QUALITY TRADEOFFS FOR PACKET VIDEO OVER DIFFERENTIATED SERVICES NETWORKS

James Macnicol, Mark Pickering, Michael Frater, John Arnold

School of Information Technology & Electrical Engineering
University College, the University of New South Wales
Australian Defence Force Academy, Canberra

ABSTRACT

Differentiated Services networks can only function effectively if an appropriate mix of traffic can be maintained in all service classes using a practical pricing model. Previous work has noted that the same video sequence being distributed to clients who are prepared to pay different amounts for the service results in traffic distributions that vary considerably between those clients. It is shown that this variation can be reduced by trading off coding distortion and the amount of protection given to packets. As a result, a simple pricing model provides stable traffic distributions over a wide range of network conditions and client budgets.

1. INTRODUCTION

One of the fundamental limitations of the Internet today is its uniform treatment of all traffic that passes through it, regardless of the type of application involved or the impact of congestion on the user experience. This is especially relevant to video streaming where systems tend to perform poorly when network throughput varies with time or when there is excessive loss or delay. As such, significant impetus exists for moving towards a system where applications can trade-off quality of service against cost. This would provide good performance for sensitive applications such as video while lower priority applications, e.g. e-mail, will be disproportionately affected when network conditions are poor.

The Differentiated Services (DiffServ) model, in either relative or proportional form [1], is attractive from an implementation perspective since it does not require per-flow processing in routers, a requirement for other approaches such as IntServ [2]. In the DiffServ model packets are tagged as belonging to one of n different *service classes* with per-byte or per-packet costs $P_0 \dots P_{n-1}$. The network guarantees that the level of service provided to a given class is (at least) no worse than the preceding class with lower cost.

A number of earlier works have considered transmission of packet video over DiffServ networks. We believe that many of these works have used impractical pricing models. For example, [3, 4] who both use the same parameters as the earlier work by Seghal et al [5], use prices P_i that are widely spaced but with corresponding packet loss rates $l_0 \dots l_{n-1}$ which do not vary to a similar degree. Indeed, cases exist where some service classes are defined which would never be used since better performance would be achieved by sending the packet multiple times in a lower service class. This is due to the (invalid) assumption that there should always be a one-to-one mapping between video packets and network packets. A second issue with these papers is that they consider only static price/loss combinations (P_i, l_i) whereas real networks have loss rates that change much more quickly than prices can [6]. Our previous work [7] has shown that given an appropriate dependence between loss rates in the different service classes, the distribution of traffic between classes can be approximately constant as the network load changes. One problem highlighted was that when the same coded video sequence was sent to multiple clients who are prepared to pay different amounts for the service that the “richer” clients would not generate as much traffic in the lower service classes. Starvation of traffic in those classes could cause the system to degenerate to performance identical to the best-effort case. In this paper we examine systems where the no-loss coding distortion (and hence the video data rate) is allowed to vary between users with different budgets and its effect on the overall traffic distribution when users with different budgets are present.

The remainder of the paper is structured as follows: Section 2 describes our model of the utility maximization process that users undertake for packet video over DiffServ. Section 3 outlines the previous results from [7] when the same video packet rate is used regardless of the user’s budget. Section 4 describes our new results when the video data rate is allowed to vary with budget and network conditions. Final conclusions are drawn in Section 5.

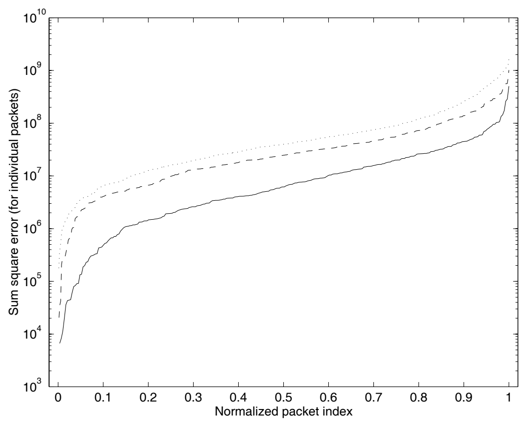

Fig. 1. Packet loss impact for three MPEG-4-coded video sequences (100 frames each)

2. VIDEO QUALITY OPTIMIZATION

Here, as in [7], we consider only the two service class case with per-packet prices (P_0, P_1) which we can consider static and loss rates $(l_0(t), l_1(t))$ which vary with time. With only two classes the results are easier to visualize; generalization to three (or more) classes is straightforward.

