

# A SIMULATION STUDY OF ADAPTIVE VOICE COMMUNICATIONS ON IP NETWORKS \*

A. Barberis\*, C. Casetti\*, J. C. De Martin†, M. Meo\*

\* Dipartimento di Elettronica, Politecnico di Torino, 10129 Torino, Italy

† CNR-IRITI, Politecnico di Torino, Italy

e-mail: {casetti,demartin,michela}@polito.it

**KEYWORDS:** IP Telephony, Voice coding, Adaptive

## ABSTRACT

This paper presents simulation results outlining the behavior of rate-adaptive voice communications over IP networks. In the considered architecture, voice coders adapt their rate to the current state of the network, so as to generate only the bandwidth that the network is capable of carrying. An algorithm is proposed for driving the transmission rate of voice sources on the basis of estimations of the network conditions, measured in terms of packet delays and losses.

The effectiveness of the proposed solution is then investigated in various scenarios, which comprise i) a dedicated network in which the available bandwidth is exclusively shared between adaptive voice connections, ii) a scenario in which adaptive voice sources compete with other TCP-like sources, and iii) an uncontrolled network environment. We have compared the performance of the rate-adaptive against the non-adaptive (i.e., fixed-rate) approach for the transport of voice over IP. Using a rate-adaptive approach, more voice communications can be carried while maintaining an acceptable quality of service, even on non-segregated networks.

## 1 INTRODUCTION

Although the value of integrating voice and data transport on the same network has been perceived for many years now, it is only recently, with the huge and still rapidly growing expansion of IP networks, that the issue has seized the attention of individuals, companies and academic institutions worldwide. To accomplish such integration, voice, so far transported

predominantly by circuit-switched networks, is digitized, the resulting samples compressed to save bandwidth, then packetized and sent through the network. At the receiver side, voice packets are reordered, decompressed and then played back to the listener.

Several challenging problems need to be solved to achieve the same level of quality offered by the traditional PSTN's in the context of a packet network like the Internet.

**Delay.** A first major issue is that of delay. In a generic IP network, in fact, delays are not bounded nor predictable. Interactive, two-way telephony, however, imposes strict limits on the allowable overall mouth-to-ear delay. Delay contributions due to the network—primarily, propagation time and average time spent in router queues—can be controlled by ad hoc engineering of the network. Other contributions to the overall delay depend, instead, on the specific Voice over IP architecture employed: parameters like packet size, speech coding algorithm, and playout buffer size all play a role in determining the final delay budget.

**Delay Variations.** Even in the context of carefully engineered networks, however, delay remains variable. The effects of such variability can be controlled by using a buffer, commonly called the *playout buffer*, that allows the receiver to wait for all packets arriving within an acceptable time limit. Smooth playback of the incoming speech signal is then possible.

**Packet Losses.** A third issue is that, in networks based on the best-effort model, packets can be lost. Generally, packet losses occur when a congested router drops the tail of its queues. Effective ways to replace missing speech packets at the receiver are, therefore, needed.

**Speech Compression.** Transport of voice over data networks is generally made possible by speech

---

\*This work was supported in part by a research contract between CSELT and Politecnico di Torino and in part by the Italian Government through the Ministry for University and Research (MURST).

compression techniques. Typical bit rates range from 64 kb/s for companding PCM to 2.4 kb/s for some vocoders. Each speech coding algorithm implements a different tradeoff between output speech quality, algorithmic delay, bit rate, computational complexity and robustness to background noise. The optimal speech coder for VoIP depends on the specific scenario. No speech coding standard has been yet developed with VoIP as primary application. Usually, fixed-rate speech coders are employed. Such coders generate a constant output bit rate, independently of network conditions. If the network cannot sustain the traffic, packets will be delayed and/or dropped.

The overall objective of voice transmission over IP networks is to find an array of technical solutions to these challenges that guarantees the desired level of perceptual quality under most network conditions.

This paper focuses on voice coders whose rate adapts to the current state of the network. If the network is congested, speech is coded at lower bit rates; if the network, instead, is lightly loaded, the speech coder is allowed to operate at higher bit rates. This approach, which may be called Adaptive Voice over IP (or AVoIP, for short), relies on variable bit rate speech coders to generate only the bandwidth that the network is capable of carrying with a given quality of service at any instant in time.

Adaptive Voice over IP has several appealing features. Firstly, an efficient use of network resources is granted since AVoIP applications do not need a rigid partitioning of the link bandwidth (as plain Voice over IP does); instead, it can exploit *statistical* rather than *deterministic* QoS guarantees: as such, AVoIP is a natural candidate for Intserv/Diffserv networks [Wroclawski97]. Also, by adapting gracefully to network congestion, AVoIP flows leave room to signaling and critical in-band flow management sharing the same network facilities. Furthermore, operators can enjoy a great flexibility: by trading off perceived QoS and monetary cost, they can increase the range of services they are offering, e.g., fixed-high-bit-rate calls, fixed-low-bit-rate calls, adaptive bit-rate calls, etc. Finally, new services can be implemented, especially in environments where both data and voice share a narrowband link. For example, technical support being provided through a modem link could use AVoIP to allow customer-salesperson interaction while patches are uploaded on the customer's workstation.

A contribution of this paper is the effort to quantify the performance of AVoIP in terms of those measures which have a major impact on the perceptual quality of service. Typical metrics such as average bit rate and packet loss probability are thus considered, as well as

more sophisticated statistics on packet losses.

The remaining part of this paper is organized as follows: Section 2 explains the rationale for Adaptive Voice over IP and describes a rate control algorithm. Section 3 summarizes the issues affecting perceptual QoS and the measurements of AVoIP performance. The simulation set up is described in Section 4 and finally results are shown in Section 5.

