

COMPUTER NETWORKS and ISDN SYSTEMS

The International Journal of Computer
and Telecommunications Networking

Theme Issue

ITC 14 Special Sessions Presentations

Guest Editor:

J.W. Roberts

KURT F. WENDT LIBRARY
COLLEGE OF ENGINEERING

APR 04 1996

UW-MADISON, WI 53706



ELSEVIER Amsterdam – Lausanne – New York – Oxford – Shannon – Tokyo

RPX Exhibit 1131



On retransmission-based error control for continuous media traffic in packet-switching networks

Bert J. Dempsey *, Jörg Liebeherr, Alfred C. Weaver

Computer Science Department, University of Virginia, Charlottesville, VA 22903, USA

Accepted 16 February 1995

Abstract

Distribution of continuous media traffic such as digital audio and video over packet-switching networks has become increasingly feasible due to a number of technology trends leading to powerful desktop computers and high-speed integrated services networks. Protocols supporting the transmission of continuous media are already available. In these protocols, transmission errors due to packet loss are generally not recovered. Instead existing protocol designs focus on preventive error control techniques that reduce the impact of losses by adding redundancy, e.g., forward error correction, or by preventing loss of important data, e.g., channel coding. The goal of this study is to show that retransmission of continuous media data often is, contrary to conventional wisdom, a viable option in most packet-switching networks. If timely retransmission can be performed with a high probability of success, a retransmission-based approach to error control is attractive because it imposes little overhead on network resources and can be used in conjunction with preventive error control schemes. Since interactive voice has the most stringent delay and error requirements, the study focuses on retransmission in packet voice protocols. An end-to-end model of packet voice transmission is presented and used to investigate the feasibility of retransmission for a wide range of network scenarios. The analytical findings are compared with extrapolations from delay measurements of packet voice transmission over a campus backbone network.

Keywords: Error control; ARQ; Continuous media; Packet-switching; Retransmission; Packet voice

1. Introduction

Continuous media services such as voice and video were traditionally carried over circuit-switching networks. Using packet-switching networks for these services has become increasingly attractive due to technology trends that have enabled high-speed multiservice networks. Since voice and video are inherently variable rate sources [4,12], statistical multiplexing gains in packet-switching networks result in a more efficient use of network resources than circuit switching. However, the distribution of continuous media is fundamentally different from traditional reliable data transfer, such as file transfer and remote login, since continuous media is sensitive to end-to-end delays and variations of the delays.

*Corresponding author. Current address: School of Information and Library Science, University of North Carolina, Chapel Hill, NC 27599, USA. E-mail: bert@ils.unc.edu.

The distribution of continuous media across a packet-switching network requires consideration of encoding schemes, end-to-end network delays, network delay variations, and packet loss, all of which significantly affect the playback quality at the receiving site:

- In recent years, considerable progress has been made in the design of efficient digital encoding techniques for analog audiovisual data [11]. The selection of an encoding scheme represents a trade-off between consumption of bandwidth in the network and playback quality at the receiving site since low-bit-rate encoding schemes result in a less precise reconstruction of the original analog signal.
- In an interactive continuous media session, human perception factors produce a requirement for bounded roundtrip delays. If roundtrip delays are too long, the interactive nature of the session is degraded.
- Statistical multiplexing introduces variations in the network delay experienced by individual packets. These variations are referred to as *delay jitter*. Delay jitter can lead to interruptions in the continuous playback of the continuous media stream at the receiver.
- Unlike data transmission, most continuous media data does not require reliable delivery, though its tolerance for packet loss is low. Techniques for robust signal processing in the presence of packet loss can significantly improve loss tolerances, but even the loss of a single packet may noticeably degrade playback quality at the receiver.

