A Service Provider's Approach for Improving Performance of Aggregate Voice-over-IP Traffic

Camelia Al-Najjar, *Student Member, IEEE* and A. L. Narasimha Reddy, *Senior Member, IEEE*

***Abstract*—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. 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 network link or path, the provider can employ a suitable linear-time encoding for the aggregated voice traffic, resulting in considerable quality improvement with little redundancy. We show that it is possible to achieve rates closer to link 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 exceed similar or other techniques applied to individual voice calls.**

***Index Terms*— DoS attacks, overlays, packet drops, performance, VoIP service.**

I. Introduction

VOICE-OVER-IP (VoIP) services have witnessed increased popularity and acceptance in recent years, but this trend comes along with concerns about voice quality over these networks. Achieving acceptable real-time interactivity of voice calls depends on loss, jitter, and delay along with other factors. The unpredictable QoS of the Internet coupled with threats from denial-of-service attacks make achieving toll quality comparable to that of PSTN systems an interesting and challenging problem.

Internet backbones are still used for the long haul for long-distance VoIP calls while some VoIP providers started using private links [1]. Businesses are incorporating voice into the existing network infrastructure to eliminate the need for switching equipment [2]. 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 [3]. Smaller providers can compete in the international market through the use of the Internet. Even major VoIP providers use Internet paths as backup to their primary lines [4]. In addition, the development of global commerce has lead many companies to enable customers to call through the Internet [2].

Security threats to VoIP should be expected even when private links are employed; it has been shown that a majority of hacking incidents come from the inside [2]. VoIP is also prone to denial-of service (DoS) attacks, worms, viruses, and Trojans [5], but here the stakes are higher; an attack can bring down the phone network, consume large bandwidth to block many calls, or even prevent emergency calls to police or fire departments [2]. It is also possible to reconfigure VoIP devices remotely and plant malicious software [4].

Furthermore, links with large and variable delay properties can cause poor performance for audio even if the average loss rate is low ([1], [6]). Forward-error-correction (FEC) protection is typically applied on individual calls on an end-to-end basis. In [7], it is shown that a blind application of signal-processing-based FEC (SFEC) (where a lower quality of speech is sent in a later packet as a backup) does not always give a beneficial result. The considerable amount of redundancy of FEC for individual calls cannot be ignored when the network carries a large number of FEC-protected voice streams [7]. Low-density-parity-check (LDPC) codes are an FEC technique which can scale up to large block sizes with reasonable (and symmetric) encoding/decoding times allowing for online usage. For clarity, in this paper, these SFEC and LDPC forms of FEC will be referred to as SFEC and LDPC codes, respectively; whereas all other FEC-based coding will be called FEC.

An alternative for error-correction is using codecs that are designed to have good robustness against packet losses, such as the iLBC [8]. In [8], a comparison shows that applying FEC to a regular codec (such as G.729) gives better speech quality than iLBC for almost all levels of loss with similar bandwidth usage. LDPC codes, used in this paper, can achieve better error recovery than Reed Solomon codes used in [8].

In this paper, we propose an approach that can be employed by the service providers to protect VoIP traffic from DoS attacks, transient traffic overloads, and the resulting packet losses. The next section presents the proposed approach and an overview of the results.

Manuscript received February 25, 2006. This work was supported in part by NSF grant ANI-0087372 and by Intel corp.

A. L. N. Reddy is with the Electrical and Computer Engineering Department, Texas A&M University, College Station, TX 77840 USA. (phone: 979-845-7598, fax: 979-845-2630, e-mail: reddy@ece.tamu.edu).

C. Al-Najjar was an MS student at Texas A&M University, College Station, TX 77840 USA. She is now working on a PhD degree from Case Western Reserve University, Cleveland, OH 44106 USA. (e-mail: caa20@case.edu).

II. Aggregate Coding of VoIP Traffic

This paper proposes a technique to improve service quality over paths that VoIP providers may not completely control or where there are threats of DoS attacks. 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. Fig. 1 illustrates the proposed scheme for a VoIP provider's network. The VoIP service provider's access routers or gateways or overlay servers can employ protection mechanisms on aggregated VoIP traffic with a common egress point of the provider's network to mitigate DoS attacks and excessive packet loss rates.