## 2 THE CASE FOR "LIGHT-WEIGHT" CONGESTION CONTROL ALGORITHMS IN IP TELEPHONY

The issue of flow and congestion control for real-time traffic is a delicate one: application-layer coding algorithms and lower-layer rate control schemes should cooperate to respond to critical situations such as network congestion. While a window-based flow control is hardly viable since it would result in stop-and-go source behavior, a "light-weight", QoS-aware congestion control scheme that drives the source encoding rate has the potential to lessen the impact of unruly real-time traffic.

The main goal of such an algorithm is reducing the load on the network when congestion (i.e., queue build-up at intermediate nodes) occurs. If a large part of network traffic is composed of variable-bit-rate regulated voice traffic, and the algorithm regulating each source is devised so as to react to specific "warnings" such as increasing delay or sudden packet loss, then a control strategy may successfully reduce the occurrence of congestion.

The topic of variable-rate adaptive schemes for multimedia streams has been addressed by several researchers in the past. In a general framework [Shenker95, Breslau98], Shenker compared the benefits of rigid and adaptive applications, pointing out that rate-adaptive applications are more robust to network congestion than other classes of adaptive applications (i.e., delay-adaptive applications, which are tolerant of occasional delay bound violations, but whose generation rate is independent of the network congestion). In [Cox80], Cox and Crochiere foreshadow adaptive schemes for variable-bit-rate coders by analyzing the behavior of a buffer control mechanism that has sources react to buffer overflows and underflows. The idea of an end-to-end feedback control is already present in the work by Bially et al. [Bially80], where network-driven voice source based on embedded speech coding schemes are studied, although the involvement of intermediate nodes is still required; in [Yin91], Yin and Hluchyj developed an analytical model for an end-to-end feedback control mechanism for variable bit-rate coders. Bolot and Turletti in [Bolot94a] proposed and tested an adap-

tive video coder based on the state of the network, while in [Bolot94b], they introduced a polling mechanism that allowed a real-time content transmitter to elicit feedback information from receivers and use it to control its own rate.

Similar in spirit to our approach, the works by Busse et al. [Busse95] and Sisalem and Schulzrinne [Sisalem98], employ a sender-based packet loss rate estimation scheme to drive the transmission rate of the real-time source; the latter, in particular, also attempts to have the adaptive real-time source behave fairly towards concurring TCP connections.

Adaptive algorithms can act upon several aspects of a voice connection, i.e. the coding algorithm, the packet length, or the playout delay at the receiver. The first two parameters can be chosen by the transmitter, while the latter can be adjusted by the receiver. The coding algorithm affects the source transmission rate, as well as other aspects of voice communication, while the packet length trades off transmission overhead and packetization time (and thus, delay). The playout delay can be controlled at the receiver by proper regulation of the playout buffer length. The size of the playout buffer can be fixed, determined at call start-up and then kept constant, or allowed to change in the course of the conversation to adapt to variations of the behavior of the network. The main trade-off is between degradation due to delay versus degradation due to packet losses: on the one hand, longer buffers decrease the number of packets discarded because of late arrival (packets are discarded at the receiver if they are “too” late), at the expense of increased delay; on the other hand, shorter buffers can curb the playout delay, but at the expense of increased packet loss rates.

In this paper, we will focus our attention on providing adaptability to IP telephony applications through variable bit-rate coding algorithms. Application-level congestion avoidance must be accomplished by looking at end-to-end metrics, such as packet loss, latency, and jitter.

Packet losses are an indication of a severe network congestion: any action taken following this notification may be belated, thus leading to a prolonged congestion. However, preemptive measures may also be undertaken by acting over an increasing-delay notification. Since the delay experienced by packets traveling from a source to a destination is a function of the number of hops, and of the queue length at intermediate nodes, a sudden increase in the delay observed by the receiver (and relayed to the sender, through periodic receiver’s reports) is often followed by packet loss within the next few round-trip times.

## 2.1 Monitoring of Network Congestion

Making the source rate—that is, the rate at which the speech compressor operates—depend on the state of the network, requires some way of estimating such state, since the IP service model does not offer congestion notification. The focus is on two main measures: the effective capacity of the link, largely determined by the bottlenecks along the connection, and the detection of temporary congestion. Based on such measures, the rate control algorithm will select bit rates compatible with the estimated capacity and respond to short-term clogging of the link.

In order to implement an effective adaptive algorithm, estimates of quantities related to the ongoing connection and needed by the adaptive algorithm are to be carried out. Such a procedure involves several aspects:

1. since delay estimates are essentially based upon a correct alignment of both the source’s and the receiver’s clocks, the skew between these clocks should be estimated;
2. while single delay estimates can be made on the reception of every RTP packet, the loss rate estimates should reflect the ratio between lost and expected packets, computed over a given time span;
3. some quantities require averaging in order to compare the current estimates with the expected behavior: moreover, the average should be weighted so as to smooth out fluctuations;
4. what information to exchange between source and receiver (i.e. who computes the estimates, who averages the quantities, etc.) should be determined.

## 2.2 The Adaptive Algorithm

The rate control algorithm described in this paper, starting from delay and loss measurements relayed by RTCP reports, tries to regulate the output rate of several voice sources. It does not rely on explicit congestion indications by either the intermediate nodes or by receivers. Rather, it uses the information that can be carried by cyclic RTCP receiver reports to let the source know the state of the on-going connection. The control algorithm runs at the source itself. An alternative scheme would be one where the algorithm runs at the receiver, where all the raw estimates are available, not only down-sampled, smoothed versions of the same. Such receiver based scheme, however, would require a signaling mechanism to convey rate increase/decrease commands to the source.