Given a video sequence composed of N_v video packets and an overall budget B, we seek a mapping from video packets to network packets which minimizes the expected sum square error (or, equivalently, which maximizes the PSNR) at the decoder. Packet i is transmitted $\alpha_i(t)$ times using packet(s) from class 0 and $\beta_i(t)$ times in class 1. It is important to stress again that $\alpha_i(t), \beta_i(t)$ can take any integer value, not just 1. This error minimization process can be represented by:

$$E_{total}(t) = Q(t) + \min_{\alpha_i(t),\beta_i(t)} \sum_{i=0}^{N_v-1} l_0(t)^{\alpha_i(t)} l_1(t)^{\beta_i(t)} S_i \qquad (1)$$

subject to the total cost constraint:

$$\sum_{i=0}^{N_v-1} \alpha_i(t)P_0 + \beta_i(t)P_1 \leq B \qquad (2)$$

where $Q(t)$ is the no-loss coding distortion (e.g. due to transform coefficient quantization by the video encoder) and S_i is the sum square error resulting from the loss of video packet i. These two components, one for coding distortion and the other for packet loss are assumed to be uncorrelated and therefore whose contributions can be summed together.

Figure 1 plots the *packet loss impact* measure S_i on a logarithmic scale for three different video sequences coding using MPEG-4 (with decoder error concealment), sorted in ascending order. Given this very uneven distribution of

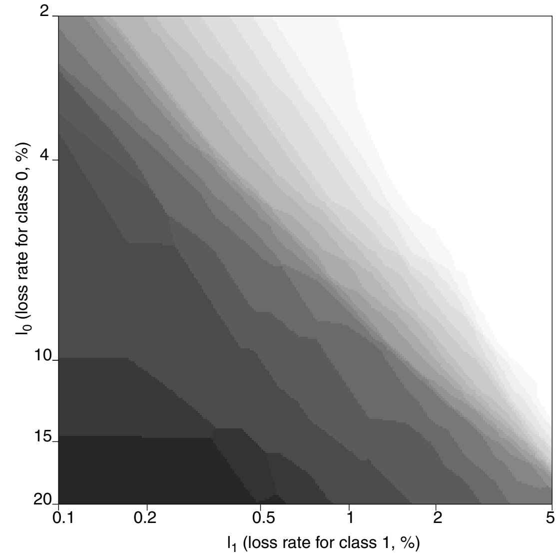

Fig. 2. Traffic distribution ρ for varying losses (l_0, l_1), maximizing quality at each sample point

packet "importance" we can see that packets to the left in Figure 1 will be allocated a single packet in the cheapest class (or perhaps dropped entirely) with more important packets being allocated a single packet from the expensive class. The most important packets may be allocated multiple network packets to minimize the possibility of significant errors being caused by the loss of such packets.

Exhaustive searches are used to determine the values for $(\alpha_i(t), \beta_i(t))$ used in the experimental results. We do not consider fast methods for computing these, but rather we are interested in the behaviour of these systems in the "optimal" case.

3. PREVIOUS RESULTS (CONSTANT VIDEO DATA RATE)

Figure 2 shows the result of minimizing (1) for a single 2s GOP from the standard "Mother & Daughter" sequence coded with MPEG-4 at 120kbps (69 packets) with prices P_0=0.6, P_1=1 and a constant budget B=60 over a range of loss rates (l_0, l_1). Plotted here is the quantity ρ which is defined as the proportion of the traffic which uses service class 0:

$$\rho = \frac{\sum_0^{N_v-1} \alpha_i}{\sum_0^{n_v-1}(\alpha_i + \beta_i)} \qquad (3)$$

The luminance values in Figure 2 are proportional to ρ i.e. black represents ρ=0 (no packets generated in class 0, all in class 1 which occurs when l_0 is high and l_1 is low) and white ρ=1. Note that for a fixed budget the overall packet rate is proportional to ρ since spending more of the budget on packets from the expensive class buys fewer packets