The benefits of low overhead can 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. Fig. 2 shows how frames from individual VoIP calls can be 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.

The lost packets will either be fully or partially recovered depending on both the input loss rate and the redundancy of the code. 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 in a block come from multiple streams, a burst of errors will affect individual streams by corrupting the few samples which are contained in these packets.

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 in an individual call can degrade the voice quality [1]. But aggregation with LDPC will spread out the losses and the quality degradation will not be localized on any of the calls.

Our results show that a single link overloaded by 10% can cause losses up to 20% for the voice traffic. The losses in most of these cases cause severe damage to a subset of the calls. When using LDPC with a similar amount of redundancy, we will show that all calls can achieve excellent quality. 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, we will show that the recovery rate and the resulting quality for LDPC are much better than that for SFEC with a comparable level of added redundancy when links are overloaded or under attack.

III. Packet Audio: Coding, Call Quality, and Error-Correction

In this paper, we will employ the E-model, an objective quality model, to evaluate transmission quality. The quality measure, which is called the R-factor [1], is evaluated with the

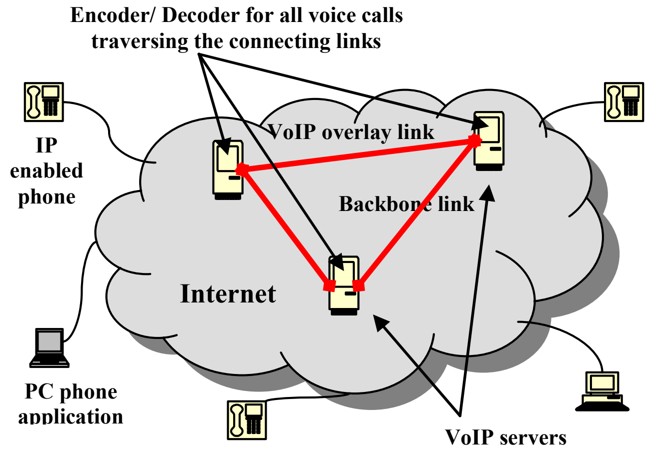

Fig. 1. Configuration of a VoIP network.

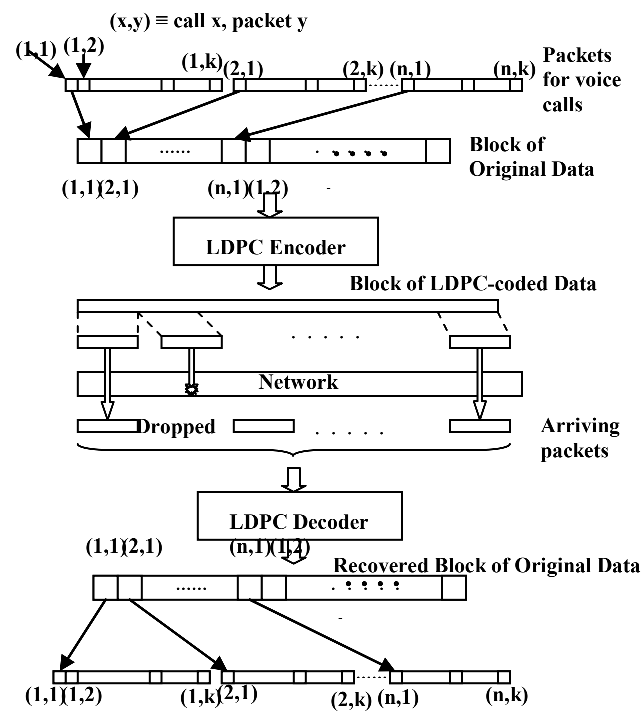

Fig. 2. Process of aggregation, encoding, transmission, and decoding for VoIP calls

equation (1) below. The R-factor is designed to resemble subjective quality measures.

$$R = R_0 - I_s - I_d - I_e + A \quad (1)$$

The terms in the equation cover the various impairments: R_0 is the effect of noise and loudness, I_s is the impairment which is concurrent with the speech signal (e.g. quantization), I_d is the impairment from one-way delay, talker/listener echo, and interactivity, I_e is the impairment from signal distortion (e.g.

low rate codecs, lost frames), and A is the advantage; when users sacrifice quality for convenience (e.g. cellular systems). The effect is additive when it is converted to the R psychoacoustic scale [1]. Toll quality corresponds to an R value of 80 or higher, whereas R values above 70 are considered of acceptable quality [1]. The factors (Ro, Is, and A) are network-independent and we will not be concerned with them here.