The adaptive algorithm follows the additive increase/multiplicative decrease paradigm that is successfully employed by many congestion control algorithms, such as TCP or ABR. The basic idea is that the source coder should reduce its rate when packet delays have been observed to have considerably increased above a ‘high-mark threshold’. Also, it should drastically decrease the rate to, say, half its value, when severe congestion is detected, (i.e. packets are lost). Besides, it should switch to higher rates if no packets are lost and if the delay decreases below a ‘low-mark threshold’.

Clearly, since we are looking at an IP telephony scenario, both parties can, at different times, be considered source and destination. In order to simplify the description, we will decouple these interactions and only consider a communication between a generic source and a receiver.

The algorithm flow chart is reported in Figure 1 and its pseudo-code is provided in Figure 2. It is periodically executed by the source, whenever a receiver RTCP report reaches it. It should be pointed out that, if no RTCP is received by the source for longer than 5 seconds, the source takes it upon itself to reduce the rate by a step; thus, it conservatively keeps its transmission speed from increasing only because a lack of statistics (i.e., no losses, stable delay) is interpreted as “favorable” network conditions.

Throughout the algorithm description, the following parameters are used (when possible, the reference values used in simulation are reported between brackets):

- **curr\_time**: current time;
- **dec\_int**: minimum time between back-to-back decrements (1 s in the simulations);
- **inc\_int**: minimum time between back-to-back increments (3 s in the simulations);
- **var\_time**: time when the last rate decrement/increment occurred;
- **Max\_Loss**: loss ratio that triggers the algorithm’s rate decrement (3% in the simulations);
- **Min\_Loss**: loss ratio that triggers the algorithm’s rate increment (7% in the simulations);
- **peer\_SLR**: smoothed loss ratio;
- **avg\_delay**: average (smoothed) delay;
- **latency\_thresh**: minimum relative latency increase that triggers a rate decrement (in the simulations, we used **latency\_thresh** = 0.1, indicating that the rate should be lowered if the latest

delay measure is 10% higher than the mean so far computed);

- **peer\_delay**: latest measure of delay.

The possible values for the bit rate vary between 8 kb/s and 64 kb/s in steps of 8 kb/s.

The numerical values of the parameters that were used in simulation often are the results of “sensible” choices, driven by experience and observation, rather than of thorough analysis; as is often the case with complex systems, a choice of parameters that optimizes the behavior of the system (i.e., maximizing utilization and/or minimizing losses) is computationally infeasible. We have however verified through extensive simulations that the algorithm is stable and behaves in the predicted way, although we cannot guarantee that it yields optimal results.

Smoothed quantities are computed through an exponential averaging function:  $\bar{x}(n) = \alpha\bar{x}(n-1) + (1-\alpha)x(n)$  where  $\bar{x}(n)$  is the averaged quantity and  $x(n)$  is the  $n$ -th sample;  $\alpha$  is a positive real number smaller than 1 ( $\alpha = 0.2$  in the simulations) and it determines speed and accuracy with which the average  $\bar{x}(n)$  follows the trend of  $x(n)$ .

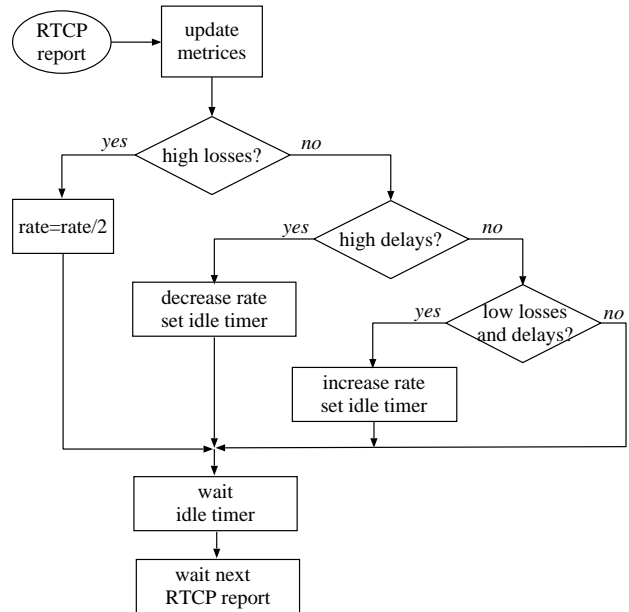


Figure 1: Flow-chart of the adaptive algorithm

Being a second-order estimate, and thus more delicate to be properly accounted for, the jitter is not used in this algorithm. However, we plan to extend the scope of our adaptive algorithms to take into account the jitter in the future.

```

halve = decrement = increment = FALSE;
/* Compare loss rates */
if (peer_SLR > Max_Loss) then
  if (curr_time > var_time + dec_int) then
    halve = TRUE;
    var_time = curr_time;
  endif
else if (peer_SLR < Min_Loss) then
  if (curr_time > var_time + inc_int) then
    increment = TRUE;
    var_time = curr_time;
  endif
endif
/* Compare delay estimates */
if (peer_delay > avg_delay + avg_delay * latency_thresh) then
  if (curr_time > var_time + dec_int) then
    decrement = TRUE;
    var_time = curr_time;
  endif
else if (peer_delay < avg_delay + avg_delay * latency_thresh) then
  if (curr_time > var_time + inc_int) then
    increment = TRUE;
    var_time = curr_time;
  endif
endif
/* Change rate according to choice (decrement has priority) */
if (halve = TRUE) then
  HalveRate();
else if (decrement = TRUE) then
  DecreaseRate();
else if (increment = TRUE) then
  IncreaseRate();
endif

```

Figure 2: Pseudo-code for the adaptive algorithm

### 3 PERCEPTUAL QUALITY OF SERVICE

Perceptual quality of service is a function of all factors affecting the speech signal from the moment it is produced by the speaker’s mouth to the moment it is perceived by the listener’s ear. Focusing only on what takes place between the output of the A/D converter at the transmitter and the input of the D/A converter at the receiver, several main factors impacting Perceptual QoS can be identified.

