], it was shown that the use of SFEC without consideration of the load on the network can worsen the congestion problem and overload the network [25].

Performance of SFEC is often evaluated using relative reward [3]. Relative reward measures the ratio of packets which are recovered with SFEC to the actual number of packets arriving at the receiver with no regard for the quality of those packets [9]. However, we will use the E-model [7] to estimate the resulting quality in order to compare SFEC with LDPC.

3.2 FORWARD-ERROR CORRECTION

In block codes, such as Reed-Solomon codes, encoding/decoding times are quadratic in the length of the block [22]. In order to encode blocks for network applications, these blocks have to span several packets; thus resulting in impractical coding times. Furthermore, for an (n, k) block-code, k of the n packets have to be successfully received in order to recover the original packets [22]. Since coding times restrict us to a small number of packets, a small burst of errors can make recovery impossible for that block. As network transmissions are unreliable, the use of such codes is difficult because of the introduced delays and the low resilience to packet loss.

Another class of FEC codes is low-density parity-check (LDPC) codes. LDPC processing times are linear with the size of the coded block and the recovery rate improves with increasing block size ([22], [20]). These properties make it feasible to use LDPC codes for real-time network applications. A large number of packets can be encoded together in a single block using efficient encoder implementations and the encoding/decoding times are not restrictive for real-time audio.

LDPC codes can be represented as sparse bipartite graphs with the encoded block on the left and the constraints on the right. Figure 3 shows how such a code would be constructed as a graph, however, an actual implementation of the code would use a much larger block size. The encoded bits are a combination of original data with the redundant bits. Constraints are the XOR of their adjacent nodes and they are implicitly assumed to be zero; therefore constraints are not sent with the block ([27], [22]).

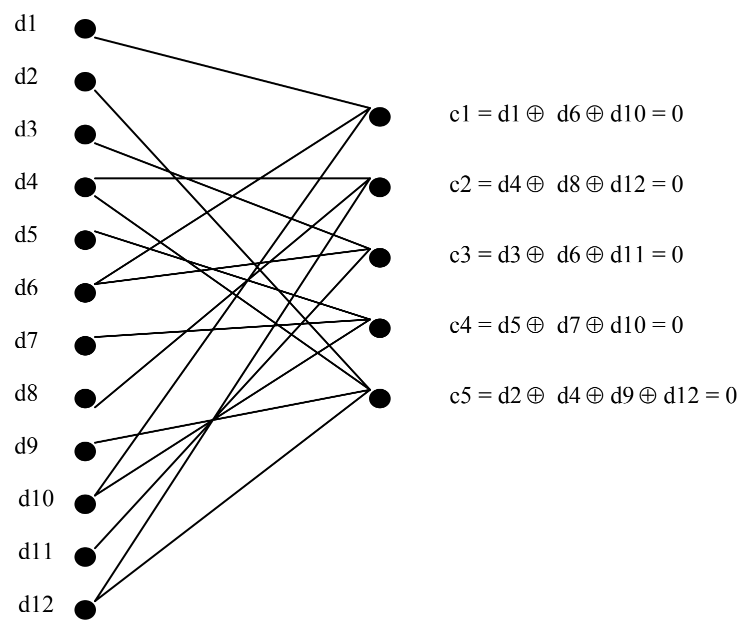

Figure 3: Sparse graph representation of LDPC codes.

When packets are lost in the network due to congestion or channel losses, the decoder at the receiver sees a number of bit errors. Since both the sender and receiver have agreed on a particular structure for the graph (or the associated code), then the decoder can recover missing bits (i.e. erasures) using the XOR relationships with its adjacent constraints in the graph ([27], [19]). As a result, if the erasures are excessive, they may not all be recovered. Loss recovery for LDPC codes is almost optimal, transmission is just below the channel capacity (capacity of the channel is (1-p); where p is the loss rate), and a random p fraction of packets which are lost can be recovered with a high probability [19]. We will define the input loss rate as the packet error rate incurred by the network, whereas the output loss rate is the fraction of packets that are not recovered after the decoding process. A block of data with original size k and which is coded with redundancy m has a rate defined as *k/(k+m).* The redundancy is *m/k* and the overhead is *(1-rate)%*.

4. AGGREGATE CODING OF VoIP TRAFFIC

The performance of LDPC codes improves as the size of the coded blocks (block size) increases to large values [19]. For these sizes, the overhead required to recover lost bits with a high probability is only slightly above the loss rate on the path. A single voice call cannot accumulate enough bits for encoding in a slot of time that is considered acceptable delay. The benefits of low overhead can however be gained when packets or bits from multiple calls are aggregated and coded together.

The block of audio samples will be encoded with LDPC and then fragmented into packets. If we assume an independent bit error rate on the channel, the packet error rate (PER) increases with the size of the packets. And those packets which are corrupted will be dropped at the network level. So both the input and output loss rates are decreased as a result of either increasing the number of packets per block or decreasing the size of the packet. However, the size of the packet can only be minimized so far without the effect of excessive communication overhead from the packet headers.

The parts of the lost packets will either be fully or partially recovered depending on both the input loss rate and the redundancy of the code. The loss rate of the decoded output will almost always be considerably less than that of the channel. In addition, because of the nature of the codes, although the lost bits are consecutive in the encoded block, they are spread out randomly among the decoded block. This may have minimal or no effect on some samples in the block. Also because packets are shared by multiple streams, a burst of errors will affect individual streams by corrupting the few samples which are contained in these packets.

Moreover, recovery of a lost packet in LDPC retrieves original quality of the frame, whereas in SFEC the recovered frame has lower quality. It has been shown that several consequent losses can harm voice quality ([7], [20]). When the losses for a single phone call are high, the quality degradation may be noticeable. But aggregation with LDPC will spread out the losses and the quality degradation will not be localized on any of the calls. With large block sizes, LDPC can compete with the redundancy levels of SFEC while achieving better resultant quality.

5. RESULTS AND EVALUATION

The following sections show the results of evaluating aggregation of VoIP calls with LDPC. In section 5.1, we demonstrate the level of recovery that these codes can achieve at various block sizes and with different number of packets per encoded block. Following that, the results of section 5.2 will show that aggregation offers many advantages compared to techniques which protect single calls. And finally in sections 5.3 and 5.4, a set of network simulations are used to compare the resulting quality of both the LDPC and SFEC encoded streams to confirm that the protection that LDPC provides to the voice streams exceeds that of SFEC applied individually to each call.

5.1 RECOVERY PERFORMANCE OF LDPC

These results are obtained using a C++ program that encodes a block with LDPC, simulates an error model for the channel, and then decodes the arriving data. The encoding process is a matrix multiplication, so we use an input block consisting of all zero bits to eliminate the need for any encoding. This doesn't affect the recovery of the code since the recovery capabilities depend only on the location of the erasures and not their actual values. The loss model used is a Gilbert-Elliot model with two states: a good state (probability of delivery is 1) and a bad state (probability of delivery is 0) [29]. The average loss rate is modifiable, but the burst length is set to an average of two packets. The structure of the LDPC code is also optimized to overcome bursts of size two. However, it is found that the use of a randomized structure does not affect performance much when the number of packets per block is high. The decoder iterates a number of times attempting to recover the lost bits and after its completion the output loss is measured. The output loss rate is averaged over a large number of simulated block coding and transmissions.

From the simulations of LDPC coding on various block and packet sizes, it is found that the performance was very similar among pairs that had the same number of packets per block regardless of the actual sizes; for example (16KB block, 128B packets) behaves similar to (128KB block, 1KB packets). For this reason, the experiments and simulations in the subsequent sections are done based on the assumption of a block size of 32 KB and a packet size of 128 B because these are reasonable sizes for VoIP.

Figure 4 illustrates the maximum loss within a single block which can be completely recovered by the decoder with respect to the redundancy of the code for 128KB block size and 1KB packets. The relationship is linear and the amount of redundancy needed is shown to be only slightly above the channel loss rate.

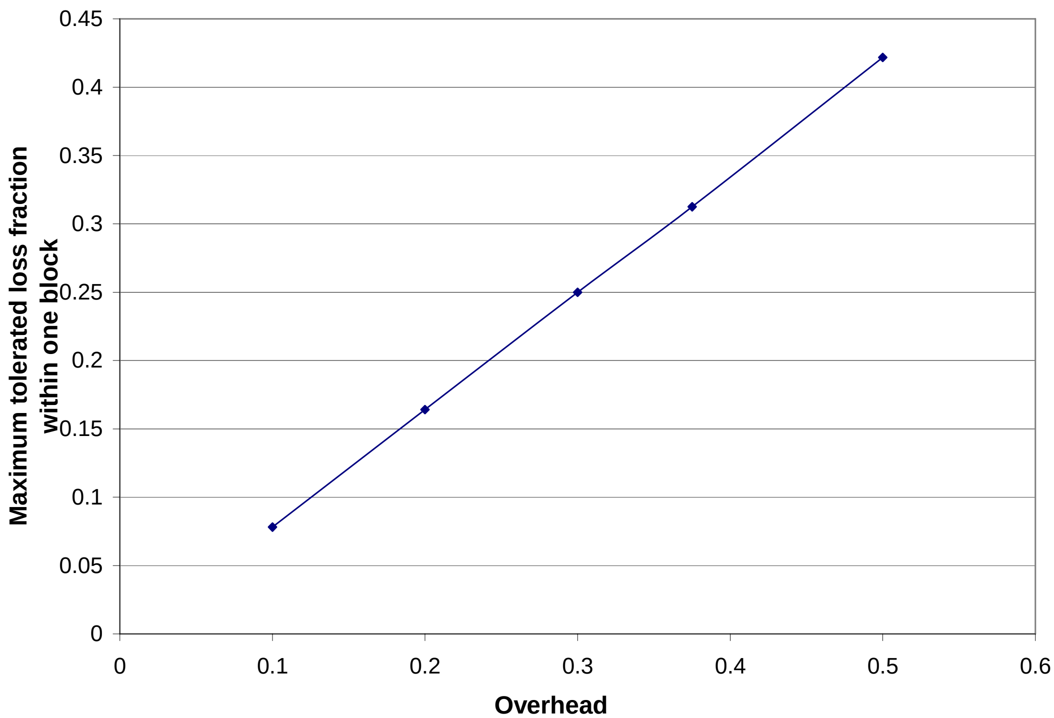

Figure 4: Overhead vs. maximum loss per single block.