In this paper we study error control for voice transmission in packet-switching networks. We selected interactive packet voice for our work since it has very stringent delay and error requirements. While our investigation focuses on voice transmission, most of our concepts apply to other forms of continuous media traffic, e.g., low-bandwidth video. We examine the feasibility of retransmission-based error recovery for continuous media traffic. In order to be effective, retransmissions of lost packets must be completed within the delay constraints of the packet stream. However, if timely retransmission can be achieved with a high probability of success, a retransmission-based approach to error control is attractive because it imposes little overhead on network resources. Note that retransmission-based error recovery can be used in conjunction with extant preventive error control schemes such as forward error correction or channel coding.

We employ analytical modeling techniques to investigate the effectiveness of retransmission for different network scenarios. To explore the relationship between our theoretical results and the dynamic behavior of voice transmission over existing networks, we present measurements of the actual delays experienced by voice transmission running over a campus backbone network. Our results indicate that retransmission-based error recovery can be effective for many end-to-end transmission scenarios in current networks.

The remainder of this paper is structured as follows. In Section 2 we review issues that must be addressed by protocols for voice distribution in packet-switching networks. In Section 3 we develop an analytic model for the end-to-end transmission of packet voice and derive a performance metric for timely retransmission in the presence of errors. We present examples where we apply the performance metric under variations of network parameters. In Section 4 we provide measurements of voice packets on a multiple-segment local area network and compare the empirically obtained data with our theoretical findings. In Section 5 we present the conclusions of the paper.

2. Protocol issues

Continuous media protocols must have mechanisms to address all factors that may degrade the quality of remote playback. In this section we briefly discuss important issues for maintaining high quality voice transmission, and discuss how these issues are resolved in extant packet voice protocols. An important consideration in the design of packet voice protocols is that speech is actually an alternating series of activity periods, or *talkspurts*, followed by silence periods with the activity periods constituting only around 40% of the total time [4].

2.1. Encoding and packetization

The packet voice source continuously collects and buffers digitized voice samples. After a fixed period of time, the so-called *packetization interval*, voice samples collected by the audio hardware are placed into a network packet, and the packet is submitted to the network. Typical packetization intervals range from 10–50 ms [19].

Given a fixed packetization interval, the encoding scheme determines the actual number of bits per packet. The ubiquitous pulse code modulation (PCM) encoding scheme for voice [6] samples every 125 μ s with 8 bits per sample to yield a 64 Kbit/s channel. Bandwidth reduction can be achieved through the use of fewer bits per sample, less frequent sampling, suppression of transmission during silence periods, and compression of the digitized data. Adaptive differential pulse code modulation (ADPCM) [7], for example, encodes only the difference between consecutive samples, reducing the number of bits per sample to 2–5 bits. Coding techniques with even lower bit rates, e.g., Linear Predictive Coding (LPC), exist, though speech fidelity is frequently poor [11].

2.2. Roundtrip delay

Behavioral studies [5,13] have shown that roundtrip delays above a certain threshold degrade the interactive nature of the conversation. Quantifying this factor is difficult since individual human users have different tolerances for delay and these tolerances vary with the application. High-quality voice applications require less than 200 ms roundtrip delays, but delays of up to 600 ms have been shown to be acceptable [13].

Since current packet-switching networks do not provide a bounded delay service, voice protocols must provide mechanisms that can cope with highly variable end-to-end delays. Adjustments of the packetization interval and buffering of voice packets at the receiver are widely used to compensate for unpredictable delays [15,19].

2.3. Delay jitter

If the network delay of voice packets is not constant, e.g., due to statistical multiplexing, discontinuity of the voice playback at the receiver can occur. We refer to these discontinuities as *gaps*. Gaps are commonly addressed through buffering at the receiving site. The first packet in a talkspurt is artificially delayed at the receiver for a period of time known as the *control time*. The control time builds up a buffer of arriving packets sufficient to provide continuous playback in the presence of delay jitter. Note however, that the control time cannot be arbitrarily large due to constraints on the roundtrip delay.

The use of a control time to compensate for delay jitter requires mechanisms to identify the beginning of talkspurts and to determine the control time. The latter is difficult since it requires knowledge of the network delay distribution. Numerous methods have been proposed for estimating the control time of a talkspurt, based on network delay measurements [15], on stochastic assumptions about the network delay [1,2], or both [16].