Firstly, in most cases the stream of speech samples produced by the A/D converter will be compressed by a *speech coder* [Gersho94] [Kleijn95]. Uncompressed telephone-bandwidth speech, indeed, requires bit rates greater than 100 kb/s — 64 kb/s if non-uniform quantization is employed, as in ITU-T standard G.711 (log PCM). Modern speech coders, instead, can operate at bit rates of 8 kb/s or lower. Such bit rate savings, however, generally come at the cost of a certain degree of perceptual degradation with respect to the original speech signal. Thus the first factor affecting the perceptual quality of service is the quality generated by the specific speech coder employed in the application; its output quality represents the upper bound to the perceptual QoS achievable by the system, that is, the quality under ideal conditions.

Note also that modern speech coders usually operate on segments of speech, called *frames*, not on individual samples. A delay is, therefore, necessarily introduced between the instant a sample enters the speech coder and the instant when the corresponding synthesized sample exits the decoder. The delay introduced by the speech coder is usually called *algorithmic delay*. As the *overall* delay (algorithmic, transmission, queueing, buffering, decoding, etc.) increases, the perceptual quality decreases, as will be discussed below.

The compressed bit-stream produced by the speech coder is then packetized and sent on the network, where two other important factors come into play: packet losses and potentially unbounded delay.

Packet losses are typical of speech transmission over packet networks, where packets can arrive too late for timely playback, or even not arrive at all. In this case a whole speech frame (typically 5 to 25 ms of speech, or more if a packet contains more than one frame) is either unavailable or declared unusable due to bit errors in the perceptually crucial bits of the frame (e.g., the most significant bits of the energy). In both cases, techniques to fill the gap corresponding to the corrupted or missing frame are needed. In the case of high rate waveform coders, such as PCM, sample-based techniques such as repetition, interpolation and extrapolation are used. In more complex coders, instead, the parameters of the previous, correctly received frame are generally used, often by repeating them, with a muting mechanism activated in case of consecutive frame erasures. Note that some speech coders are quite sensitive to frame erasures; under 3% random frame erasures, for instance, the quality of the widely used ITU-T G.729 8 kb/s standard drops by about 0.5 MOS, a very significant drop. For 5% frame erasures losses, the quality drops even lower, to less than 3 MOS. In general, the impact of packet losses varies depending on the statistics of the frame erasures (random, burst) and the specific speech coder at hand.

As mentioned before, another very significant factor influencing perceptual QoS is delay. Besides the delay introduced by the speech coder at the transmitter, additional delay is introduced at several other points during transmission, some of it unbounded due to the IP best-effort model. The perceptual effects of delay on telephony have been extensively studied in the past. ITU-T Recommendation G.114 reports the results of several subjective tests, which led to the conclusion that for most users delay of less than 150 ms are not objectionable. Delays between 150 ms and 400 ms, instead, are deemed tolerable, provided that the service provider is aware of the effects that such levels of delay have on quality. Delays above 400 ms are not acceptable for general network planning. Note

that, unlike most other speech transmission scenarios, in Voice over IP delay can vary during the course of the same conversation; the perceptual effects of such variability have not been extensively studied yet.

To guarantee a certain level of perceptual QoS is, therefore, equivalent to guaranteeing that:

1. the chosen speech coder is capable of delivering the desired speech quality;
2. there is enough capacity to reliably transmit the bit rate generated by the speech coder;
3. packet losses are below a certain threshold;
4. overall mouth-to-ear delay is contained within a given level.

### 3.1 Perceptual Quality and AVoIP

When the compressed bitstream is transmitted over a time-varying channel, estimating and controlling the perceptual quality becomes particularly demanding. In the specific case of voice over IP, the IP network can be seen as a lossy channel with time-varying *capacity*, *packet loss rate* and *delay*.

The instantaneous capacity is determined in part by structural characteristics of the connection, such as its bottlenecks, in part by traffic dynamics. In any case, the basic assumption is that transmission rates above the current capacity result in packet losses and/or increased delays, that is, in potentially significant degradation of the perceptual quality.

If a fixed rate source coder is employed and the capacity drops below the operating bit rate, the perceptual quality will be more or less severely affected, with no way of controlling the degradation. In this scenario, countermeasures can only be implemented at system level, e.g., increasing the capacity of the system and implementing mechanisms that ensure that the capacity will never, or very rarely, drop below a certain safety level.

If, instead, a network-driven variable bit rate speech coder is used, the time-varying nature of the channel can be, at least to some extent, matched, resulting in greater control of the perceptual QoS and in more efficient use of network resources. More specifically, an adaptive approach makes possible to choose the optimal bit rate for the current instantaneous channel conditions. Under the assumption that lower bit rates are correlated with lower levels of packet losses and/or delay, there is a threshold below which the decrease in quality entailed by the lower bit rate is more than offset by a lower level of degradation due to losses and delay. In other words, the adaptive approach makes possible to choose the best trade-off between bit-rate, packet losses and delay and thus to operate, at least

in principle, at the instantaneous optimal operating point.