In Figure 5, we show the results for a 32 KB block and 128B packets. The graph demonstrates the redundancy needed to guarantee an output loss level of 0% and 0.5% at different average input loss rates. This figure reaffirms the results from Figure 4 of the correlation between the redundancy and the network losses. (The particular implementation of the encoder program which was used restricted the possible values of the redundancy especially in the higher range (0.5 – 1). For the 0% output loss curve, any loss less than 0.1% was considered as having zero loss rate. Otherwise, the next lower rate, which would give a zero loss rate, will overstate the amount of redundancy needed to achieve full recovery.)

From Figure 5, we also note that any particular code (with a fixed redundancy level) can tolerate an increase in loss rate on the channel by 2-4% without much deterioration. For example, an LDPC code with redundancy (0.33) can completely recover losses up to 10%. However, the loss rate can be increased up to 14% with barely any degradation to the audio stream (at an output loss of 0.5%).

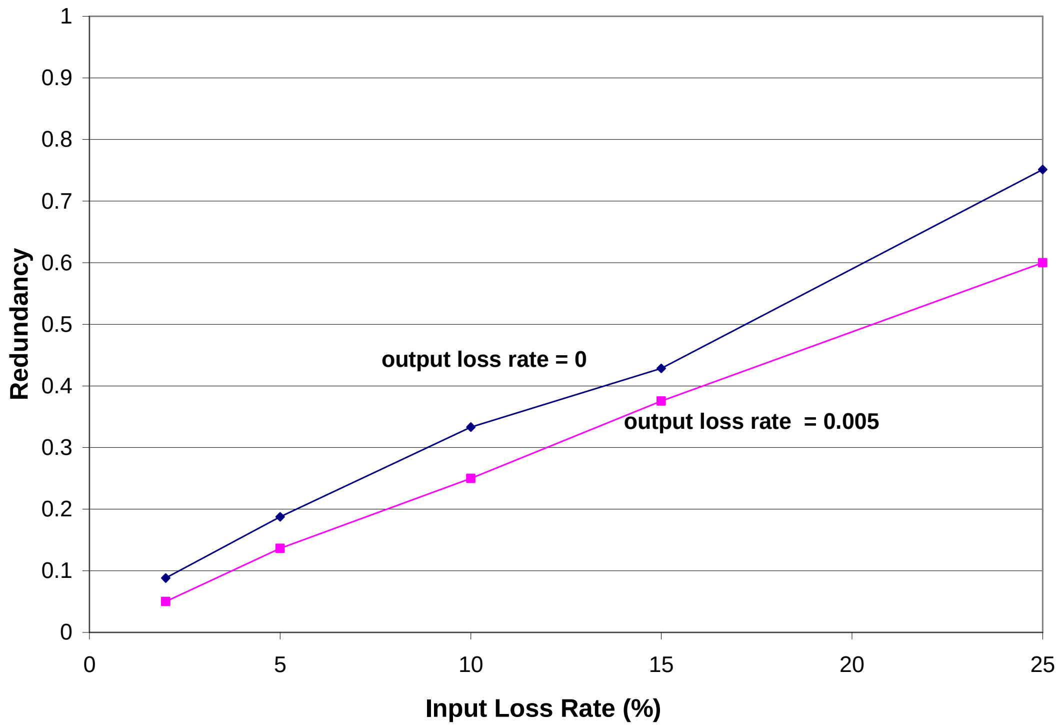

Figure 5: Input loss rate vs. redundancy (32KB block, 128B packet).

5.2 AGGREGATION VS. SINGLE CALL

The performance of other FEC block codes on single calls is poor in comparison with LDPC. From [31], it was shown that using FEC with 50% overhead, a 38% input loss rate was only reduced to 15%. To bring the loss to a tolerable level of 5%, an overhead of 75% was needed [31]. On the other hand, LDPC over multiple calls with 50% overhead code was found to recover to an output loss rate of 4% in our experiments.

SFEC codes differ from LDPC codes in that their recovery rate is independent of the level of redundancy added, and the redundancy for SFEC affects only the recovered audio quality. The output loss rate is measured here based on all packets received. Since the SFEC resultant stream contains both original and low-quality frames, there is no consideration for quality in this metric. The quality of recovered packets in SFEC depends on the particular low-quality codec used (such as LPC or GSM). As a result, to achieve the same resultant quality, LDPC can tolerate a higher loss rate than SFEC. This effect is evident by a comparison of block erasure codes and SFEC for wireless LANs in [31]; although decoder output loss rates were almost twice as much in FEC than SFEC for an input loss rate of 38%, the quality for SFEC was rated with a PESQ of 2 while FEC had a PESQ score of 3.5.

All the results for this section are for 32KB blocks and 128B packets. Figure 6 and the larger scale image of Figure 7 show how the redundancy affects the output loss rate at each particular input loss level (the data lines are labeled by their input loss rate). The horizontal lines belong to the SFEC coding of an individual call while the other curves are the results for LDPC. Even though we depict SFEC's output loss rate as constant up to 0% redundancy, it is to be noted that below a certain level of redundancy, the recovered sample will be of unacceptable quality. Furthermore, the redundancy level of SFEC cannot be reduced indefinitely. The size of the compressed low-quality packet also places a restriction on the minimum redundancy. If we assume an ADPCM-coded original stream, then using an LPC codec (4.8 kbps) for redundancy, would result in a redundancy of 0.15. If a better quality codec, such as GSM, is used for the backup stream, the redundancy in SFEC will be 0.4.

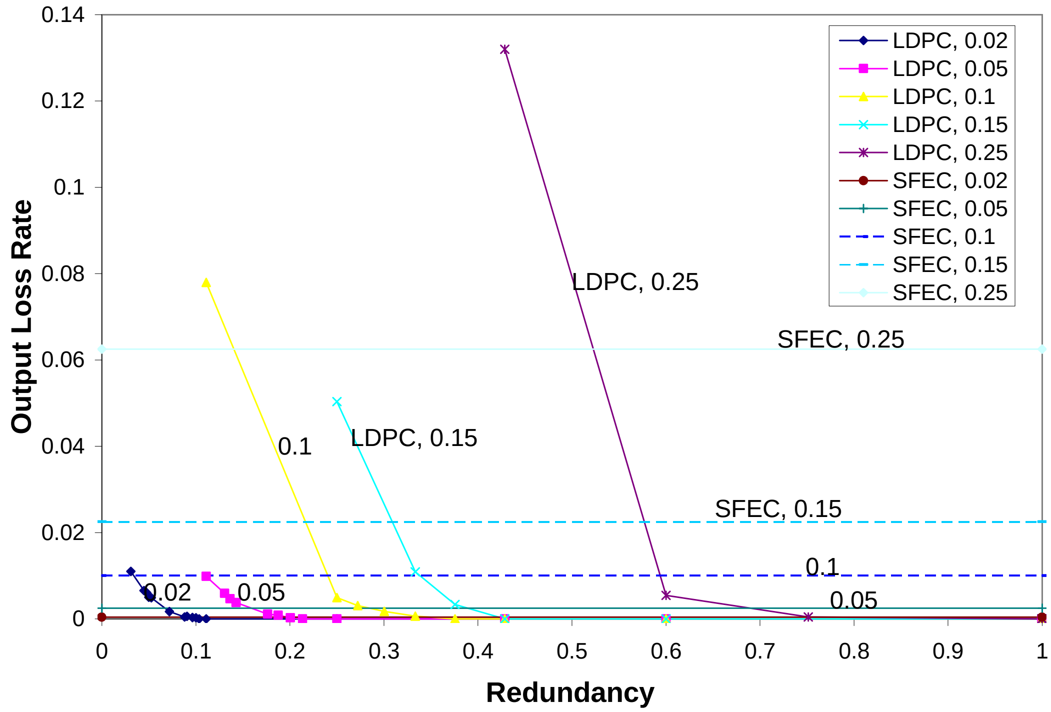

Figure 6: Output loss vs. redundancy for different input loss rates.

From the graphs, it is observed that LDPC surpasses SFEC's output loss rate at some level of redundancy for each different input loss rate. This shows the strength of the proposed approach in reducing the output loss rates to levels significantly lower than those feasible with SFEC. LDPC can give better audio quality than SFEC at the same recovered rate since SFEC's recovered sample is of lower quality.

The redundancy of SFEC is independent of the loss rate. At higher loss rates, the output loss rate cannot be reduced below certain level using SFEC. However, better audio quality can be achieved with LDPC even at higher loss rates when higher redundancy can be employed. We only considered complete recovery of all the bits in the lost packets in our experiments with LDPC codes. This results in the higher loss rates for LDPC codes at lower redundancies.

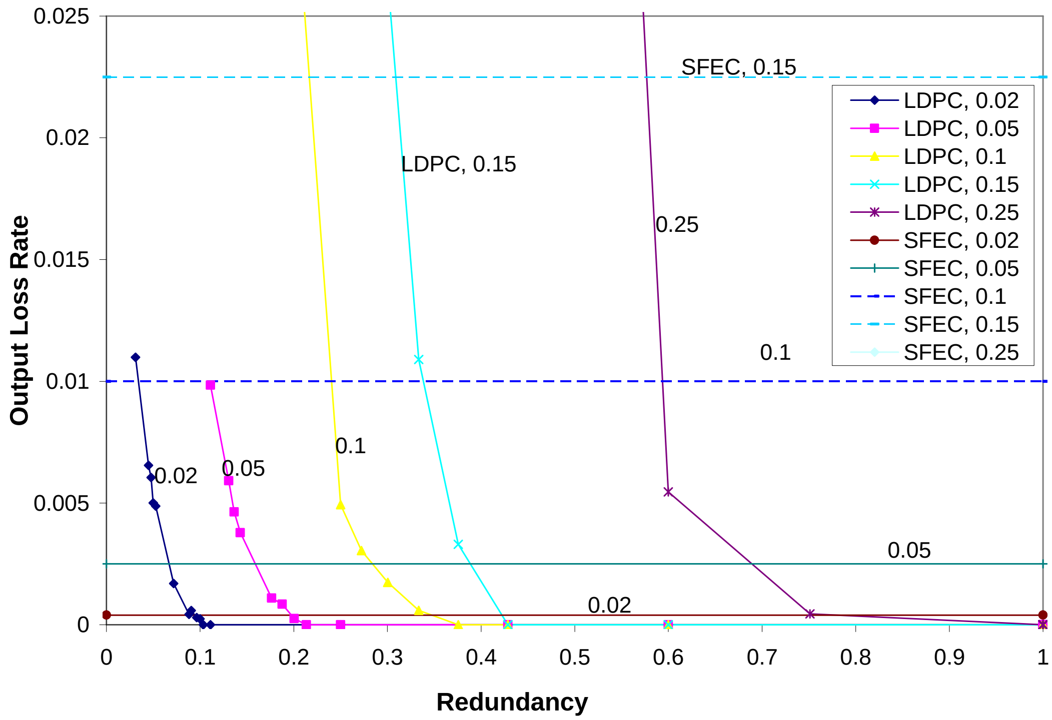

Figure 7: Output loss vs. redundancy for different input loss rates (closer view).

