


encoded trace of the Star Wars movie [3] ($B = 300$ kbit.)



For the three scenarios, we compare the bandwidth per stream $\frac{c}{n}$ as a function of c , where n is the maximum number of streams that can be accommodated for zero loss. (See figure.) Our scheme achieves slightly less SMG than the unrestricted case, because buffers are not shared. Nevertheless, we are able to extract most of the gain, especially for high service rates.

5 Discussion

RCBR solves the two major problems raised in Section 2 by a) doing away with one-shot traffic descriptors for VBR sources and b) dealing explicitly with sustained peaks in the source rate. Two additional advantages of RCBR are that it is easy to implement, since CBR service is well understood and we are merely adding a fast renegotiation component to it. By building on CBR service, we make the network design simple, since neither complex scheduling disciplines nor large network buffers are required for CBR traffic.

It has recently come to our notice that Chong et al [2] have independently arrived at conclusions similar to ours, that is, compressed VBR traffic has sustained peaks, and thus is best carried by a renegotiated service. They have concentrated in their work on coming up with good predictors for rate changes. In contrast, our work has used large deviations analysis for insight into the nature of statistical multiplexing gain and is coupled with a systems perspective.

References

1. P.E. Boyer and D.P. Tranchier. A reservation principle with applications to the ATM traffic. *Computer Networks and ISDN Systems*, 24:321–334, 1992.
2. S. Chong, S.Q. Li, and J. Ghosh. Predictive dynamic bandwidth allocation for efficient transport of real-time VBR video over ATM. *IEEE Journal on Selected Areas of Communications*, 13:12–23, January 1995.
3. M. W. Garrett and Walter Willinger. Analysis, Modeling and Generation of Self-Similar VBR Video Traffic. In *ACM Sigcomm '94*, pages 269–280, University College London, London, UK, August 1994.
4. M. Grossglauser, S. Keshav, and D. Tse. A novel variable bit rate service for integrated services networks. *submitted to the ACM Sigcomm'95*.
5. E. P. Rathgeb. Policing of Realistic VBR Video Traffic in an ATM Network. *International Journal of Digital and Analog Communications Systems*, 6:213–226, 1993.
6. D. Tse, R. Gallager, and J. Tsitsiklis. Statistical multiplexing of multiple time-scale Markov streams. *Preprint*.

for traffics such as compressed video, because the slow time-scale is significantly longer than the delay requirement. On the other hand, when a large number of independent multiple time-scale streams are multiplexed together, much more gain can be obtained beyond looking at the equivalent bandwidth of each stream in isolation. This gain is due to the fact that with high probability, not too many sources can be in a high-rate sub-chain at any one time. Thus, the bulk of the gain is obtained through averaging between sources with respect to the slow time-scale dynamics rather than through smoothing by the buffer.

4 A New Scheme for VBR: RCBR

RCBR augments standard CBR service with a fast renegotiation mechanism. Each stream has a dedicated access buffer to absorb the fast time-scale dynamics, and at any given time, it is given a CBR rate. The source renegotiates a new CBR rate whenever there is a slow time-scale transition. (A signaling mechanism for this purpose is described in [1].) In this way, the slow time-scale statistical multiplexing gain described above can be captured, while providing more protection than direct multiplexing without restriction. For stored data stream applications, we have developed an algorithm to compute the renegotiation schedule in advance. For interactive applications, we propose that an active component monitor the user-network buffer and initiate renegotiations based on the buffer occupancy level. While we cannot present details for lack of space, analysis shows that with careful coding and cleanly designed end-systems and switch controllers, renegotiation overheads are small.

Due to the inherent uncertainty in the bandwidth requirements of calls over their lifetime, there is always a possibility that a renegotiation for a higher bandwidth can fail due to unavailable capacity. The probability of renegotiation failure is kept low by admission control, such that there is always excess capacity in the system. The amount of slack can be determined by estimating the renegotiation failure probability using large deviations techniques. For users who do not want to take any risk, they might reserve resources at or close to the peak rate. Note that even if the renegotiation fails, the source can keep whatever bandwidth it already has. For interactive applications, the signaling system can also ask the user or application to temporarily reduce its data rate, such as using an adaptive codec. Thus, we believe that there are a number of techniques we can call upon to deal with renegotiation failures. For more details of the analysis, refer to [4].

We have compared the statistical multiplexing gain (SMG) achievable through our scheme with two other scenarios. The first scenario multiplexes n streams without any restriction on a server with rate c and buffer size nB . This is used to determine the maximum achievable SMG for the given sources. The second scenario has a leaky bucket regulator policing the stream entering the network. The average rate a policed by the leaky bucket is chosen so that the maximum backlog is B . The third scenario represents RCBR, with a buffer of size B per stream. The streams we have used are n random shifted versions of an MPEG-1

a data stream to a descriptor. The situation becomes even worse for interactive applications, where the data stream is not known in advance.