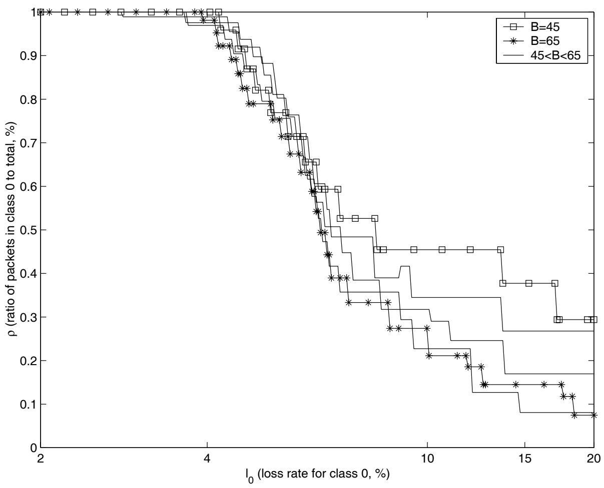

Fig. 3. ρ vs. budget using same video bit rate

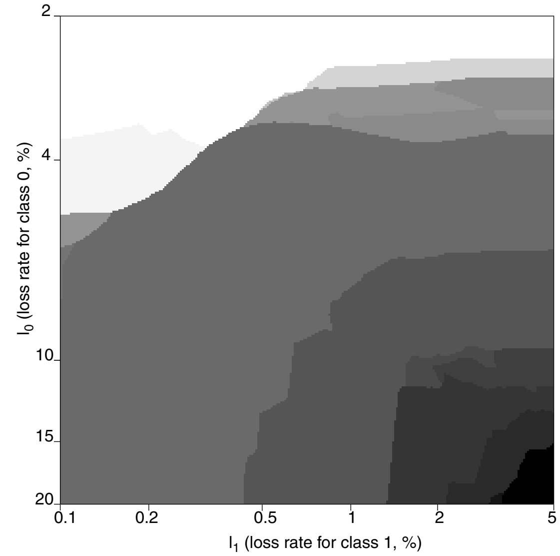

Fig. 4. Optimal video bit rate for constant budget

overall. For typical networks l_0 and l_1 do not vary independently and ideally we would like any changes in $(l_0(t), l_1(t))$ to *decrease* the value of ρ (or at least for it to remain constant) as network load increases. An increase in ρ implies an increase in the overall sending rate which only worsens network congestion.

A major issue identified in [7] relate to large-scale networks with many users. While each user is charged the same price per packet, different users have different total budgets B which affects the proportion of packets they buy from each class. It is vitally important that under all circumstances there is significant demand for packets in all service classes. Otherwise the DiffServ network behaves identically to a best-effort network.

Re-computing the results in Figure 2 for a higher total budget results in increased demand for high-cost/low-loss packets from class 1 which then decreases the value of ρ. This is illustrated in Figure 3 which plots ρ for various budget levels between B=45 and B=65. Since it is not possible to directly plot ρ as a function of l_0,l_1 and B at the same time, the horizontal axis used in Figure 3 corresponds to the top-right/bottom-left diagonal in Figure 2. Large changes in ρ take place along this path and so we can observe these changes taking place without independently varying l_0 and l_1. It can be clearly seen that as the budget increases, demand for packets in class 1 rises significantly. Many users with such a high budget causes degraded performance for all due to starvation of traffic in class 0.

4. RESULTS WITH VARIABLE VIDEO DATA RATE

The approach in [7] that was summarized in the previous section treats the term $Q(t)$ in (2) as constant over time and only adjusts the allocation of video packets to network packets when dealing with changes in budget or network conditions. If a user's budget is high and/or the network has low loss rates the summation term in (2) is much smaller than $Q(t)$. In such cases it may be advantageous to increase the video data rate (e.g. by using a finer quantizer on transform coefficients) while decreasing the amount of "protection" afforded to those packets. Conversely, for high loss rate and/or low budgets the video packet rate should be dropped to allow for greater protection from network losses.

Firstly, to investigate the value of adjusting the video data rate under changing network loss conditions (but constant budget), the same video segment as used in Section 3 was coded at different rates from 70kbps through 140kbps (in steps of 2.5kbps) and the optimal video rate was found for each combination of (l_0, l_1). The luminance of Figure 4 is proportional to the optimal rate calculated from this experiment (black is 75kbps, white is 135kbps). As expected, the video rate is low for very high losses (bottom right corner) and high for low loss conditions (top left corner).

A second set of tests was conducted to demonstrate a similar dependence between video packet rate and total budget. Figure 5 plots the optimal video rate as both total budget and loss varies. The horizontal axis here represents the top-left/bottom-right diagonal of Figure 4 with both l_0 and l_1 increasing towards the right and shows the same behaviour as seen in that Figure when the budget is constant: the optimal video rate decreases as the overall loss increases. As the budget is increased we can see that the optimal video rate also increases which confirms the earlier hypothesis that any additional network packets purchased with a greater budget should be used to send new video packets rather than providing extra protection to existing ones.