Assuming that the voice samples are accumulated from voice calls over a period of 50 ms, the block size is restricted by the link speed. As the link speed increases, the number of packets per block for LDPC coding also increases giving better recovery rate. When using SFEC for a single voice call, the rate of the audio stream restricts the block size and therefore it cannot benefit from the increased link speed.

Figure 8 shows the change in the output loss rate with redundancy of the code for aggregate-coding of a combination of calls over different link speeds as well as a single voice call, when the input loss rate is 10%. The link speeds are calculated based on a block length of 50 ms. The loss rate for SFEC is fixed and is independent of the input loss whereas the LDPC codes generally improve as redundancy is increased.

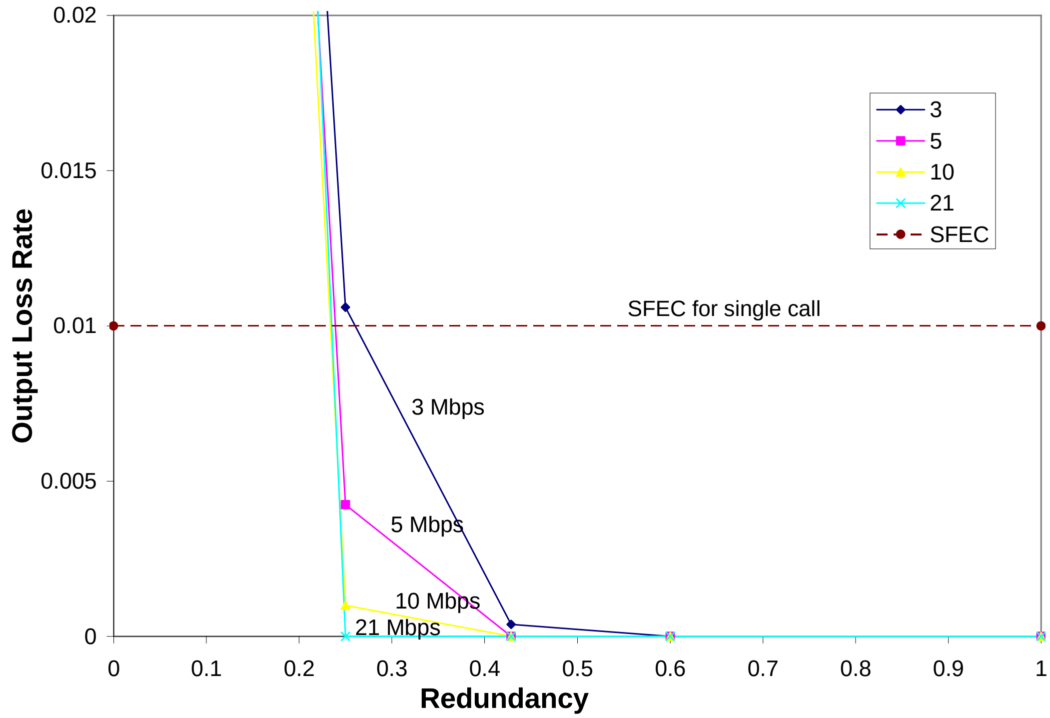

Figure 8: Output loss vs. redundancy for different link speeds; assuming loss rate = 0.1, delay = 50 ms.

At redundancy levels above 0.45, LDPC coding completely eliminates any errors regardless of the block size at the different link speeds. In the range of redundancies (0.25 – 0.45), LDPC can still outperform SFEC. In this range, it can be seen that better performance is achieved at higher link speeds (hence a higher block size under the same delay constraints). For redundancy levels less than 0.25, SFEC performs better, but the recovered signal will have lower quality. And below a certain redundancy, the SFEC recovered quality will be deemed unacceptable. As mentioned before, SFEC (with LPC) has a redundancy of 0.15 and SFEC (with GSM) has a redundancy of 0.4. For the assumed input loss rate of 10% and an SFEC output loss rate of 1%, the R-factor for the resultant stream will be 85 for SFEC-LPC (86 for SFEC-GSM) using the E-model equation from section 2.1. To recover to the same R value for LDPC, the stream can tolerate up to 3% losses (as opposed to the 1% loss for SFEC). This loss rate can be achieved by an LDPC code of redundancy 0.2. The code has comparable redundancy to that of SFEC-LPC (0.15) and a much lower redundancy than SFEC-GSM (0.4).

Furthermore, for the same block length, shorter delays can be provided at higher link speeds at a similar level of protection against packet losses. Figure 9 illustrates this for a delay of 10ms.

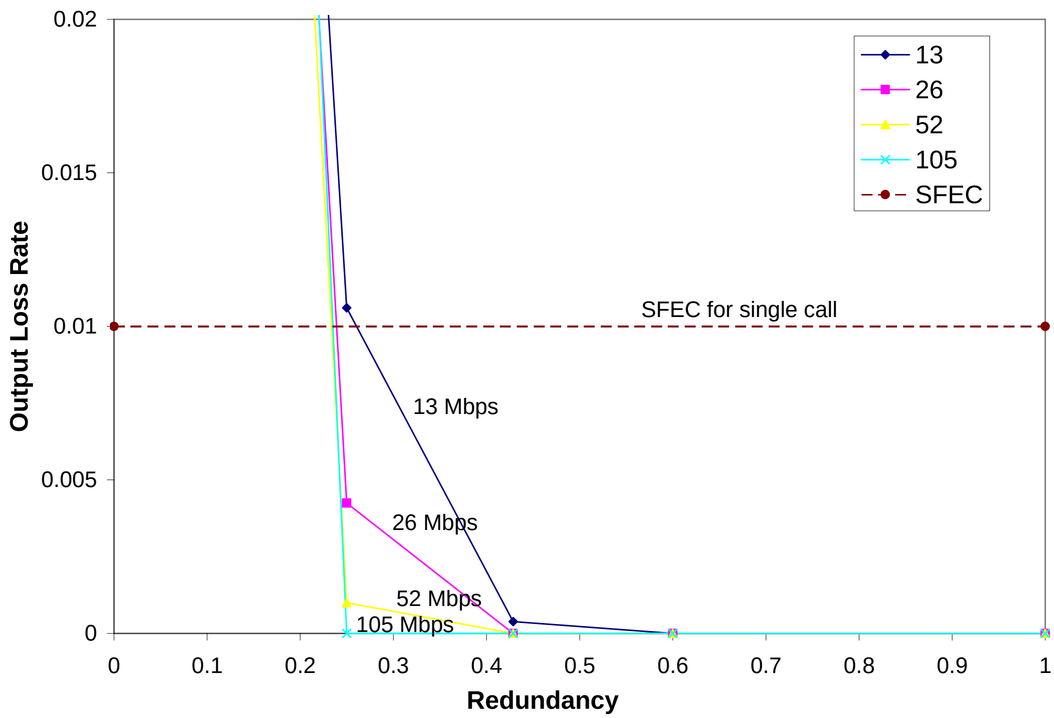

Figure 9: Output loss vs. redundancy for different link speeds; assuming loss rate = 0.1, delay = 10 ms.

5.3 EVALUATION OF THE PERFORMANCE BY NETWORK SIMULATIONS

In this second part of the evaluation, we will assess the ability of both SFEC and the proposed LDPC scheme to recover from packet losses at congested links in the network. The opposing flows which cause the overload represent either a DoS attack or a transient overload of the link. Through network simulations with NS-2 simulator [15], several such scenarios will be developed in order to study the performance of both coding techniques. The simulation setup consists of a single link with a speed of 10 Mbps and a delay of 10 ms, and this link holds both the VoIP traffic and the opposing traffic. A VoIP stream is represented as a constant bit-rate (CBR) flow with a rate of 64 kbps with packet sizes of 128 bytes. There are 82 VoIP flows which consume approximately half the bandwidth on the link. The flows are delayed with respect to one another so that the arrival of voice packets is approximately evenly spaced. The opposing traffic, also consisting of CBR traffic, takes different forms in various stages of this evaluation from being a single flow with equally spaced packets to a set of flows which periodically send out their packets in bursts. For clarity, the voice flows will be referred to as VoIP while the opposing traffic will be called CBR.

From the resulting traces, the loss rate of the recovered stream is found for both SFEC and LDPC. For SFEC, we consider the case of a single redundant stream (a low-quality frame is attached to the third subsequent packet) and three redundant streams (one low-quality frame is attached to each of the three subsequent packets). The quality of the outgoing stream is then estimated based on the assumption of an ADPCM-encoded audio stream with a redundant low-quality coding of either GSM or LPC. Table 2 shows the rate of the resulting stream (ratio of the redundancy to the total bit-rate) and the intrinsic quality (R) for all the cases considered. The intrinsic quality is the R-factor of the encoded stream without any loss impairment. The calculated rate in the table is based on the ADPCM’s 32 kbps bit-rate to assess the best-case scenario when using an SFEC stream with a constant resulting bit-rate of 64 kbps. SFEC3 has the least amount of redundancy (13%) since it has the highest rate (0.87), whereas SFEC2 has the most with 55% of redundancy.

Table 2: Characteristics of the evaluated SFEC streams.

		Original Stream			Redundant Stream				
SFEC Encoding		Codec	Bit-rate	R	Copies	Codec	Bit-rate	R	Rate
SFEC1	ADPCM /GSM1	ADPCM	16-40 kbps	39 -89	1	GSM	13 kbps	70	0.71
SFEC2	ADPCM /GSM3	ADPCM	16-40 kbps	39 -89	3	GSM	13 kbps	70	0.45
SFEC3	ADPCM /LPC1	ADPCM	16-40 kbps	39 -89	1	LPC	4.8 kbps	50	0.87
SFEC4	ADPCM /LPC3	ADPCM	16-40 kbps	39 -89	3	LPC	4.8 kbps	50	0.69

The recovered stream will be mostly ADPCM coded and the remaining part will be of the low-quality coded. Based on the proportions of these two parts, the quality of this stream is estimated. In addition, if neither the original nor the redundant packet is recovered, there will still be some packet losses. The quality of the resulting voice is estimated based on the E-model while using the intrinsic quality and loss impairment values of PCM as in [20] to get the best-possible resultant quality.