Table I summarizes the characteristics for several codecs used in packet-audio ([1], [9]). GSM and LPC both have a total encoding delay of 20 ms, whereas delays in PCM and ADPCM are negligible in comparison [9]. Table I also shows the intrinsic R value, for each codec, without consideration for losses (i.e. $R_o - I_s$). For a PCM stream, ($R_o - I_s$) has a value of 94.3 ([10], [1]) and R can then be calculated as ($R = 94.3 - I_d - I_e$). Delays up to 175 ms have little effect on R and I_d is considered negligible in this range [10]. The loss impairment values (I_e) of PCM are given in [1]. In [1], the R-factor approximately declines by 4 for a loss of 1% and drops by 2 for every 1% after that. We employ these results later in computing the quality of the calls in simulations.

In this paper, we compare LDPC based aggregate coding of audio traffic with SFEC based protection of individual calls. In SFEC, a redundant low-quality version of a packet (e.g., LPC or GSM) is piggy-backed onto a subsequent packet. At the receiver, the low quality frame replaces the original frame when it is lost ([6], [11]). SFEC also uses error concealment techniques when neither the original nor the redundant packet is recovered [7]. It is possible to vary the number of redundant copies or the delay of the redundant frame from the original one to produce a number of SFEC methods with different redundancy rates and recovery levels [11].

In LDPC, a number of calls are aggregated by the provider and encoded for protection as shown in Fig. 2. LDPC encoding/decoding times for large blocks are not restrictive for real-time audio, compared to other forms of FEC such as Reed–Solomon codes. When packets are lost in the network due to congestion or channel losses, the decoder can recover missing bits (i.e. erasures) using the code. When the erasures are excessive, they may not all be recovered.

We define the input loss rate as the packet error rate incurred by the network, and 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)%*.

TABLE I
PROPERTIES OF DIFFERENT CODING STANDARDS

Codec	*PCM*	*ADPCM*	*GSM*	*LPC*
Rate (kbps)	64	16 - 40	13	4.8
Intrinsic R	94.3	39 - 89	70	50

IV. RESULTS AND EVALUATION

The following sections show the results of evaluating aggregation of VoIP calls with LDPC. First, we study the level of recovery that these codes can achieve at various block sizes and with different number of packets per encoded block. Following that, a set of network simulations are used to compare the resulting quality of both the LDPC and SFEC encoded streams. Due to the limitation of space, some results are summarized briefly without accompanying experimental data. For more detailed results, the reader is referred to [12].

A. 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 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) [13]. 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. The output loss rate is averaged over a large number of simulated block coding and transmissions.

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 [14]. Our results in fact showed that a linear relationship exists between the loss rate in the channel and the level of redundancy needed for recovery from errors.

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). With the assumption of a voice stream of 64kbps, a 32KB block combines 256 calls. On the other hand, if a lower bit-rate codec such as G.729 (8 kbps) is used, an 8KB block size and 32B packets could be used for the same delay and the same number of aggregated calls. Although the results in the following section are for the (32KB, 128B) combination, they also apply to any block configured with 256 packets per block since losses are seen for entire packets. We also note that from the simulations, the 128 packets per block encoding does not show much degradation from the case with 256 packets per block, enabling aggregation of smaller number of calls. When fewer calls are available for aggregation, multiple packets from each call can be included to form a larger block with accompanying higher delay.