Although most speech coding standards are fixed rate, several variable bit rate speech coders are available. A recent example of a variable bit rate solution is the new GSM Adaptive Multi-Rate speech coding standard [Ekudden99]. In GSM-AMR one of 8 rates ranging from 4.75 kb/s to 12.2 kb/s is selected, depending on the instantaneous condition of the wireless channel, a scenario with several points of contact with the one considered in this paper. Another example is the new ISO MPEG-4 audio standard, which includes a variable bit rate CELP speech coder operating at bit rates between 3.85 kb/s and 12.2 kb/s [CELP98].

As a measure of perceptual quality, the following performance indices are taken into consideration in this paper:

- average and instantaneous bit rate of the voice coder
- average and instantaneous delay experienced by packets in traversing the network (it comprises the queueing and transmission delays)
- the distribution, expressed in terms of the probability mass function, of the number of consecutive packets (belonging to the same voice stream) which are successfully received
- the distribution of the number of consecutive packets which are lost

The last two measures assess the ease in reconstructing a gap of missing frames in a corrupted voice stream.

## 4 SIMULATION SETUP

All simulations whose results are presented in this paper were run using a simple bottleneck network topology. More complex topologies, while more realistic, have the drawback of providing results that do not lend themselves to reliable interpretation. The bottleneck link is supposed to have a limited bandwidth equal to 256 kb/s, while secondary links are assigned a sufficiently high bandwidth so that all traffic traveling from a source to a destination only experiences congestion on the bottleneck link. Link latencies are 3 ms for the bottleneck link and 1 ms each for the secondary links (resulting in a 10-ms round-trip-delay). The buffering capacity of the bottleneck link router is equal to 16 kB, and buffer congestion is dealt with using a drop-tail strategy. Sources are placed at both ends of the bottleneck link, although the results we examine only refer to the set of sources

transmitting in one direction (the presence of reverse-direction sources may however affect the reception of RTCP packets carrying Receiver Reports relating to forward-direction communications). Simulated voice sources generate a continuous traffic flow throughout the simulation.

Our study was targeted at studying three main traffic scenarios:

- *homogeneous traffic*: only VoIP sources are active, thus simulating a segregated network environment, where resources are entirely devoted to voice traffic;
- *interfering FTP traffic*: VoIP sources compete for bandwidth with greedy FTP connections sharing the same network path;
- *interfering Pareto/UDP traffic*: a single VoIP source and interfering traffic modeled as Pareto *On/Off* connections;

The simulation tool is the well-known `ns` simulator [ns], modified to suit our purposes.

## 5 SIMULATION RESULTS

The performance metrics presented in this section are meant to compare the performance of VoIP communications with and without adaptivity, under the scenarios outlined in section 4. Therefore, every experiment was replicated several times, using adaptive-rate VoIP sources in one instance, and constant-bit-rate (CBR) VoIP sources otherwise. Rate adaptation was achieved by having the sources transmit variable-size packets at fixed rate (i.e., on rate increases, the source starts sending larger packets, without changing the packet inter-transmission time, fixed at 125 ms). All CBR sources send 512-bytes-long packets, while the transmission time is determined according to the source rate (and is kept constant). In this paper we will examine comparisons between AVoIP and fixed-rate sources transmitting at 8, 16, 24 and 32 kb/s. The respective performances were compared in terms of loss rates and delivery delays; the loss process was further analyzed through the observation of several statistical properties, characterizing the quality of the communication perceived by end users, i.e.:

- *back-to-back lost packets*: average, variance and pdf of the number of packets that were consecutively dropped by router *r1* because of the onset of congestion.
- *back-to-back received packets*: average, variance and pdf of the number of packets that were delivered and played back on time, with no packet drops in between.

### 5.1 Homogeneous Traffic

The first set of experiments focuses on several voice sources competing against each other on a bottleneck link, without interfering traffic. This situation might represent a Virtual Private Network of a small VoIP operator, carrying all IP phone calls on the same network infrastructure. Also, it might depict the aggregate behavior of a subset of traffic in a Diff-Serv environment: following the guidelines set forward in [Blake98], all voice traffic could be marked so as to receive an adequate service level, through per-hop-behaviors (PHB) negotiated between customer and provider.

Results were obtained for 10 and 20 active sources in each direction.

Figure 3 provides a first glimpse as to the behavior of the algorithm for a sample AVoIP source over a 250-seconds time span, when 20 AVoIP sources share the bottleneck link: from top to bottom, the source rate, the end-to-end delay and the loss percentage are shown as a function of time. A straightforward observation is that, aside from the very beginning, the algorithm is mostly operated through loss indications resulting in rate halving, while single-step rate reductions due to delay indications rarely take place (an example can be seen approximately 150 s into the experiment). This behavior appears to be peculiar of the considered scenario, since single-step reductions are more frequent when non-adaptive interfering traffic is present, as will be shown below.

Further indications can be drawn by comparing loss and delay metrics when adaptive only and constant-bit-rate (CBR) only traffic is simulated: Tables 1 and 2 detail these comparison for 20 and 10 sources respectively (the rate between brackets for AVoIP sources refers to the average rate at which sources have been transmitting during the simulation). Several observations are in order:

- AVoIP sources, while not necessarily providing an optimal performance (i.e., maximal network utilization and no losses) in all three cases, succeed in finding a suitable 'operating point' without no prior knowledge of the available bandwidth (as is often the case); the high loss percentage and high delays experienced when 20 sources share the link are determined by the optimal operating point being placed between 8 and 16 kb/s, i.e. between two adjacent 'states' that either yield no loss or cause more than 30% of lost traffic; the algorithm thus suffers from the shortened dynamics;
- the loss rates of AVoIP sources are of interpolating nature, i.e., they roughly correspond to the loss rates of CBR sources operating at a