To evaluate the performance of LDPC, the traces are fed to the encoder/decoder program and the resulting losses are found as stated in the previous section. The quality is found by accounting for the loss impairment using the measured PER. Although the measured bit-error rate (BER) is always less than the PER and the use of loss concealment could help recover from lost bits to form a lower quality sample, it is not considered in the quality estimations since there is no way to quantify it without tests.

In section 5.3.1, the different scenarios will be presented and then in 5.3.2 the results of the comparison of the two methods. And finally, in section 5.3.3, we will discuss the effect of changing the different parameters on the voice traffic.

5.3.1 The Opposing Traffic

The small packet size of the VoIP streams makes it more resilient to losses than a flow with larger packet sizes as the link is overloaded. When the CBR flow is set to a packet size of 1000 B with the VoIP traffic at a size of 128 B, it is found that losses experienced by VoIP were negligible even at very high link loads. Figure 10 shows the actual loss rates seen in the network and the loss rate after the SFEC recovery mechanism is employed (assuming one-copy SFEC). At these high loads, the majority of the losses are suffered by the CBR traffic because the queue in the link is constantly full or near-full which causes the larger size packets to be constantly dropped while the smaller ones find enough buffer space to be queued. According to the recorded loss rates, no coding may be needed at all and receiver-only techniques may be sufficient to overcome any remaining losses. It should be noted that the loss values in the figure are averaged over the VoIP flows, and although the majority have negligible loss rates, it was only one or two flows which suffered very high losses. If we are to use coding in such a case, applying SFEC coding will lower the loss more (as shown in the figure), but using LDPC would allow for lower redundancy judging from the results of section 5.2. And at the same time the losses will be spread out so that none of the flows would suffer and all calls would have excellent quality.

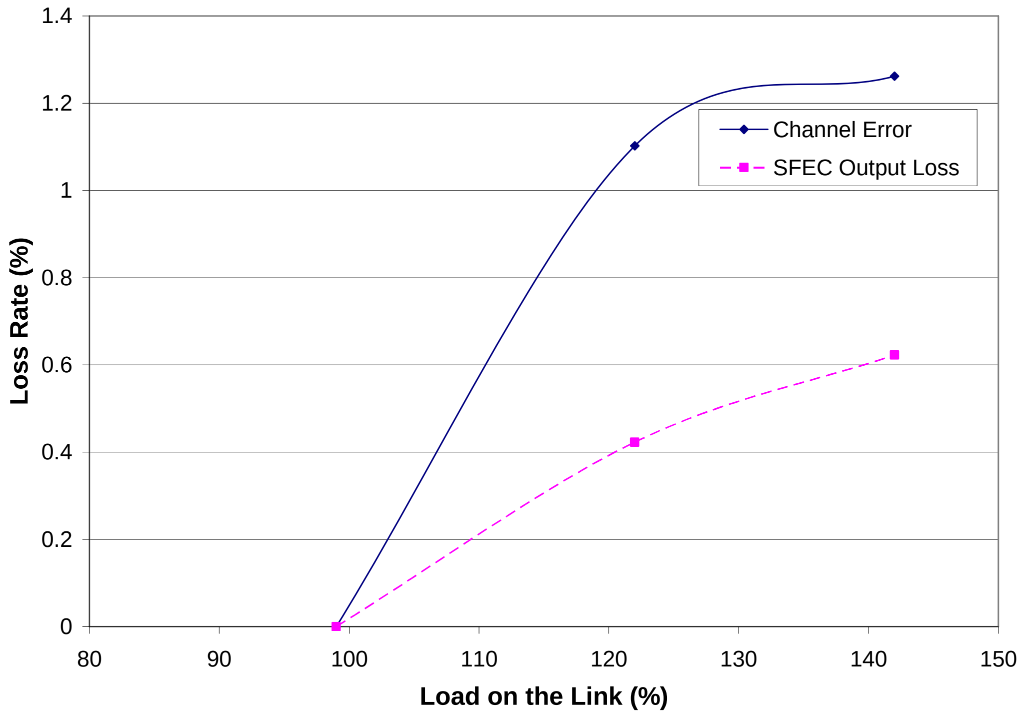

Figure 10: Resilience to opposing traffic with large packet sizes.

In the next set of experiments, the packet size for all the flows on the link is set to 128 B with a single opposing CBR flow. The losses suffered by both VoIP and CBR are comparable in this case. In Figure 11, we show the losses in the network as well as the resultant losses when using both SFEC and with LDPC at different rates. The measured SFEC losses are for a single redundant stream (i.e. SFEC1 {rate 0.71} or SFEC3 {rate 0.87} from Table 2). Since losses above 5% are considered damaging to the quality of voice [1], there is no need to consider resultant loss levels larger than 5%. LDPC code with a rate 0.5 always achieves lower losses within this range. LDPC 0.71 recovers from losses up to 25% and will give near-original quality while the SFEC experiences some losses which somewhat deteriorate the quality of the voice. LDPC 0.833 sustains excellent quality until about 15% after which the losses become too excessive. In all cases, an overload of 20% can be overcome with either technique. However, these load and loss levels are rather high for any network with over-provisioned links or with privately-owned links. It is reasonable to consider that these losses can occur transiently either as a result of a temporary overload or due to an attack.

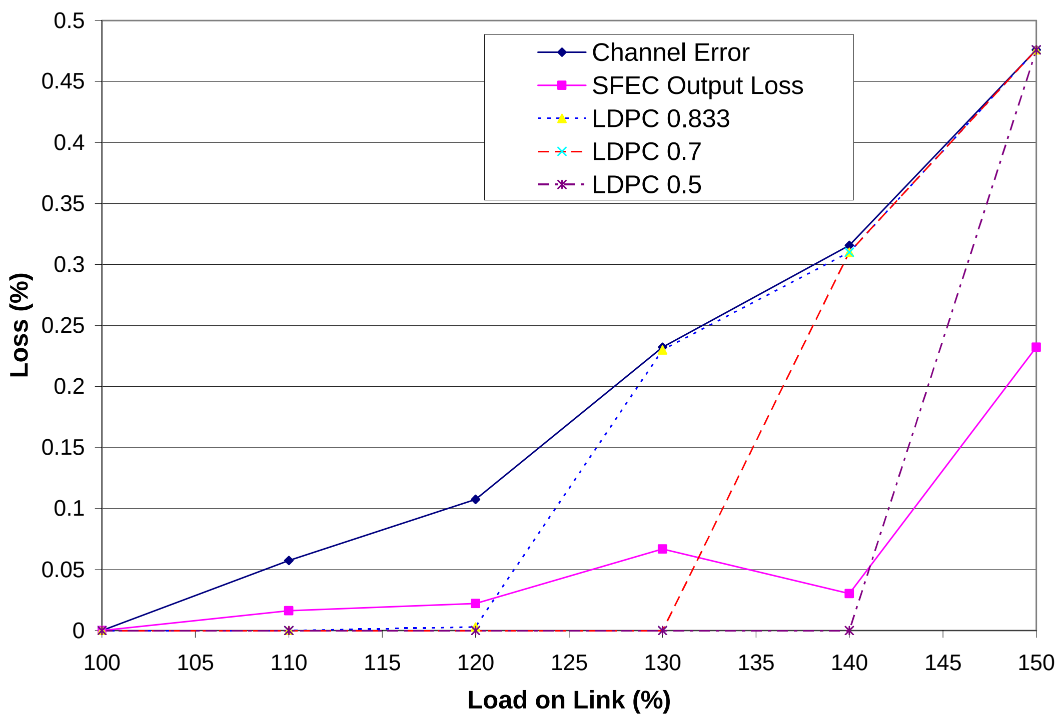

Figure 11: Recovery of SFEC and LDPC with a single CBR flow at different loads.

Table 3 shows the resulting quality of the voice estimated using the results for 20% overload. The table shows the level of redundancy needed, the resulting loss, the quality (R) averaged over all the VoIP flows, and the percentage of the calls which have toll quality ($R \geq 80$). It can be seen that LDPC 0.842 with almost comparable redundancy to the SFEC3, can achieve toll quality for all the flows although the average quality is slightly lower than that of SFEC3. However, with the addition of a small amount of redundancy (LDPC 0.833), the calls will all have excellent quality exceeding the rating of SFEC3.

Table 3: Resulting quality for a load of 20% (11% loss) for a single CBR flow.

Coding	Rate	Loss	Average R	% Toll
SFEC1	0.710	0.026	89.0	91.5
SFEC2	0.450	0.021	89.9	97.6
SFEC3	0.870	0.026	87.4	91.5
SFEC4	0.690	0.021	88.3	97.6
LDPC	0.833	0.003	92.9	100.0
LDPC	0.842	0.044	82.2	100.0

In the following, we will be considering the case of multiple CBR flows in an attempt to simulate a DoS attack which would try to maximize losses for the voice traffic with as little load as possible. In these simulations, the CBR flows are similar to the voice flows in the packet size as well as the bit-rate. The flows would all start at the same instant and send their packets so that they all crossed the link in one burst. This would recur at regular periods equal to the inter-packet interval in the flows. Having the flows all send at the 'exactly' the same time instant reduces the impact of an attack and probably may not be feasible in a real-life situation. As a large number of packets arrive at the queue at once, the queue fills up and the rest of the attack packets are dropped. Unless the voice packets are unlucky enough to arrive in the short period of time when the queue is full (50 – 100 μs), it is unlikely that it will be affected at all by this traffic no matter how high the load on the link becomes. And this is in fact what the simulations show; losses for VoIP are negligible even at 40% overloading when the attack traffic arrives simultaneously. However, the situation arises again in some of the simulations where few calls experience irrecoverable losses while the rest are not affected at all. As stated earlier in this section, LDPC – with lower redundancy than SFEC - can alleviate this situation while SFEC barely recovers from any losses within these streams. Since the average losses on the channel are low, the LDPC encoding will be able to recover fully. If any losses are remaining, they will be spread out among the VoIP calls and all calls will show excellent quality.

Subsequently, the CBR flows are spaced by the time it takes to transmit a single CBR packet from the queue. This appears to increase the impact of attack since most of the losses from overloading are suffered by the voice traffic. Table 4 below shows the loss rates seen by both VoIP and CBR flows in a set of scenarios that illustrate the effect of altering the characteristics of the CBR traffic.

Table 4: Various scenarios where coding techniques are evaluated.