B. 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 NS-2 simulator [15], several such scenarios are studied to evaluate both the techniques. The simulation setup consists of a single link with a speed of 10 Mbps and a delay of 10 ms, and this link carries 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. The VoIP flows 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 whereas 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 two cases: (a) a low-quality frame is attached to the third subsequent packet and (b) 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 a 32 kbps ADPCM-encoded audio stream with a redundant low-quality coding of either GSM or LPC. Table II shows the rate of the resulting stream and the intrinsic quality (R) for all the cases considered. The recovered stream will be mostly ADPCM coded and the remaining part will be of low-quality code. Based on the additive property of the R scale, we estimate the quality of this stream from the proportions of these two parts. 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 using the loss impairment values of PCM in [1].

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.

1) *The Opposing Traffic*

Some initial experiments showed that when the opposing flow has a large packet size (1000B) with respect to the voice packets, the voice traffic sees negligible losses even at overloads of 40%. However, when a single CBR flow with the same packet size as voice (128B) is used, comparable losses are seen for both the voice flows and the opposing flow. Therefore, we will be using a 128B packet size in the following experiments.

TABLE II
CHARACTERISTICS OF THE EVALUATED SFEC STREAMS

	Redundant Stream				
SFEC Encoding	***Copies***	***Codec***	***Bit-rate (kbps)***	***R***	**Rate**
SFEC1-GSM	1	GSM	13	70	0.71
SFEC3-GSM	3	GSM	13	70	0.45
SFEC1-LPC	1	LPC	4.8	50	0.87
SFEC3-LPC	3	LPC	4.8	50	0.69

TABLE III
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
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

Here, we will consider 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. Each CBR flow sends packets of specified size at regular rate. The packets of different flows are spaced apart by a time equal to the ratio (packet size/link capacity). Table III 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.

The losses are considerable in all the cases even with low overloads of 10%. In the table, it is illustrated that through varying some of the parameters of the flows, the CBR flows can maximize losses on the voice traffic. The following are some observations from Table III. (a) At similar load levels, more voice packets are dropped when the bursts are closer together (from scenarios 1 and 4) by reducing packet size without changing the bit-rate of the individual CBR flows. (b) A smaller CBR packet size at the same burst separation causes more losses for voice than a larger one (from 3 and 4). (c) 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. (d) By choosing appropriate CBR characteristics, voice traffic can be made to experience the majority of the losses on the channel (from 4 and 5). Table IV shows the distribution of quality evaluations among the 82 calls traversing the link for different scenarios with SFEC1-GSM 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.

2) *Comparison of the Recovery Level of SFEC and LDPC*

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 focused on a subset of the flows; as in scenario 1 of Table III. In Table V, the resultant losses using various coding levels are presented for scenario 1. For LDPC, rates in the range of 0.8 to 0.85 are considered; resulting in lower redundancy levels than three of the SFEC coding methods we are evaluating. 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, approximately 15% of the voice calls exhibit poor quality. The results show that achieving an acceptable level of quality for all flows is possible with LDPC at a comparable or better redundancy level than all the SFEC methods. 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.

In scenario 4 of Table III, 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 VI. Moreover, approximately 17% of the calls have poor quality. With LDPC at a 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 SFEC1-LPC and lower redundancy levels than the three other SFEC options.

In the final situation of Table VII (scenario 6 in Table III), the losses are spread out among the VoIP flows and the results show 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 (SFEC1-LPC) in its rate and recovered quality, whereas it surpasses all forms of SFEC in the required redundancy to achieve same or better quality.

TABLE IV
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

TABLE V
SCENARIO 1: LOSSES FOCUSED ON SUBSET OF FLOWS (LOSS 9.9%)

Coding	**Rate**	**Loss**	**Average R**	**% Toll**	**% Acceptable**
SFEC1-GSM	0.710	0.0636	82.5	85.4	85.4
SFEC3-GSM	0.450	0.0338	86.3	87.8	87.8
SFEC1-LPC	0.870	0.0636	81.8	85.4	85.4
SFEC3-LPC	0.690	0.0338	85.3	85.4	85.4
LDPC	0.850	0.0721	76.2	0.0	100.0
LDPC	0.833	0.0431	82.3	100.0	100.0
LDPC	0.800	0	94.3	100.0	100.0