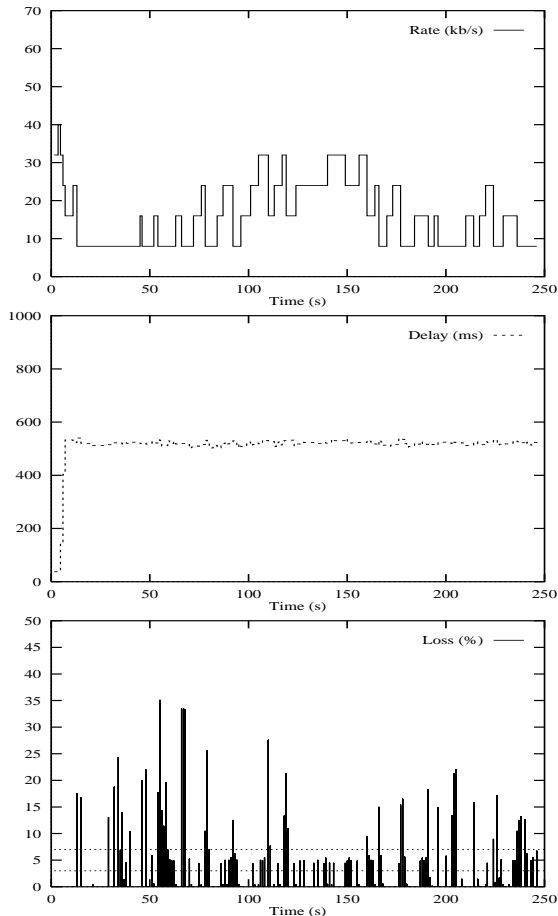


Figure 3: Sample AVoIP source behavior under homogeneous traffic (20 sources)

rate equal to the average AVoIP rate (this assumption is based on a linear interpolation of CBR sources behavior as summarized by Tables 1 and 2); again, this result is peculiar of the scenario we are examining, as will be discussed below;

- when large numbers of sources are sharing the link, the delivery delays for AVoIP sources are quite high; beside the lack of dynamic range discussed above, this is a consequence of the algorithm not acting upon *high* delays but, rather, upon *highly varying* delays (Figure 3 highlights that the delay variation is not high).

Much though the speech quality perceived by end users cannot be ascertained in a simulation experiment like the ones we conducted, nonetheless we tried to overcome this limitation through a simple statisti-

Source Type	Loss Perc.	Delay
AVoIP (12.44 kb/s)	20.4 %	0.52 s
CBR 32 kb/s	65.0 %	0.53 s
CBR 24 kb/s	53.6 %	0.53 s
CBR 16 kb/s	31.7 %	0.51 s
CBR 8 kb/s	0 %	0.06 s

Table 1: AVoIP and CBR comparison (20 sources)

Source Type	Loss Perc.	Delay
AVoIP (18.81 kb/s)	0.02 %	0.05 s
CBR 32kb/s	29.77 %	0.52 s
CBR 24kb/s	7.18 %	0.51 s
CBR 16kb/s	0 %	0.05 s
CBR 8kb/s	0 %	0.04 s

Table 2: AVoIP and CBR comparison (10 sources)

cal analysis of the flow of received packets. First of all, we compared average and variance of the length of *loss bursts*, i.e. back-to-back packets that were not received, and that, conceivably, led to 'playback holes' hard to recover even with the techniques described in section 3. Similarly, we kept trace of the number of consecutive packets that were correctly received. These quantities are reported, respectively, in Tables 3, 4 for the case of 20 sources (it should be noted that CBR sources at 8 kb/s are missing: since they do not register any losses, their statistics would have been meaningless). These results point out that, on average, AVoIP sources lose 1-2 packets every 6-7 received packets. Compared to 16 kb/s CBR sources, which lose, on average, 2 packets and more every 4 received packets, one can assume that the perceived speech quality is far superior for AVoIP sources.

## 5.2 Interfering FTP traffic

The second scenario we have considered features the staggered activation of 3 greedy FTP file transfers (they are turned ON at 10-second intervals starting at times 50 s and 260 s, and are turned OFF at 10-second intervals starting at time 170 s). FTP sources send 1,000-bytes-long packets and the TCP connections they are supposed to establish have a maximum window of 20 packets. 10 voice sources are active throughout the simulation.

The behavior of a sample AVoIP source can be observed in Figure 4 over a 400-second simulation run. By comparing the rate curve with the one shown in Figure 3, we have an exemplification of how the delivery delay can drive the control algorithm: it provides a smooth transition to lower bit rates (as can be seen on several instances during the simulation) by single-



Source Type	Average	Variance
AVoIP	1.70	2.54
CBR 32kb/s	3.83	22.69
CBR 24kb/s	3.08	17.94
CBR 16kb/s	2.14	8.53

Table 3: Average and Variance of length of packet loss bursts for AVoIP and CBR sources

Source Type	Average	Variance
AVoIP	6.15	57.68
CBR 32kb/s	2.01	6.24
CBR 24kb/s	2.46	16.93
CBR 16kb/s	4.00	39.00

Table 4: Average and Variance of length of packet reception bursts for AVoIP and CBR sources

step decreases, as opposed to the steeper transitions observed when only AVoIP sources were present, in a "stationary" scenario. We can thus conclude that our algorithm is all the more effective when network conditions change (i.e. when sources switch on and off).

The comparison with the behavior of a CBR source in this scenario allows us to investigate some statistical properties of the packet reception pattern, an observation that can be linked to users' satisfaction with the quality of the received speech. Figure 5 shows a comparison of the probability distributions of the number of back-to-back received packets, for the AVoIP case and for the CBR case (note: bursts longer than 200 packets are cumulatively accounted for at the x-axis value 200). Of course, longer, uninterrupted streams of received packets can be translated into a better perceptual quality, and AVoIP sources (whose measured average rate was 11.6 kb/s) appear to have the edge when compared to CBR sources.