The single largest application using VBR service is expected to be compressed video. The second problem with current VBR service is that typical compressed video traffic does not match the moderately bursty traffic model envisaged by designers of VBR service. It has been observed by several researchers [3, 5] that, independent of the coding algorithm, there are fairly long durations, when the data rate of the video source is continuously near the peak rate. For such traffic, if a leaky-bucket like descriptor is used, one is faced with a series of poor choices. Assume for the moment that the token rate is chosen close to the average rate in order to maximize the statistical multiplexing gain. Now, if the token bucket is chosen to be small, then during sustained peaks, either the data buffer at the regulator has to be very large, or there will be many losses. If the loss rate is to be small, and data buffers are made large, this leads to expensive regulators and long delays for the sources. On the other hand, if the token bucket is made large enough to rapidly drain bursts, then the network and receiver will need large data buffers to prevent cell loss during a burst. Further, even a compliant source has considerable freedom to disrupt the network core by sending in data in very large bursts (on the order of tens of Megabytes). We call this loss of *protection*, since other endpoints would be unprotected from bursts from a compliant but ill-conditioned source.

Thus, the phenomenon of sustained peaks leads either to a) loss of smoothing gain or b) large loss rate or c) large delays and expensive regulators or d) loss of protection. Given the current framework, there is no way to avoid all four problems simultaneously. This is a simple consequence of the fact that the sustained peaks in workload violate the design assumptions for VBR service.

3 Statistical Multiplexing Multiple Time-Scale Streams

Recent work provides the theoretical basis for understanding the gain achievable by multiplexing traffic sources exhibiting sustained peaks [6]. In this work, each variable-rate stream is modeled as a process modulated by a multiple time-scale Markov chain: a chain which consists of several sub-chains between which the transitions have very small transition probability. The dynamics within each sub-chain model fast time-scale behavior (such as correlations between adjacent frames) while the transitions between the sub-chains model slow time-scale behavior (such as scene changes). The sustained peak observed by several researchers corresponds to remaining in a high-rate sub-chain for a long time in this multiple time-scale model.

There are several key results of this work, based on large deviations theory. First, when one computes the *equivalent bandwidth* of an individual multiple time-scale stream, it is found that one has to allocate the maximum of the equivalent bandwidths of the fast-sub-chains. This essentially means that one has to allocate a rate near the sustained peak, and it underscores the fact that the statistical multiplexing gain due to smoothing using buffers is of limited use

The Case Against Variable Bit Rate Service

M. Grossglauser, S. Keshav and D. Tse

AT&T Bell Laboratories, Murray Hill, NJ 07974, USA

Abstract. There are two major problems with current designs of VBR service for compressed video traffic: a) it is inherently hard to describe the traffic and b) sustained peaks in the source rate substantially degrade the performance. We propose a new scheme, Renegotiated CBR (RCBR), that augments CBR service with fast renegotiation and buffer monitoring. Analysis shows that the scheme is easily implementable and solves the problems above. RCBR works well because it makes intelligent use of the presence of multiple time-scales in the traffic.

1 Introduction

Variable Bit Rate (VBR) service has been designed in *anticipation* of future traffic, particularly compressed video traffic. Such traffic is modeled as having an intrinsic long-term average rate, but with periods in which data is generated in a burst at some peak rate. Current designs essentially augment a standard CBR service with the ability to admit limited bursts into the network. The hope is that by admitting bursts users would get smaller delays. At the same time, if the burst size is limited by an intelligent admission control scheme and policing, limited buffering in the network would be sufficient to give strong guarantees on delays and losses. Note that it is important to be able to characterize both the worst case burst size and the long term average rate for the scheme to work well.

We argue that this design for VBR service suffers from two major flaws (Section 2). In fact, a different modeling of VBR traffic leads to some insight into a better service definition (Section 3). This motivates Renegotiated CBR (RCBR) service, where the network core is CBR, but rapid renegotiation of CBR rates allows the network operator to extract almost all of the statistical multiplexing gain inherent in the traffic (Section 4). Analysis and experiments indicate that RCBR is stable, efficient and has low overhead. We believe that our approach is successful because it correctly models compressed video traffic in terms of multiple time-scales.

2 Problems with Current VBR Service

There are two major problems with current service definitions. First, many researchers assume that users of VBR service can describe their traffic using a small number of descriptors, such as the token bucket size and token drain rate for the leaky bucket scheme. We have found that it is hard to determine these parameters even for a stored data stream, since there isn't a unique mapping from