TABLE VI
SCENARIO 4: LOSSES FOCUSED ON SUBSET OF FLOWS (LOSS 18.9%)

Coding	**Rate**	**Loss**	**Average R**	**% Toll**	**% Acceptable**
SFEC1-GSM	0.710	0.1708	77.8	82.9	82.9
SFEC3-GSM	0.450	0.1707	77.8	82.9	82.9
SFEC1-LPC	0.870	0.1707	77.4	79.3	82.9
SFEC3-LPC	0.690	0.1707	77.4	79.3	82.9
LDPC	0.769	0.0595	78.9	0.0	100.0
LDPC	0.727	0	94.3	100.0	100.0

TABLE VII
SCENARIO 6: LOSSES SPREAD OUT AMONG THE FLOWS (LOSS 10.6%)

Coding	**Rate**	**Loss**	**Average R**	**% Toll**	**% Acceptable**
SFEC1-GSM	0.710	0.0088	89.3	97.6	100.0
SFEC3-GSM	0.450	0.00002	91.7	100.0	100.0
SFEC1-LPC	0.870	0.0088	87.3	97.6	100.0
SFEC3-LPC	0.690	0.00002	89.6	100.0	100.0
LDPC	0.875	0.0694	76.8	0.0	100.0
LDPC	0.850	0.0391	83.2	100.0	100.0
LDPC	0.833	0.0286	85.4	100.0	100.0
LDPC	0.800	0	94.3	100.0	100.0

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.

First, 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 results showed that the losses are no longer exclusive to a set of flows; they 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. Although recovery with SFEC will improve as the rate deviates from its base value, LDPC is still as effective in such a situation. LDPC provides 100% toll quality calls, while for SFEC, some of the calls have slightly lower quality. SFEC1-GSM approaches the recovery rate of LDPC, but with almost 15% additional redundancy.

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 and it is changed by ±20% of its original value. It was observed that any change in this value decreases the loss rate seen by the VoIP traffic. 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. 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. Fig. 3 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. 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.

More experiments with different scenarios are studied, the results can be found in [12]. In all cases of variations, the effects were similar. The loss rate seen by VoIP flows is only reduced slightly, but the losses are more evenly spread out among the flows. As a result, both SFEC and LDPC can recover a lot better. The recovery of LDPC is still the same –if not better- than that of SFEC.

4) *Network Simulations with Multiple Links*

All the previous experiments were performed over a single link. Here, we consider attacks or transient overloading on multiple links, in particular three overloaded links.

In all the cases, the three links were overloaded by 10%. The opposing traffic on the link is either bursty (as multiple CBR flows as described in previous sections) or a single CBR flow. The two simulated cases are: one where the bursty traffic is on all three links and the other where bursty traffic only exists on the last link. Table VIII compares these two cases with each other and with the single link results from the previous sections. These results are all for bursty traffic with burst spacing of 16 ms (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 bursty traffic. Only network simulations of the form of case 3 will be considered in the following. Table IX shows the resulting call quality for both SFEC and LDPC with different rates for 16 ms burst spacing. Almost 30% or more of the calls coded with SFEC have poor quality whereas acceptable quality can be achieved for all calls when LDPC is used with a rate of 0.75. Similar results are seen for the 7 ms burst spacing.

TABLE VIII
RESULTING QUALITY WITH 16 MS-SPACED BURSTY 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 - bursty traffic on all links***	18.7	72.0	72.0
Case 3	***Three links - bursty traffic on last link only***	18.4	63.4	63.4

TABLE IX
RECOVERY OF LDPC AND SFEC FOR THE 16 MS CASE.