### 5.3 Interfering Pareto/UDP traffic

The last scenario we have examined features a single voice source and interfering traffic modeled as Pareto On/Off sources over UDP connections. As many as 10 such sources are active at the same time (the activation pattern is similar to the one described in the previous section). It should be pointed out that the disactivation of a source is potentially disruptive of the LRD effect introduced by the Pareto distribution, since the source is switched off even in the midst of an On period.

Since UDP traffic is not adaptive (while FTP traffic employs TCP congestion control), we expect this scenario to be more testing than the previous one.

The Pareto sources have average On and Off time




Figure 4: Sample AVoIP source behavior with interfering FTP sources (10 AVoIP sources + 3 FTP sources)

equal to 2 s, average nominal rate equal to 128 kb/s, packet size fixed to 1 kB, and Pareto  $\alpha$  shape parameter equal to 0.5.

Interesting observations can be drawn from Figure 6, where rate, delay and loss dynamics for the only voice source are plotted. The intervention of the varying-delay indications, resulting in single-step decrements, are visible not only when UDP sources are activated, but also when they are deactivated (thin rate peaks around 200 s): the algorithm curbs the aggressiveness of the AVoIP source as UDP traffic is slowly leaving the buffer at the router. No losses are detected after the first UDP sources are deactivated.

Top plot in Figure 7 shows the probability distribution of the length of loss bursts of the AVoIP source, and it should be compared with similar plots in the same figure referring to CBR sources at 16 and 32 kb/s. To put things in perspective, the AVoIP

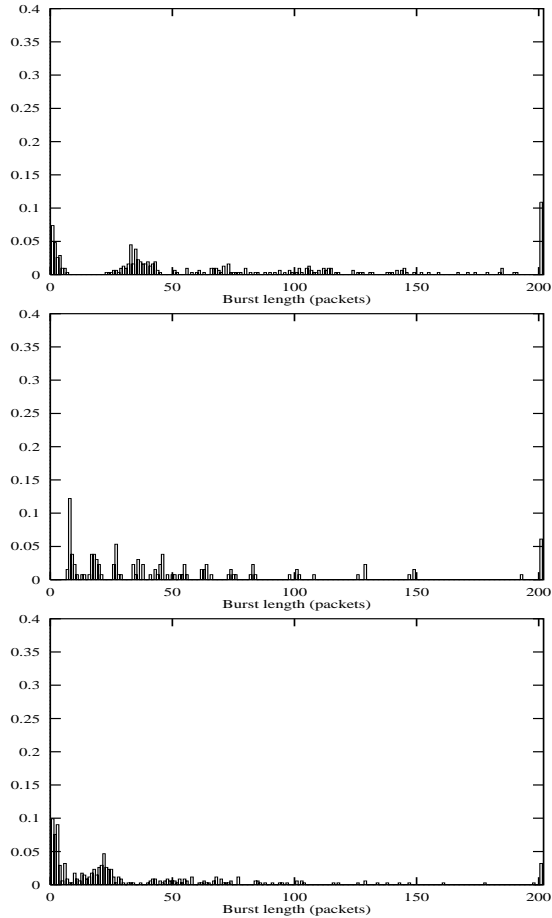


Figure 5: Distribution of number of back-to-back received packets by the AVoIP source, 8- and 16-kb/s CBR sources (from top to bottom)

source reached an average of 27.88 kb/s during the simulation. By comparing the distribution of losses, we can conclude that long, disruptive playback gaps (5 or more lost packets) were rarely observed, while this event was more frequent for CBR sources, even at rates much lower than the average rate achieved by the AVoIP source.

## 6 CONCLUSIONS

We have presented a study on the effectiveness of using network-driven adaptive voice sources for interactive voice communications over IP networks. In the proposed architectural solution variable bit rate voice coders adapt their rate to the time varying network conditions by means of a control algorithm whose aim is maximizing the utilization of the available band-

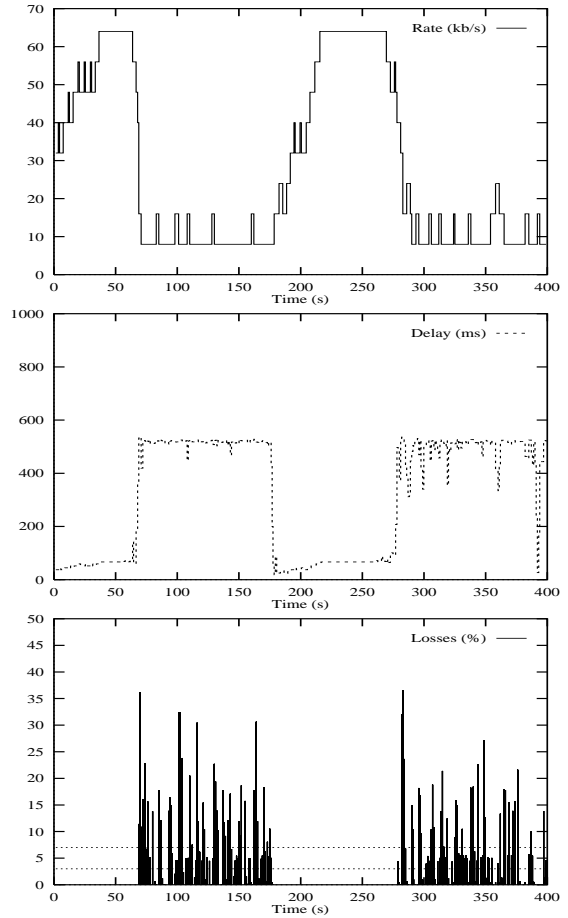


Figure 6: Sample AVoIP source behavior with interfering FTP sources (10 AVoIP sources + 3 FTP sources)

width while reducing and preventing the occurrence of packet losses.