Scenario #	1	2	3	4	5	6	7
Link Load (%)	110	120	110	110	120	110	120
Packet Size (B)	128	64	128	64	64	128	64
Bit-rate (kbps)	64	32	128	64	64	146	51
Inter-flow Spacing (ms)	16	16	8	8	8	7	10
Packet Separation (us)	102	51	102	51	51	102	51
# of CBR Flows	90	211	45	90	106	40	133
CBR Loss (%)	8.5	4.9	5.2	0.3	0.2	9.1	6.7
VoIP Loss (%)	9.9	32.4	13.5	18.9	38.4	10.6	30.0

The losses are considerable in all the cases even with low overloads of 110%. In the table, it is illustrated that through varying some of the parameters of the flows, the CBR flows can minimize their losses while maximizing losses on the voice traffic. The following are some observations from Table 4.

- At similar load levels, more voice packets are dropped when the bursts are close together (from scenarios 1 and 4) by varying the bit-rate of the individual CBR flows.
- A smaller CBR packet size causes more losses for voice than a larger one (from 3 and 4)

- When the inter-arrival time of the voice packets in each individual flow (16 ms) is a multiple of the burst interval, the losses for VoIP flows are more (from comparing 4 and 5 to 6 and 7) but the losses are also concentrated on the particular flows whose arrival times coincide with the arrival time of the burst. In cases similar to scenario 4, the distribution of losses among the flows was almost always in two categories: those which experience no losses and those which endure irrecoverable losses. In case of scenario 6, all the flows have equivalent loss levels and these are the cases where SFEC is most effective.
- By combining the effect of smaller packet sizes and a smaller burst interval, losses on the CBR traffic can be made negligible while voice traffic suffers more (from 4 and 5).

Table 5 shows the distribution of quality evaluations among the 82 calls traversing the link for different scenarios with SFEC1 encoding. This illustrates the point mentioned earlier about some scenarios having losses focused on a subset of flows or those which have losses spread out randomly among the flows.

Table 5: Distribution of quality for the 82 calls on the link.

Scenario	**1**	**3**	**4**	**6**
Burst spacing	**16**	**8**	**8**	**7**
Excellent (%)	85.4	82.9	79.3	73.2
Very Good (%)	0.0	0.0	3.7	24.4
Acceptable (%)	0.0	0.0	0.0	2.4
Many Dissatisfied (%)	1.2	0.0	0.0	0.0
Most Dissatisfied (%)	0.0	0.0	0.0	0.0
Not recommended (%)	13.4	17.1	17.1	0.0

5.3.2 Comparison of the Recovery Level of SFEC and LDPC

Some of the scenarios exhibit similar behavior based on the selection of some parameters as explained in the section 5.3.1. In this section, a representative set of cases will be selected in order to examine the performance of both SFEC and LDPC in these setups. An overload of 10% will be considered here as an example and in order to not overstate the benefits of LDPC.

In the first situation, losses are spread out evenly among the flows as in scenario 6 of Table 4. In Table 6, the resultant losses using various coding levels are presented for scenario 6. For LDPC, redundancy levels in the range of 0.8 to 0.85 are considered, and these are better rates than three of the considered SFEC coding methods. It is also quite close to the rate of SFEC3, which has the lowest level of redundancy. Using more redundancy in SFEC does not improve the situation at all because the losses were focused on a subset of the flows. So all the SFEC options give similar quality levels. When a single flow experiences a large number of losses, it is less probable that either the original coded-frame or the low-quality-coded redundant frame is received. For SFEC, it can be seen that in all cases, approximately 15% of the voice calls will exhibit poor quality. In addition, the table shows that achieving an acceptable level of quality (same as the PSTN quality [20]) for *all* flows is possible with a comparable or better redundancy level than all the SFEC methods. Furthermore, reducing the rate slightly from 0.85 to 0.833 (i.e. adding 1.7% more redundancy) will give all the voice flows toll quality.

Table 6: Scenario 1: losses focused on subset of flows (loss 9.9%).

Coding	Rate	Loss	Average R	% Toll	% Acceptable
SFEC1	0.710	0.0635951	82.5	85.4	85.4
SFEC2	0.450	0.0337716	86.3	87.8	87.8
SFEC3	0.870	0.0635951	81.8	85.4	85.4
SFEC4	0.690	0.0337716	85.3	85.4	85.4
LDPC	0.850	0.0721328	76.2	0.0	100.0
LDPC	0.833	0.0431331	82.3	100.0	100.0
LDPC	0.800	0	94.3	100.0	100.0

In this second scenario, the losses suffered are higher (19%) although they are still directed at a subset of the flows. Here, SFEC did not achieve *any* improvement by increasing the redundancy; the loss level has not changed much and neither has the call quality, as shown in Table 7. Moreover, approximately 17% of the calls have poor quality. When considering LDPC of rate of 0.727, which has a better rate than three of the four evaluated SFEC methods, can result in 100% recovery from errors, with all the calls achieving perfect (like-original) quality. If the higher rate code of 0.769 is observed, it can still maintain all calls at an acceptable level while having 10% more redundancy than SFEC3 and lower redundancy levels than the three other SFEC options.

Table 7: Scenario 4: losses focused on subset of flows (loss 18.9%).

Coding	Rate	Loss	Average R	% Toll	% Acceptable
SFEC1	0.710	0.170736	77.8	82.9	82.9
SFEC2	0.450	0.170736	77.8	82.9	82.9
SFEC3	0.870	0.170736	77.4	79.3	82.9
SFEC4	0.690	0.170736	77.4	79.3	82.9
LDPC	0.769	0.059484	78.9	0.0	100.0
LDPC	0.727	0	94.3	100.0	100.0

In the final situation of Table 8, the losses are spread out among the VoIP flows and the table shows that there is some benefit from using three redundant copies of SFEC rather than one. Losses are recovered to an acceptable level, all the resulting flows have acceptable quality, and almost all these are of toll quality. However, the least amount of redundancy to achieve this is at a rate of 0.87. In comparison, LDPC of a similar rate can also result in acceptable quality but not of toll quality. Adding 2% redundancy (i.e. at rate 0.85) brings all the calls to toll quality. In this case, the performance of LDPC is comparable to the lowest redundancy SFEC code (SFEC3) in its rate and recovered quality, whereas it surpasses all forms of SFEC in the required redundancy to achieve same or better quality.

Table 8: Scenario 6: losses spread out among the flows (loss 10.6%).

Coding	Rate	Loss	Average R	% Toll	% Acceptable
SFEC1	0.710	0.008768	89.3	97.6	100.0
SFEC2	0.450	0.000024	91.7	100.0	100.0
SFEC3	0.870	0.008768	87.3	97.6	100.0
SFEC4	0.690	0.000024	89.6	100.0	100.0
LDPC	0.875	0.069391	76.8	0.0	100.0
LDPC	0.850	0.039055	83.2	100.0	100.0
LDPC	0.833	0.028601	85.4	100.0	100.0
LDPC	0.800	0	94.3	100.0	100.0

5.3.3 Effect of Changing Parameters

In this section, we consider how the variation in some parameters for the CBR flows within a certain tolerance can still cause the damaging effect on voice traffic that was shown in the previous sections. Since neither the bit-rate nor the spacing between packets or flows can be deterministically guaranteed, the simulation results in this section show that a small variation in the parameters can still allow an attack to cause damage to VoIP quality in a real network.

In Figure 12, we show the effect of having a drift in the bit-rate of the flows between the CBR and VoIP flows. This is tested in the context of cases where the bit rate is held at 64kbps or multiples of it. In scenario 1, the bit rate is increased or decreased from 64kbps in the CBR flows. As a result, the flows are no longer spaced at exact intervals of 16 ms. The figure relates the percentage difference of the packet inter-arrival time (with respect to 16ms for scenario 1) to the loss rate experienced in the network. There is no obvious pattern in the measured average loss rate. However, the losses are no longer exclusive to a set of flows. The losses are almost the same for all the VoIP flows. As the percentage of variation is increased, the SFEC method becomes more effective since it can now recover more of the lost packets. This is illustrated in Figure 13. As a result of a change in the bit-rate, a considerable loss rate is still experienced in the network, but the losses are distributed more evenly among the flows and damage to voice traffic declines.

Although recovery with SFEC will improve as the rate deviates from its base value, LDPC is still as effective in such a situation. Table 9 shows the quality of the resultant streams for one-copy SFEC and LDPC. The average rating and the redundancy are comparable. As before, LDPC provides 100% toll quality calls, while for SFEC, some of the calls have slightly lower quality. SFEC1 approaches the recovery rate of LDPC, but with almost 15% additional redundancy.

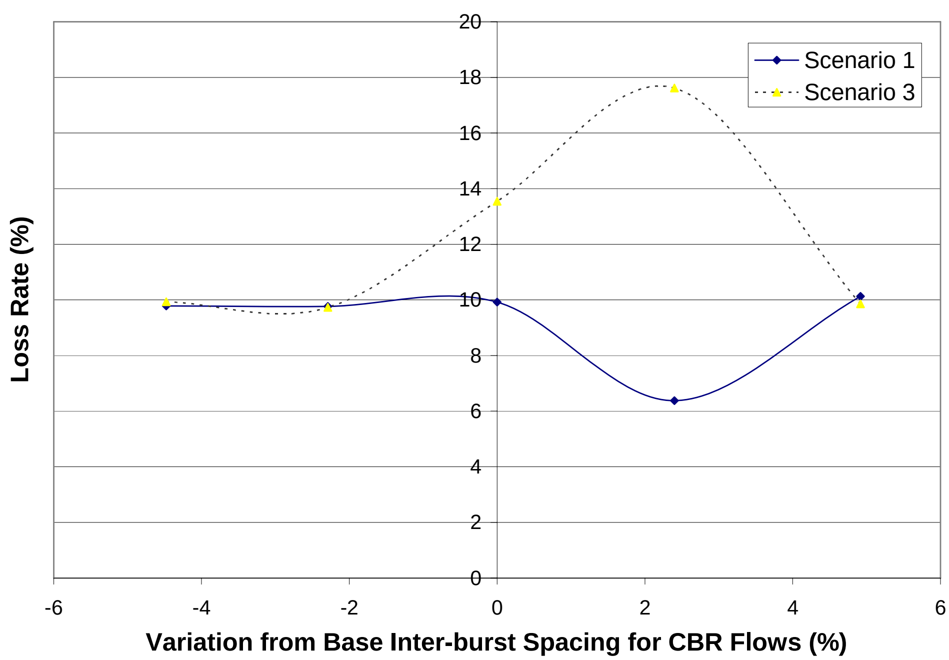

Figure 12: The effect of a drift in the sending rate of the flows on loss rates.