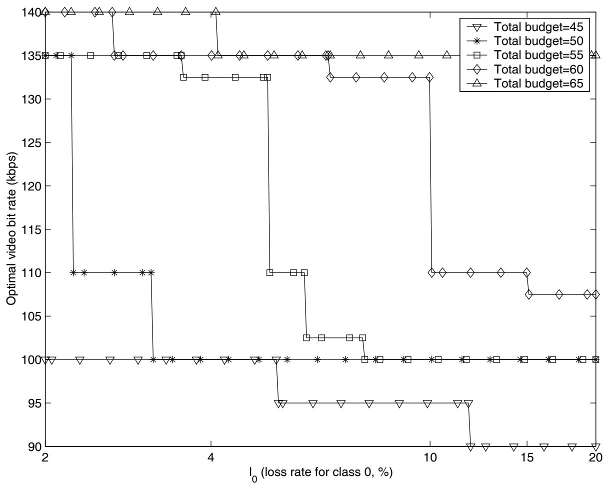

Fig. 5. Optimal video bit rates for different budgets

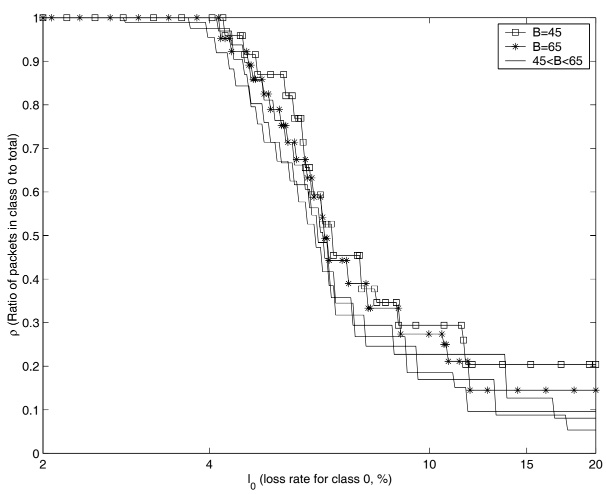

Fig. 6. ρ vs. budget using optimal rates

Finally, Figure 6 reproduces Figure 3 except in this case the video rate is allowed to also change. We can see that the distribution of traffic between classes is now much less sensitive to the overall budget. Indeed, there is no consistent trend in the variation of ρ as the budget is increased. The volatility shown here (which is also visible in Figures 4, 5 and to a lesser extent Figure 2) is likely due to the fact that the search for the optimal video quality is done over a relative small set of video packets taken from a single GOP (using a larger sample would require an impractical amount of computation for the exhaustive search). Future work should investigate the modelling of the packet loss impact (e.g. using an exponential distribution) as an alternative to small empirically-determined samples. This is also likely to be the basis for practical techniques that do not require excessive computational resources.

5. CONCLUSION

Unless an appropriate balance of traffic is maintained between the various traffic classes a DiffServ network regresses to behaving identically to a best-effort network. It is important to ensure that under a variety of network conditions utility-maximizing users always maintains this balance. Our previous work indicated that if a single coded video sequence was used regardless of network conditions or budget that users would choose different traffic distributions. This work has shown, however, that if the video rate is allowed to vary in response to these factors, so that the effects of quantization error and packet loss error can be allowed to balance out, then these differences are significantly reduced. From this result we predict that a DiffServ network carrying video traffic can function successfully under a simple, static pricing regime.

6. REFERENCES

[1] Y. Chen, C. Qiao, M. Hamdi, and D. H. K. Tsang, "Proportional Differentiation: A Scalable QoS Approach," *IEEE Communications Magazine*, vol. 41, no. 6, pp. 52–58, June 2003.

[2] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: an Overview," RFC 1633, June 1994.

[3] F. Zhai, C. E. Luna, Y. Eisenberg, T. N. Pappas, R. Berry, and A. K. Katsaggelos, "A Novel Cost-Distortion Optimization Framework for Video Streaming Over Differentiated Services Networks," in *Proc. ICIP*, Barcelona, Spain, September 2003.

[4] F. Zhang, M. R. Pickering, M. R. Frater, and J. F. Arnold, "Optimal QoS Mapping for Streaming Video Over Differentiated Services Networks," in *Proc. ICASSP*, Hong Kong, April 2003, vol. V, pp. 744–747.

[5] A. Sehgal and P. Chou, "Cost-Distortion Optimized Streaming Media Over DiffServ Networks," in *Proc. ICME*, Lausanne, Switzerland, August 2002, pp. 857–860.

[6] J. Altmann, H. Daanen, H. Oliver, and A. Sánchez-Beato Suarez, "How to Market-Manage a QoS Network," in *Proc. INFOCOM*, New York, June 2002, pp. 284–293.

[7] J. Macnicol, M. Pickering, M. Frater, and J. Arnold, "Pricing of Packet Video For Relative Differentiated Services Networks," in *Australian Telecommunications, Networks and Applications Conference*, Melbourne, December 2003.