Coding	Rate	Loss	Average R	% Toll	% Acceptable
SFEC1-GSM	0.710	0.07855	76.5	63.4	70.7
SFEC1-LPC	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

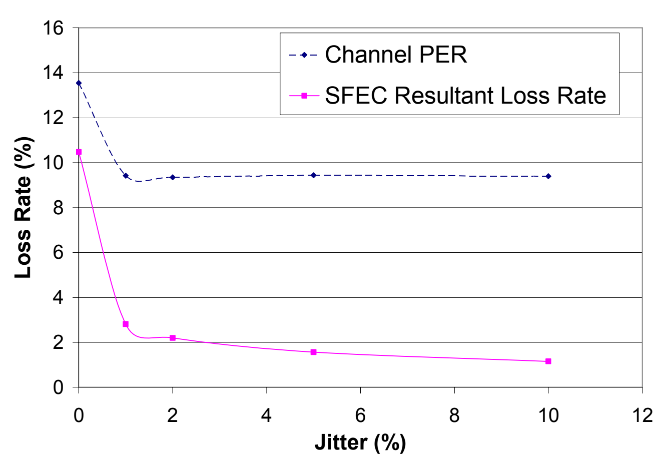

Fig. 3. Effect of jitter in packet inter-arrival times on loss rate (scenario 3).

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, jitter reduced the damage of the attack considerably. However in these experiments, the loss rate remained fairly constant and the damaging effects are not diminished until the level of jitter reaches approximately 5%. Table X and XI illustrate the resultant loss and quality for different encoding schemes at 16 ms and 7 ms 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 at 0.5% of jitter, SFEC encoding has 50% of the calls experiencing un-acceptable quality. The recovery of SFEC improves as jitter is increased and it too recovers 100% of the calls to 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.

And finally, we study how the characteristics of the loss bursts change from a single voice flow compared to the aggregate voice traffic. 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 XII 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 ms 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.

These results show that LDPC coding of aggregate traffic is consistently more effective in protecting against attacks than SFEC applied to individual calls.

TABLE X
THE EFFECT OF ADDING JITTER TO THE 16 MS BURST-SPACING CONFIGURATION.

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

TABLE XI
THE EFFECT OF ADDING JITTER TO THE 7 MS BURST-SPACING CONFIGURATION.

	Jitter (%)	***0***	***1***	***3***	***5***
	Channel loss	***0.196***	***0.189***	***0.188***	***0.140***
SFEC1-GSM (0.71)	***Loss***	0.070	0.074	0.057	0.028
	Acceptable (%)	73.2	63.4	100.0	100.0
SFEC1-LPC (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

TABLE XII
LOSS RUN MEASUREMENT COMPARISON FOR BOTH SFEC AND LDPC.

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

V. Discussion

The proposed approach is dependent on the existence of appropriate high speed links or paths that carry a sufficient 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.

The fact that similarly shaped traffic to the voice calls causes the most damaging effect reinforces the idea that DoS attacks are very probable and damaging for VoIP even if the networks are well-provisioned or calls do not traverse the internet. A slight overload can cause considerable damage over several links in the path. Over-provisioning is not sufficient to tolerate DoS attacks.

The selection of the codec for the original stream only affects the base quality that we use to estimate the resulting voice quality. Changing this value (by using a different codec) does not affect the results of our comparison. In addition, the selection of the codec for the redundant stream defines the redundancy level in the SFEC code. The low bit-rate codecs used for voice (e.g. CELP) use at least as much bandwidth as LPC which is used here. On the other hand, the more compressed codecs for the original stream will have a bit-rate lower than the assumed 32 kbps rate of ADPCM. As a result, alternate SFEC schemes will have higher redundancy than those used in the comparison here and LDPC will outperform these schemes even more than it is illustrated in our results

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 [11]. And in [16], the measured end-to-end delay for an SFEC method was found to be larger than the delays incurred by the FEC coding. LDPC codes have even lower processing delays than the block codes considered in [16]. Moreover, the minimum delay between packets for SFEC cannot be arbitrarily decreased due to the minimum frame size required by the low-quality codec and due to packet header overhead.