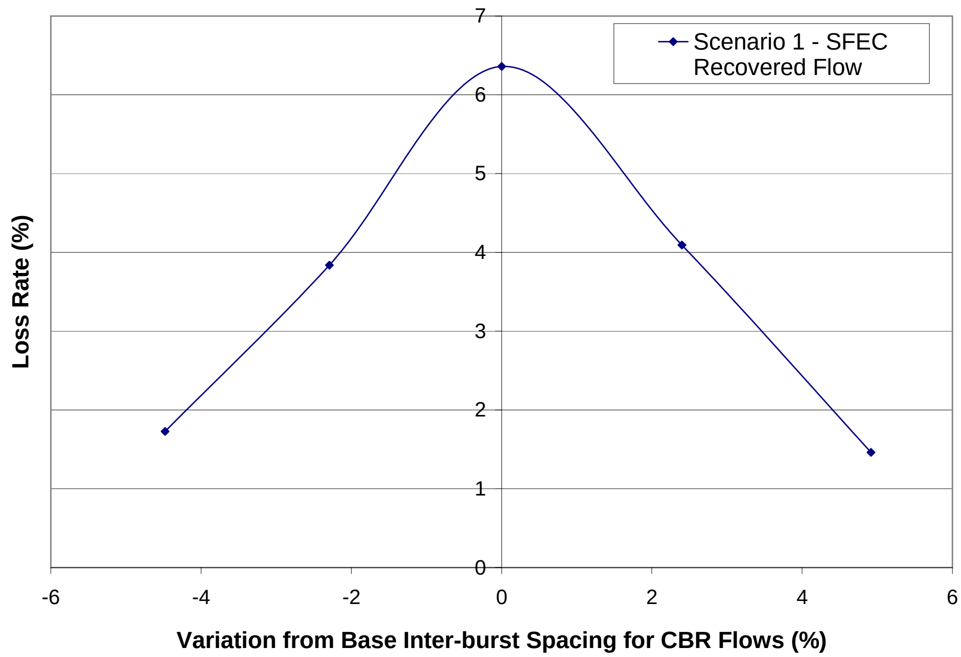

Figure 13: The effect of drift on recovered SFEC loss rate.

Table 9: Performance of SFEC vs. LDPC with -5% variation (loss = 9.8%).

Coding	Rate	Loss	Average R	% Toll	% Acceptable
SFEC1	0.710	0.01727	86.3	98.8	100.0
SFEC3	0.870	0.01727	84.6	92.7	100.0
LDPC	0.850	0.01400	88.5	100.0	100.0

Next the effect of changing the inter-flow spacing for the CBR flows is considered. The spacing between flows of the CBR traffic was originally set to the transmission time of a single packet in the CBR flow at the speed of the link. Figure 14 displays the effect of changing this value by ±20% of its original value. This is measured for two cases: one where the bit-rate of CBR is a multiple of the VoIP bit-rate and one in which it is not a multiple.

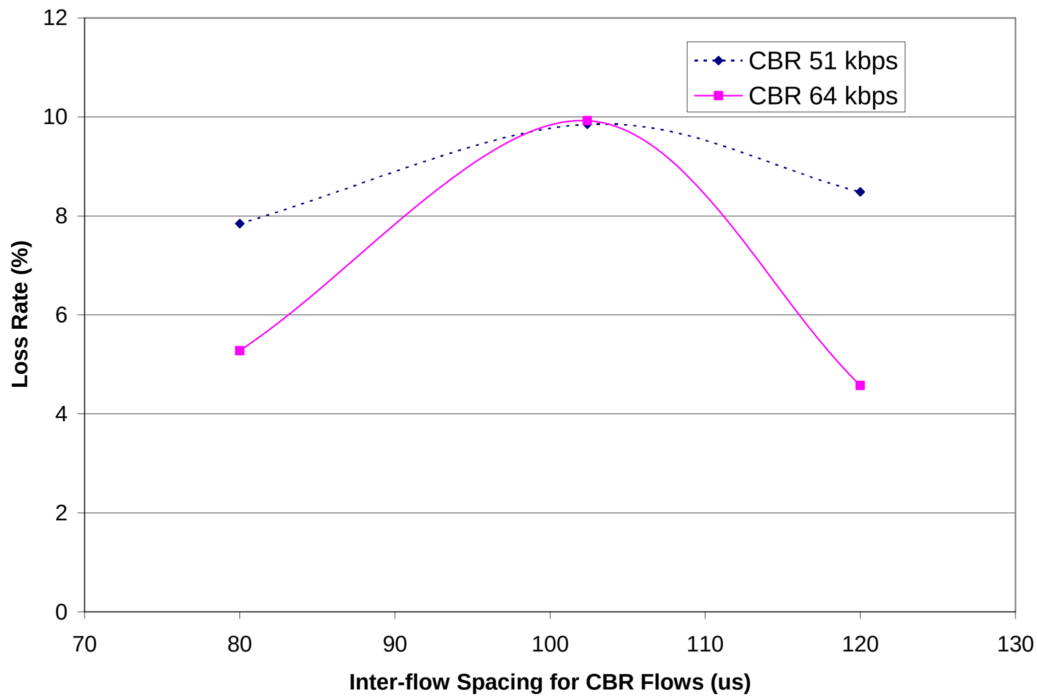

Figure 14: Effect of varying inter-flow spacing for two CBR setups.

It can be seen that any change in this value decreases the loss rate seen by the VoIP traffic. The loss rate is affected more when the CBR traffic is 64Kbps and the same effect can be seen on the recovered rate with SFEC. As seen before, the variation causes the losses to be spread out among more calls and improves the recovery rate of the audio stream from errors.

And finally, we study the effect of jitter in the inter-packet spacing for both the CBR and the VoIP flows. Jitter (or delay variation) is defined as the difference in the inter-arrival time of two packets [8]. This can be a result of the variable delays which the packets experience in the network because of the unpredictable queuing times. The simulation here changes more than one parameter of the setup, but the aim here is to introduce some randomness. Figure 15 illustrates the results of introducing jitter values ranging from 0 to 10% in scenario 3 which exhibited the highest loss rate for VoIP traffic (14%) at 110% loading of the link. The loss rate was reduced to 10% once jitter was introduced and it remains steady for higher jitter values. The ensuing loss rate from

SFEC encoding also dropped after the introduction of jitter and then it gradually drops off as jitter increases.

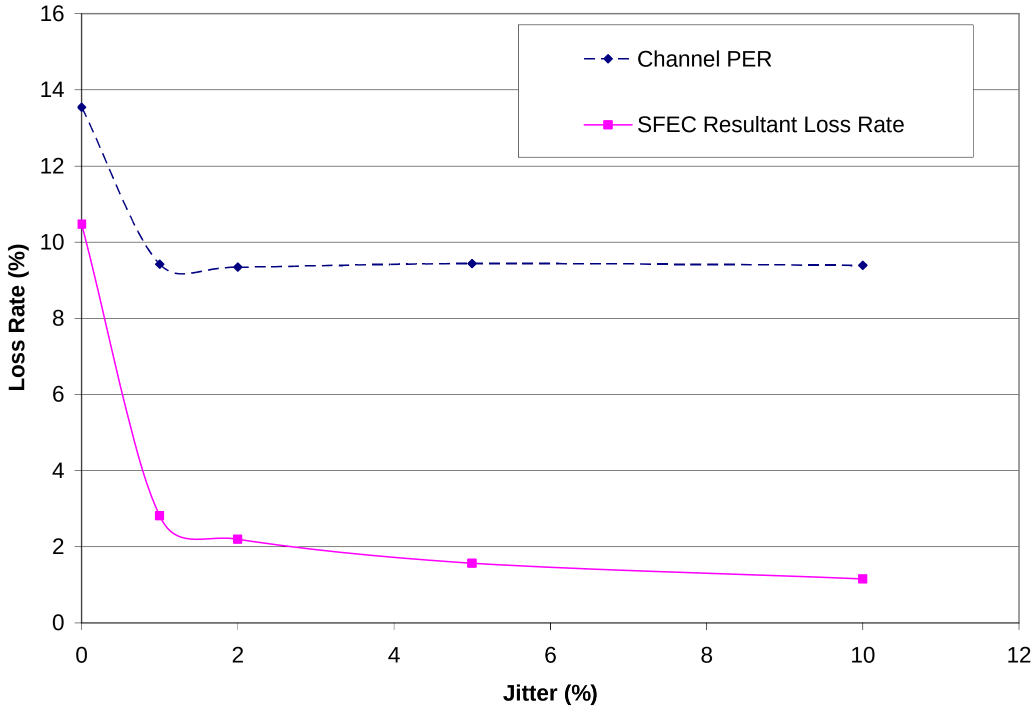

Figure 15: Effect of jitter in packet inter-arrival times on loss rate (scenario 3).

Table 10: Effect of a jitter of 5% on loss recovery in scenario 3.

Coding	Rate	Loss	Average R	% Toll	% Acceptable
SFEC1	0.710	0.01567	86.2	100.0	100.0
SFEC3	0.870	0.01567	84.6	100.0	100.0
LDPC	0.850	0.00263	93.0	100.0	100.0
LDPC	0.875	0.02979	85.1	100.0	100.0

Table 10 shows the relation between the recovery of SFEC and LDPC when ±5% jitter is introduced in the flows. In this case, they all achieve toll quality for the flows although a 10% loss was seen by the flows before recovery. Once more, the lower loss rate is a result of the randomness which causes the losses to be uniformly distributed and increases the effectiveness of the error correcting schemes.

5.4 NETWORK SIMULATIONS WITH MULTIPLE LINKS

All the previous experiments were performed over a single link. Voice calls may traverse many links in a VoIP network and attacks may occur on one or more of these links. To assess the damaging effect of attacks, this section shows the results for a three-link overloaded connection over which the voice traffic crosses.

In all the cases, the three links were overloaded to 110%. Two cases of attack traffic are simulated: one where the attacks are on all three links and the other where attack traffic only exists on the last link. The attack traffic is in bursts as described in previous sections, and all other traffic is simulated as a regular UDP flow. Table 11 compares these two cases with each other and with the single link results from the previous sections. These results are all for attack traffic with burst spacing of 16 msec (i.e. similar to that of the voice traffic). In both cases 2 and 3, the losses are almost doubled from that of case 1. Because all three links are slightly overloaded, the voice flows are losing packets on all three links and especially as a result of the attack traffic. In case 3, even though the attack traffic is only on the last link the percentage of calls with poor quality is very similar (and slightly more) than that of case 2 where all the links carry attack traffic. This result is displayed in table 12. Since the damaging effect for cases 2 and 3 are almost the same and it would be easier to direct attack traffic to one link rather than multiple links. Only network simulations of the form of case 3 will be considered in the following.