2.4. Error control

The impact of packet loss on voice quality varies since interpolation can mitigate the effects of lost samples and not all samples contain equally important information. In any case, the tolerance to packet loss is low and even the loss of even a single packet may be perceptible during playback.

Packet-level error control for packet voice streams must be designed so as to provide the best possible quality for the stream. Conventional error control techniques are unacceptable since they do not consider the delay sensitivity of voice data. Hence, researchers have dismissed a retransmission-based approach to error control, focusing instead on open-loop techniques that recover or limit the effects of losses.

Forward error correction (FEC) [3,17] provides robustness in the presence of packet loss by adding redundant information to the original samples. If only a small number of packets is lost, the added redundancy enables a reconstruction of the original voice data at the receiver. The ability to recover lost information strongly depends on the degree of redundancy. In addition to considerable processing overhead, FEC-based error control results in increased network bandwidth consumption. Thus, FEC contributes to network congestion, and, since losses in the network are most often due to congestion, may even be the cause of packet loss.¹

Channel coding refers to a class of approaches that separate the voice signal into multiple data streams with different priorities. The priorities are used by network switches to selectively discard low priority packets which carry information that is less crucial in reconstructing the voice signal. Channel coding techniques have been shown to provide a graceful degradation of playback quality in a variety of loss scenarios [10,20,23]. For PCM-encoded voice, packet loss rates of over 5% on the channel carrying the least significant information have been reported as tolerable when using small (32-byte) packets [21]. A drawback to channel coding, however, is that the network is required to support selective discarding of packets during periods of congestion.

3. An analytical model for retransmission of packet voice data

In this section we present an analytical retransmission model for error control of a voice packet stream in a packet-switching network. Through analysis of the model we are able to quantify the gain achieved by the retransmission of packets lost in the network. Given an encoding scheme, voice quality depends primarily on maintaining the continuity of the playback of each talkspurt at the receiver. The loss of voice quality results from discontinuities due to delay jitter or packet loss. Since they both cause gaps, delay jitter and packet loss cannot be considered separately. Timely retransmission is of little value if discontinuities due to delay jitter degrade the quality of the playback. For this reason we define the performance metric for measuring transmission quality as the probability of continuous playback, i.e., a playback without gaps, of an entire talkspurt. Under a given packet loss scenario, this metric accounts for the quality degradation due to delay jitter *and* to untimely retransmission. The quality degradation due to delay jitter alone is found by computing the metric under the assumption of error-free transmission.

Our model considers all protocol issues reviewed in the previous section. Since packets in different talkspurts rarely interfere with each other, we model the transmission of a single talkspurt within a packet voice stream. The sender introduces packets into the network with a deterministic spacing as given by the packetization interval. It is assumed that the network preserves the transmission order of the packets in the voice stream. Since connection-oriented networks, such as ATM networks, guarantee in-order delivery, and, in practice, most connectionless networks rarely reorder the packets in a stream communication,² the assumption of in-sequence delivery is widely applicable.

In Section 3.1 we give a detailed description of an end-to-end retransmission model. In Section 3.2 we develop analytic expressions for the effectiveness of retransmission of lost voice data. In Section 3.3 we present numerical examples of our results.

¹ Note that retransmission-based error control also consumes network bandwidth and increases network congestion. Unlike retransmission, however, forward error correction always imposes additional overhead on the network and thus may induce congestion-based losses in the network where none would otherwise occur.

² In connectionless networks a traffic stream such as a voice transmission often passes over the same set of routers. In LANs and MANs, it is often the case that only a single physical path exists between a pair of communicating endsystems. Even if there are multiple network paths between the sender and the receiver, current routing algorithms modify their routing tables relatively infrequently, resulting in infrequent route changes. One study on wide-area Internet traffic, for example, reports packet reorderings to have occurred for less than 0.05 % of the packets transmitted in its experiments [18].

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.