Assuming that the voice samples are accumulated from voice calls over a constant period (50ms in this paper), 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. 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. Furthermore, for the same block length, shorter delays can be provided at higher link speeds at a similar level of protection against packet losses. 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. Fig. 4 shows the change in the output loss rate with redundancy of the code for aggregate-coding of a combination of calls with LDPC over different link speeds when the input loss rate is 10%. The link speeds are calculated based on a block length of 50 ms and 128B packets. The loss rate for the LDPC codes generally improves as redundancy is increased. 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 performs better as the link speed is increased (hence larger block size under the same delay constraints). As the redundancy levels are increased below 0.25, the loss rate deteriorates until the code can no longer recover any errors because its capacity has been reached. However, this limit is close to channel capacity.

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 an additional delay required to accumulate sufficient number of packets into a block before processing it. But the arrival times between blocks do not vary much, and any variance that exists is only because the unavoidable queuing delays on the network. Whereas in SFEC, inter-arrival times may 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.

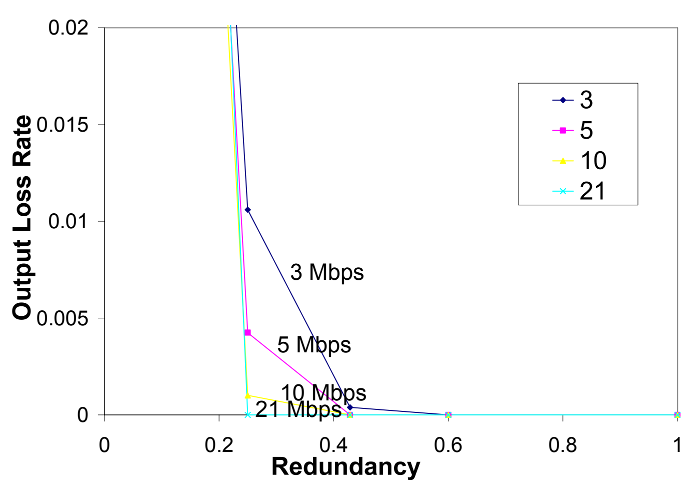

Fig. 4. Output loss vs. redundancy for different link speeds; assuming loss rate = 0.1, delay = 50 ms

A block of code in LDPC consists of a large number of packets (128 and above), whereas 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 [6]. It is emphasized that we only considered fully recovered packets for LDPC in our evaluations and the results could be improved with such concealment techniques.

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 [14], but may require hardware at higher speeds.

VI. Conclusion

We proposed a technique (for service providers) which employs LDPC codes on aggregated VoIP traffic on network links, to improve the resiliency of voice traffic against packet losses. Our experiments have shown that LDPC codes on aggregated traffic 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 consistently in various network conditions.

Acknowledgment

The authors wish to thank Prof. Krishna Narayanan and Abhiram Prabhakar of Texas A&M University for their help in understanding the LDPC codes.

References

[1] 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, pp. 150-159.

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

[3] 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

[4] 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

[5] B. Reynolds and D. Ghosal, "Secure IP telephony using multi-layered protection", presented at the NDSS Symposium, San Diego, CA, 2003.

[6] 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.

[7] 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, Feb. 2000, pp. 307 – 313.

[8] W. Jiang and H.Schulzrinne, "Comparisons of FEC and codec robustness on VoIP quality and bandwidth efficiency," in *Proc. of the ICN Conf.*, Atlanta, GA, Aug. 2002.

[9] J. D. Gobson, *Multimedia Communication: Directions and Innovation*. Academic Press Series in Communications, Networking and Multimedia, San Diego: Elsevier, 2001, pp. 41.

[10] 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.

[11] 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, pp. 232 – 239

[12] C. Al-Najjar, "An approach for improving performance of aggregate voice-over-IP traffic," M.S. Thesis, Dept. Elect. & Comp. Eng., Texas A&M Univ., College Station, TX, 2005, [Online]. Available: http://dropzone.tamu.edu/techpubs/2005/TAMU-ECE-2005-04.pdf

[13] E. Vergetis, R. Guerin, and S. Sarkar, "Benefiting from path diversity with bursty losses," Univ. of Pennsylvania, Philadelphia, PA, 2004, unpublished [Online]. Available: http://www.seas.upenn.edu/~vergetis/EV_ Diversity_04. pdf

[14] 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, pp. 150 – 159.

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

[16] 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, pp. 161 – 170.