In section 5.3, the most damage is incurred by attack traffic with a burst period matching that of the voice traffic. However, when simulations are performed over three links, the loss rate for 7 msec spacing (19.6%) is somewhat larger than that for 16 msec periods (18.4%). The effect on quality of individual calls is also similar in both cases. Table 13 shows the distribution of call quality in case of GSM coding for SFEC. Tables

14 and 15 compare the resulting call quality for both SFEC and LDPC with different rates for 16 msec and 7 msec burst spacing, respectively. In both cases, almost 30% or more of the calls coded with SFEC have poor quality while acceptable quality can be achieved for all calls when LDPC is used with a rate of 0.75.

Table 11: Resulting quality with 16 msec-spaced attack traffic.

		Actual Loss	% of calls with toll quality	
			LPC	GSM
Case 1	Single link with attack traffic	9.9	85.4	85.4
Case 2	Three links - attack traffic on all links	18.7	72.0	72.0
Case 3	Three links - attack traffic on last link only	18.4	63.4	63.4

Table 12: Quality distributions for the three cases of table 11.

	Case 1	Case 2	Case 3
Excellent (%)	85.4	68.3	42.7
Very Good (%)	0.0	3.7	20.7
Acceptable (%)	0.0	0.0	6.1
Many Dissatisfied (%)	1.2	0.0	3.7
Most Dissatisfied (%)	0.0	0.0	0.0
Not recommended (%)	13.4	28.0	26.8

Table 13: Comparison of the different burst spacing (equal to voice packet spacing or not).

	16 msec	7 msec
Excellent (%)	42.7	45.2
Very Good (%)	20.7	28.0
Acceptable (%)	7.3	0.0
Many Dissatisfied (%)	2.4	0.0
Most Dissatisfied (%)	0.0	0.0
Not recommended (%)	26.8	26.8

Table 14: Recovery of LDPC and SFEC for the 16 msec case.

Coding	Rate	Loss	Average R	% Toll	% Acceptable
SFEC1	0.710	0.07855	76.5	63.4	70.7
SFEC3	0.870	0.07855	74.1	63.4	69.5
LDPC	0.750	0.03447	84.1	100.0	100.0
LDPC	0.769	0.09216	72.0	0.0	100.0

Table 15: Recovery of LDPC and SFEC for the 7 msec case.

Coding	Rate	Loss	Average R	% Toll	% Acceptable
SFEC1	0.710	0.07037	77.6	73.2	73.2
SFEC3	0.870	0.07037	74.8	68.3	73.2
LDPC	0.750	0.01695	87.8	100.0	100.0
LDPC	0.769	0.08215	74.1	0.0	100.0

In section 5.3.3, the addition of jitter to the single-link simulations is shown to increase the effectiveness of error-correcting techniques and as a result better call quality is achieved at the same loss rates. Here in this section, different levels of jitter will be added to the three-link simulation as described before to see how the damaging effect of attacks changes with jitter.

In a single link experiment, any introduction of jitter would reduce the damage of the attack dramatically. However in these experiments, the loss rate remained fairly constant and the damaging effects are not diminished till the level of jitter reaches 5%. Table 16 and 17 illustrate the resultant loss and quality for different encoding schemes at 16 msec and 7 msec burst-period, respectively. The loss rate on the channel is fluctuating as jitter is increased and so is the resultant quality. LDPC code with a rate of 0.75 achieves acceptable quality for all calls at different jitter levels. With comparable rates to LDPC, SFEC encoding has 50% of the calls experiencing un-acceptable quality with 0.5% of jitter. The recovery of SFEC improves as jitter is increased and it too achieves 100% of the calls with acceptable quality. The higher LDPC rate of 0.769 can still achieve acceptable quality for all the calls. The few cases where no calls achieve acceptable quality, the actual average R-rating was around 69 which is just below the threshold for acceptable quality (i.e. at 70). But because the losses are high and they are spread out among the flows, the average rating is on the verge of the acceptable range.

Table 16: The effect of adding jitter to the 16 msec burst-spacing configuration.

	Jitter (%)	0	0.5	1	3	5
	Channel loss	**0.184**	**0.196**	**0.166**	**0.196**	**0.189**
SFEC1 (0.71)	**Loss**	0.079	0.088	0.058	0.061	0.050
	Acceptable (%)	70.7	47.6	85.4	98.8	100.0
SFEC3 (0.87)	**Loss**	0.079	0.088	0.058	0.061	0.050
	Acceptable (%)	69.5	35.4	73.2	81.7	100.0
LDPC (0.75)	**Loss**	0.034	0.046	0.000	0.063	0.006
	Acceptable (%)	100.0	100.0	100.0	100.0	100.0
LDPC (0.769)	**Loss**	0.092	0.108	0.001	0.107	0.067
	Acceptable (%)	100.0	0.0	100.0	0.0	100.0

Table 17: The effect of adding jitter to the 7 msec burst-spacing configuration.

	Jitter (%)	0	1	3	5
	Channel loss	**0.196**	**0.189**	**0.188**	**0.140**
SFEC1 (0.71)	**Loss**	0.070	0.074	0.057	0.028
	Acceptable (%)	73.2	63.4	100.0	100.0
SFEC3 (0.87)	**Loss**	0.070	0.074	0.057	0.028
	Acceptable (%)	73.2	52.4	91.5	100.0
LDPC (0.75)	**Loss**	0.017	0.006	0.086	0.000
	Acceptable (%)	100.0	100.0	100.0	100.0
LDPC (0.769)	**Loss**	0.082	0.062	0.108	0.000
	Acceptable (%)	100.0	100.0	0.0	100.0

And finally, we study how the characteristics of the loss bursts change from a single voice flow compared to the aggregate flow of 82 voice calls. Loss run refers to the number of consecutive packet losses. It is measured here for different network simulations (by changing the burst-spacing for attack traffic and the jitter levels) for both a single voice call (as seen by SFEC) and for the 82 calls combined (as seen by LDPC). Table 18 shows that the average and maximum loss run for LDPC is always less than that seen by the SFEC method. Because individual calls in SFEC may be targeted separately from the rest of the voice flows by an attack coinciding with it (as seen from the previous results), the maximum loss run values are very high in these cases. From the measurements for LDPC, it can be seen that these long loss runs are eliminated. SFEC does not tolerate very high burst lengths because the original and redundant samples cannot be separated by a large number of packets due to delay limitations. For example, for 128 B packets and 50 msec of allowable delay, no more than three packets of separation can be tolerated. LDPC, on the other hand, can handle as much packet losses as the particular code can tolerate. As long as enough packets from a block are received, recovery is still possible.

Table 18: Loss run measurement comparison for both SFEC and LDPC.

		Loss Run		
		Average	Maximum	Minimum
16 msec no jitter	LDPC	1.191	12	1
	SFEC	1.539	499	1
16 msec 1%	LDPC	1.353	9	1
	SFEC	1.644	15	1
16 msec 3%	LDPC	1.445	9	1
	SFEC	1.541	15	1
16 msec 5%	LDPC	1.421	7	1
	SFEC	1.424	12	1
7 msec no jitter	LDPC	1.249	8	1
	SFEC	1.352	192	1
7 msec 1%	LDPC	1.418	8	1
	SFEC	1.762	19	1
7 msec 3%	LDPC	1.418	7	1
	SFEC	1.529	18	1
7 msec 5%	LDPC	1.283	6	1
	SFEC	1.305	8	1

6. DISCUSSION

The proposed approach is dependent on the existence of appropriate high speed links or paths that carry a large number of voice calls. VoIP providers use several such paths in their overlay networks or eventually combine voice streams on the same backbone links. Businesses also may tunnel their voice traffic from an internal network over the Internet or using private circuits to other branch offices or partners. This would enable encoding and decoding to be done on such paths. Deployment of the proposed approach is easier for overlays and label-switched paths (LSPs), than a widespread approach over IP where aggregation can only be done hop-by-hop.

There is an inherent delay in an FEC-based scheme due to the time needed to aggregate the block to be coded and due to the processing delays at both the transmitter and the receiver. Although the SFEC method does not require much processing for the original stream at the receiver, the encoding of the redundant stream is still CPU-intensive [2]. And in [31], the measured end-to-end delay for an SFEC method was found to be larger than the delays incurred by the FEC coding method. LDPC codes have even lower processing delays than the block codes considered in [31]. Moreover, the minimum delay between packets for SFEC is restricted from decreasing due to the minimum frame size required by the low-quality codec and the overhead from packet headers. Through aggregation of packets from many audio streams in our approach, a reasonably large block size can still have a small delay for individual calls. With a higher bandwidth or a larger number of aggregated flows, the packet size and the sending rate remain flexible.

It is also possible to apply LDPC on flows which are already protected by SFEC. The effect would be to reduce the loss rate experienced over the link(s) where LDPC is applied. If the flow can tolerate the increased delay from encoding and decoding, applying LDPC on certain high loss or bursty links may be beneficial.

Using LDPC can also reduce the jitter compared to the SFEC. In LDPC there is additional delay required to accumulate all parts of the block before processing it. But the arrival times between blocks does not vary much, and any variance that exists is only because the unavoidable queuing delay on the network. Whereas in SFEC, inter-arrival times depend on whether the original frame or the redundant frame are used in the played out stream. For example, if the separation of the original from the redundant packets is 50 ms, the delay will increase by 50 ms every time a lost frame is recovered and the jitter for the flow is increased by 50 ms than that of an unprotected stream.

A block of code in LDPC consists of a large number of packets (128 and above), while SFEC can only have the redundancy spaced a few packets away (3 packets is the maximum when 128B packets are used) to stay within the delay requirements. This enables LDPC to be more tolerant to burst errors than SFEC. In addition, the results show that the average loss run experienced by LDPC is slightly less than that for SFEC.

The unrecoverable losses of the LDPC method will result in missing bits spread out across all the samples; but these samples belong to different voice calls from those that have been aggregated. So, the corrupted samples are not lost completely and by using similar error-concealment techniques as those employed for the reconstruction of the SFEC stream; these will become lower quality samples. It has been shown that the degradation in audio quality due to replacement of a lower quality sample is not very high [11].