The performance of the proposed approach is evaluated in various scenarios, which comprise a network dedicated to the exclusive use of adaptive voice sources as well as heterogeneous uncontrolled network environments.

Results are also compared against the performance of non-adaptive sources; it is shown that the adaptive approach is more effective, being able to carry more voice communications while maintaining an acceptable quality of service, even on non-segregated networks.

## REFERENCES

- [Wroclawski97] J. Wroclawski. Specification of the Controlled-Load Network Element Service.

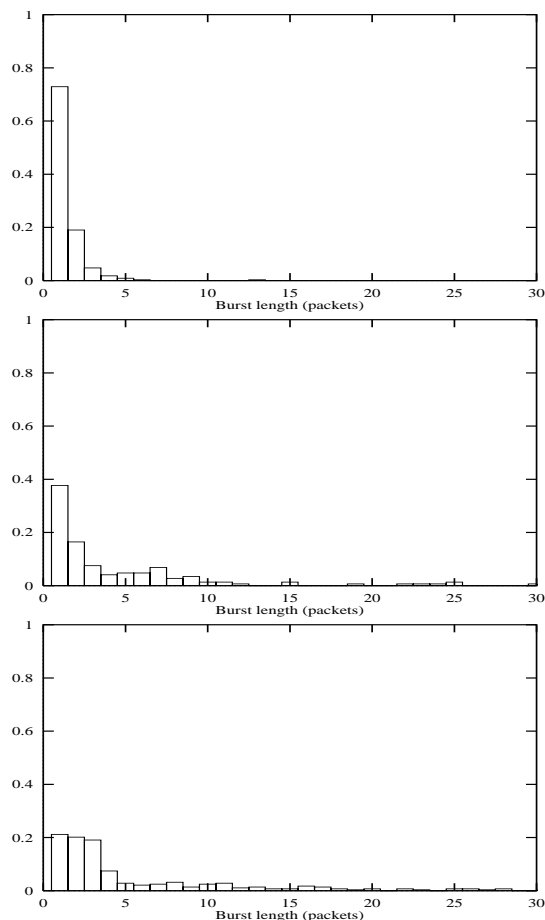


Figure 7: Distribution of number of back-to-back lost packets by the AVoIP source, 16- and 32-kb/s CBR sources (from top to bottom)

RFC 2211, September 1997.

- [Shenker95] S. Shenker. Fundamental Design Issues for the Future Internet. *IEEE Journal on Selected Areas in Communications*, September 1995.
- [Breslau98] Lee Breslau and Scott Shenker. Best-Effort versus Reservations: A Simple Comparative Analysis. In *ACM Computer Communication Review*, vol. 28, no. 4, pp. 3–16, Sep. 1998.
- [Cox80] R.V. Cox and R.E. Crochiere. Multiple User Variable Rate Coding for TASI and Packet Transmission Systems. *IEEE Transactions on Communications*, 28(3):334–344, March 1980.
- [Bially80] T. Bially, B. Gold, and S. Seneff. A Technique for Adaptive Voice Flow Control in Integrated Packet Networks *IEEE Transactions*

*on Communications*, Vol. COM-28, pp. 325–333, Mar. 1980.

- [Yin91] N. Yin and M.G. Hluchyj. A Dynamic Rate Control Mechanism for Source Coded Traffic in a Fast Packet Network. *IEEE Journal on Selected Areas in Communications*, 9(7):1003–1012, September 1991.
- [Bolot94a] J.C. Bolot and T. Turletti. A Rate Control Mechanism for Packet Video in the Internet. In *Proceedings of IEEE INFOCOM'94*, Toronto, Canada, June 1994.
- [Bolot94b] J.C. Bolot, T. Turletti, and I. Wakeman. Scalable Feedback Control for Multicast Video Distribution in the Internet. In *Proceedings of SIGCOMM'94*, London, UK, August 1994.
- [Busse95] I. Busse, B. Deffner and H. Schulzrinne. Dynamic QoS Control of Multimedia Applications Based on RTP. In *Proceedings of 1st Int. Workshop on High Speed Networks and Open Distributed Platforms*, St. Petersburg, Russia, June 1995.
- [Sisalem98] D. Sisalem and H. Schulzrinne. The Loss-Delay Adjustment Algorithm: A TCP-Friendly Adaptation Scheme. In *Proceedings of International Workshop on Network and Operating System Support for Digital Audio and Video (NOSS-DAV)*, Cambridge, England, July 1998.
- [Gersho94] A. Gersho. Advances in Speech and Audio Compression, *Proceedings of the IEEE*, June 1994, Vol. 82, Issue: 6.
- [Kleijn95] W.B. Kleijn and K.K. Paliwal (editors). *Speech Coding and Synthesis*. 1995, Elsevier Science.
- [Ekudden99] E. Ekudden, R. Hagen, I. Johansson, and J. Svedberg. The Adaptive Multi-Rate Speech Coder, In *Proceedings IEEE Workshop on Speech Coding Proceedings*. Haikko Manor, Porvoo, Finland, June 1999.
- [CELP98] ISO/IEC. Information technology - coding of audiovisual objects, part 3: Audio, subpart 3: Celp. 14496-3 FDIS, December 1998.
- [ns] ns-2, network simulator (ver.2), LBL. URL: <http://www-mash.cs.berkeley.edu/ns>.
- [Blake98] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss. An Architecture for Differentiated Services. RFC 2475, December 1998.