Our approach requires encoding and decoding on network paths. The encoding and decoding of LDPC codes can be done in software up to 100 Mbps [19], but may require hardware at higher speeds.

7. CONCLUSION

We proposed to employ LDPC codes on aggregated VoIP traffic on providers' links to improve the resiliency of voice traffic against packet losses. Our experiments have shown that LDPC codes can provide significantly lower output loss rates with comparable overheads when compared to FEC schemes employed for individual calls. Our experiments also show that LDPC coding of aggregated traffic can employ much lower overhead at lower input loss rates while providing a similar protection of tolerating burst errors. The network simulations demonstrate that LDPC can maintain a higher quality for the voice calls than SFEC while requiring comparable or lower redundancy. Using many different network setups and various forms of opposing traffic, the LDPC coding is shown to be quite resilient to packet losses and provides an acceptable, if not excellent, level of quality at all times.

One direction for the future would be experimenting with the performance of LDPC in the presence of quality-of-service protection on the VoIP links; for example, when differentiated services or admission control are employed.

REFERENCES

[1] Y. Boger, “Fine-tuning voice over packet services,” RADCOM Ltd., Paramus, NJ, accessed March 2005, [Online]. Avaliable: http://www.radcom.com/radcom/WhitePapers/voip2.htm

[2] J.-C. Bolot and A.V. Garcia, “Control mechanisms for packet audio in the Internet,” in *Proc. IEEE INFOCOM*, San Francisco, CA, March 1996, vol. 1, pg. 232 – 239 [Online]. Avaliable: http://ieeexplore.ieee.org/iel2/3539/10628/00497898.pdf?tp=&arnumber=497898&isnumber=10628

[3] J.-C. Bolot and A.V. Garcia, “The case for FEC-based error control for packet audio in the Internet,” *ACM Multimedia Systems*, 1997, [Online]. Avaliable: ftp://ftp-sop.inria.fr/rodeo/avega/papers/96.Audio_FEC.ps.gz

[4] P. Calyam, M. Sridharan, W. Mandrawa, and P. Schopis, “Performance measurement and analysis of H.323 traffic,” presented at the PAM Workshop, Antibes Juan-les-Pins, France, April 2004, [Online]. Avaliable: http://www.osc.edu/ oarnet/itecohio.net/beacon/h323.pdf

[5] Cisco Systems, “Cisco security advisory: Multiple vulnerabilities in Cisco IP telephones,” Cisco Systems, San Jose, CA, May 2002, [Online]. Available: http://www.cisco.com/warp/public/707/multiple-ip-phone-vulnerabilities-pub.shtml

[6] Cisco Systems, “Vonage: Customer profile,” Cisco Systems, San Jose, CA, 2002, [Online]. Available: http://www.cisco.com/warp/public/cc/techno/tyvdve/sip/prodlit/vcpag_cp.pdf

[7] A. D. Clark, “Modeling the effects of burst packet loss and recency on subjective voice quality,” presented at the IP Telephony Workshop, New York, NY, April 2001, [Online]. Avaliable: http://www.telchemy.com/reference/iptel2001.pdf

[8] C. Demichelis and P. Chimento, "IP packet delay variation metric for IP performance metrics (IPPM)," IETF, Request for Comments RFC 3393, November 2002, [Online]. Avaliable: http://www.faqs.org/rfcs/rfc3393.html

[9] M. Friedman, “VoIP security chief warns of ‘life and death’ security threats,” Networking Pipeline, CMP Media LLC, Manhasset, NY, April 2005, [Online].

Avaliable: http://www.informationweek.com/story/showArticle.jhtml?articleID=160700565

[10] J. D. Gobson, *Multimedia Communication: Directions and Innovation*, Academic Press Series in Communications, Networking and Multimedia, pg. 41, San Diego: Elsevier, 2001, [Online]. Avaliable as e-book from: http://netlibrary.com

[11] V. Hardman, M. A. Sasse, M. Handley, and A. Watson, “Reliable audio for use over the Internet,” presented at the INET Conf., Honolulu, Hawaii, June 1995, [Online]. Avaliable: http://citeseer.ist.psu.edu/hardman95reliable.html

[12] C. Hoene, B. Rathke, and A. Wolisz, “On the importance of a VoIP packet”, presented at the ISCA Tutorial and Research Workshop on the Auditory Quality of Systems, Akademie Mont-Cenis, Germany, 2003, [Online]. Avaliable: http:// www.tkn.tu-berlin.de/publications/ papers/paper10. pdf

[13] C. Hoene and E. D.-Lalla, “Predicting performance of PESQ in case of single frame losses,” presented at the MESAQIN Conf., Prague, CZ, June 2004, [Online]. Avaliable: http://www.tkn.tu-berlin.de/publications/papers/hoene_ mesaqin_2004 .pdf

[14] C. Hoene, S. Wietholter, and A. Wolisz, “Predicting the perceptual service quality using a trace of VoIP packets,” presented at the QofIS Workshop, Barcelona, Spain, September 2004, [Online]. Avaliable: http://www.tkn.tu-berlin.de/tkn/ publications/papers/hoene 2004qofis.pdf

[15] Information Sciences Institute (ISI), “Network simulator” (NS-2), ISI, University of Southern California, Marina del Rey, CA, accessed January 2004, [Online], Available: http://www.isi.edu/nsnam/ns/

[16] A. Klein, “Security analysis: Traditional telephony and IP telephony,” SANS Institute, as part of GIAC practical repository, Bethesda, MD, March 2003, [Online]. Available: http://www.sans.org/rr/whitepapers/telephone/924.php

[17] T. J. Kostas, M.S. Borella, I. Sidhu, G. M. Schuster, J. Grabiec, and J. Mahler, 3COM, “Real-time voice over packet-switched networks,” *IEEE Network*, vol. 12, issue 1, pg 18-27, January – February 1998, [Online]. Avaliable: http://

ieeexplore.ieee.org/iel4/65/14452/00660003.pdf?tp=&arnumber=660003&isnumber=14452

[18] J. Liesenborgs, “Voice over IP in networked virtual environments,” School for Knowledge Technology, Hasselt, Belgium, Thesis, May 2000, [Online]. Available: http://research.edm.luc.ac.be/jori/thesis/thesis.html

[19] M. G. Luby, M. Mitzenmacher, M. A. Shokrollahi, D. Spielman, and V. Stemann “Practical loss-resilient codes,” in *Proc. of the ACM Symposium on Theory of Computing*, El Paso, TX, 1997, pg. 150 – 159, [Online]. Avaliable: http://citeseer.ist.psu.edu/luby98practical.html

[20] A. Markopoulou, F. Tobagi, and M. Karam, “Assessment of VoIP quality over Internet backbones,” in *Proc. of the IEEE INFOCOM*, New York, NY, June 2002, vol. 1, pg. 150 - 159, [Online]. Avaliable: http://citeseer.ist.psu.edu/luby98practical.html

[21] E. McPhillips, “Voice over IP equipment and service markets,” HP Laboratories, Bristol, UK, Tech. Rep. HPL-IRI-1999-005, 1999, [Online]. Available: http://www.hpl.hp.com/techreports/1999/HPL-IRI-1999-005.html

[22] M. Mitzenmacher, “Digital fountains: A survey and look forward,” in *Proc. of the IEEE Information Theory Workshop*, San Antonio, TX, October 2004, pg. 271 – 276, [Online]. Avaliable: http://ieeexplore.ieee.org/iel5/9641/30470/01405313.pdf?tp=&arnumber=1405313&isnumber=30470

[23] B. Munch, “IP telephony security: Deploying secure IP telephony in the enterprise network,” Juniper Networks, META Group white paper, Stamford, CT, January 2005, [Online]. Available: http://www.juniper.net/company/events/metavoip.pdf

[24] C. Padhye, K. Christensen, and W. Moreno, “A new adaptive FEC loss control algorithm for voice over IP applications,” in *Proc. of the IEEE IPCCC*, Phoenix, Arizona, February 2000, pg 307 – 313, [Online]. Avaliable: http://ieeexplore.ieee.org/iel5/6712/17967/00830332.pdf?tp=&arnumber=830332&isnumber=17967

[25] M. Podolsky, C. Romer, and S. McCanne, “Simulation of FEC-based error control for packet audio on the Internet,” in *Proc. of the IEEE INFOCOM*, San Francisco, CA, April 1998, vol. 2, pg. 505 – 515, [Online]. Avaliable: http://ieeexplore.ieee.org/iel4/5315/14598/00665068.pdf?tp=&arnumber=665068&isnumber=14598

[26] B. Reynolds and D. Ghosal, “Secure IP telephony using multi-layered protection”, presented at the NDSS Symposium, San Diego, CA, 2003, [Online]. Avaliable: http:// www.isoc.org/isoc/conf.s/ndss/03/Proc./papers/3.pdf

[27] A. Shokrollahi, “LDPC codes: An introduction,” Digital Fountain, Inc., Fremont, CA, April 2003, [Online]. Avaliable: http://algo.epfl.ch/pubs/abstract.php?db=pubs_ldpc.txt&v=ldpc_intro.pdf

[28] K. Steinklauber, “VoIP security in small businesses,” SANS Institute, as part of GIAC practical repository, Bethesda, MD, May 2003, [Online]. Available: http://www.giac.org/practicals/GSEC/Klaus_Steinklauber_GSEC.pdf

[29] E. Vergetis, R. Guerin, and S. Sarkar, “Benefiting from path diversity with bursty losses,” University of Pennsylvania, Philadelphia, PA, 2004, [Online]. Available: http://www.seas.upenn.edu/~vergetis/ EV_Diversity_04.pdf

[30] E. Weiss, “Security concerns with VoIP,” SANS Institute, as part of the Information Security Reading Room, Bethesda, MD, August 2001, [Online]. Available: http://www.sans.org/rr/whitepapers/voip/323.php

[31] Z. Zhou, P. K. McKinley, and S. M. Sadjadi, “On quality-of-service and energy consumption tradeoffs in FEC-encoded audio streaming” in *Proc. of the IEEE IWQoS Workshop*, Montreal, Canada, June 2004, pg. 161 - 170 [Online]. Avaliable: http://ieeexplore.ieee.org/iel5/9152/29055/01309376.pdf?tp=&arnumber=1309376&isnumber=29055

VITA

Camelia Al-Najjar received her B.S. degree in electrical engineering in 2002 from the University of Jordan, Amman, Jordan. After being granted a Fulbright scholarship, she attended Texas A&M University - College Station and obtained her M.S. degree in computer engineering in 2005. Her permanent mailing address is P.O. Box 711447, Amman 11171, Jordan.