Proceedings of the First Polish-German Teletraffic Symposium (PGTS2000), Dresden, Septempber 2000.

ON AGGREGATION STRATEGIES FOR MULTIMEDIA TRAFFIC

Klaus Dolzer, Wolfgang Payer
University of Stuttgart, Institute of Communication Networks and Computer Engineering
Pfaffenwaldring 47, 70569 Stuttgart, Germany
e-mail: {dolzer, payer}@ind.uni-stuttgart.de

Abstract

An actual trend in networking is the aggregation of traffic streams. This allows not only to keep the core equipment simple – while increasing the complexity at the network edge – but also to overcome scalability issues in the core network. However, traffic flows within aggregates are not distinguishable any more and therefore experience all fairly the same service quality which can be harmful in multi-service networks. In this paper several different strategies for traffic aggregation are compared in an enterprise network scenario. In order to apply realistic traffic sources, we carried out measurements of multimedia traffic streams, the results of which are presented and discussed in detail with respect to different traffic characteristics. Finally, we performed simulations on several different aggregation strategies in order to show that it is possible to satisfy the individual QoS requirements of all real-time flows while maintaining an overall high load – even with only two classes.

Keywords

Flow Aggregation, Traffic Characterization, Multimedia, Measurement, Simulation

1 INTRODUCTION

Due to scalability issues in architectures which originally did not support traffic aggregation like IntServ [2][13], the community started to put a lot of effort in defining sophisticated aggregation strategies which are either integrated in established architectures or which are in consideration for the design of new architectures. IP tunnels within IntServ [10] is an example for existing architectures while the behaviour aggregates concept in DiffServ [1] and the label stack concept in MPLS [12] are examples for arising architectures. Unfortunately this architectures provide mainly protocols to support the traffic aggregation but questions like ‘which traffic types should be aggregated?’ and ‘how much bandwidth should be reserved for an aggregate in order to obtain a sufficient QoS?’ are degrees of freedom. With a quickly growing number of applications that are based on different more or less stringent QoS requirements, the decision of which flows should be aggregated and which should be kept separate gets increasingly difficult.

From the point of view of teletraffic theory, the smallest coefficient of variation of the waiting time can be achieved if only flows with the same traffic characteristics are aggregated. Because of the many different applications with specific QoS

This work was partly sponsored by Siemens AG, Germany.

requirements, the resulting number of classes would still be too large. Thus a trade off between a small number of classes, allowing a high multiplex gain and a simple network but without isolation within the classes, and a high number of classes, offering good isolation between flows but also causing scalability problems in the core network, has to be found (see [4]).

The remainder of this paper is organized as follows: In section 2 the simulation scenario as well as the traffic types that are likely to occur in such a scenario are introduced. A detailed discussion of the measured traces of interactive video with respect to frame length distributions, correlation effects and temporal behaviour is given in section 3. Finally, in section 4, different aggregation strategies are introduced and compared with respect to their waiting time distribution functions as an important measure for QoS of multimedia traffic.

2 SYSTEM AND TRAFFIC MODEL

Fig. 2.1 depicts an enterprise network which is formed by LANs that are interconnected over a QoS-supporting WAN, which we expect not to essentially affect traffic characteristics. As we assume the traffic between the LANs within this network to be symmetric, we only consider the outgoing traffic of one LAN.

In our scenario all flows which belong to the same aggregate are forwarded into one unbounded FIFO queue. The classification of the packets is done according to the corresponding aggregation strategies which are introduced in section 4. The isolation between the aggregates is achieved with an SCFQ scheduler [5], a GPS (generalized processor sharing) approximating service discipline with low processing power requirements which is widely applied (e. g. [6]).

As our main focus is on the evaluation of realistic multimedia traffic sources in a realistic scenario, we assume to find a mix of audio, interactive and streaming video, and data traffic. The audio traffic is not only generated by IP telephony, but likewise by video conferences, which also produce the interactive video traffic (H.261 or H.263 coded). The streaming video traffic (MPEG coded) might come from (net) news, distant learning or online multimedia documents and databases. Finally, the data traffic represents e-mails, FTP and web access.

We modelled the audio traffic (G.711 and G.723) as CBR sources (with mean rates of 84.2 kbps and 21.8 kbps respectively).

In order to take the not negligible autocorrelation between successive video frames into account, we used traces of interactive and streaming video traffic as sources to our system. Concerning interactive video traffic, we carried out own measurements in our network lab in order to obtain detailed traces (see also [9]). A discussion of these traces is part of section 3. Considering, that MPEG traffic is already exhaustively examined, we used the well-known MPEG traces from O. Rose [8] as streaming video sources.

As the focus of the presented simulation study is on the obtained QoS of real-time traffic, the data traffic is modelled as worst-case background traffic. Therefore we used M/Pareto traffic ($\alpha = 1.2$, mean file size 3 KB, packet length 576 and 1500 Byte respectively) to represent WWW get requests and responses, e-mail and FTP, and sources with poissonian arrival to represent acknowledges (e. g. [3]).

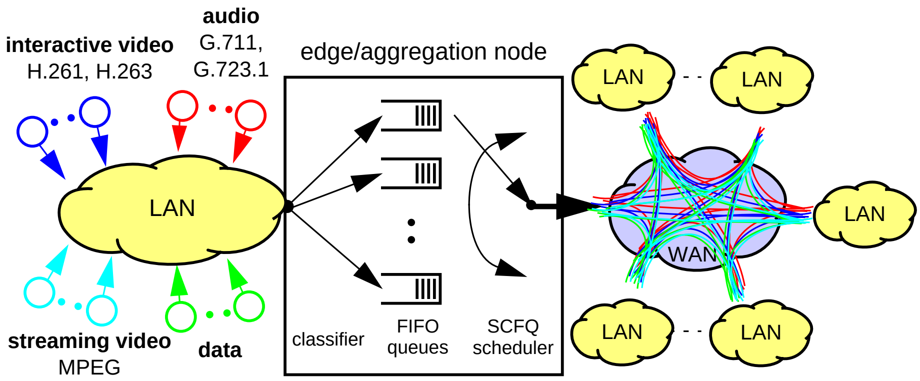

Figure 2.1: Edge/aggregation node in an enterprise network

3 MEASUREMENTS & EVALUATION

As video material for the measurements we decided to use a political talk show because of its large proportion of head and shoulder scenes (typically for video conferences). Played by a video recorder, captured with a framegrabber card and coded, using the videoconferencing tool *vic* (available from [11]), the video was transmitted over an empty 10 Mbps Ethernet segment to another PC. On its way, we captured the complete header of each bypassing packet.

In order to obtain realistic traces, we only changed the settings *codec* (H.261 or H.263), *quality* and *rate control* from the user interface. Considering (visible) *quality*, we chose two specific values to represent high and low quality following a visual comparison of different settings. For the lower quality (10), we once restricted the maximum bandwidth to 384 kbps and once let the codec use as much as it needed (up to 3 Mbps). For the higher quality (3) we chose a restriction to 768 kbps and none as well. Concerning the adjustable frame rate, a setting of 15 frames/s provides a reasonable smooth picture sequence. Thus we obtained eight different traces from the same video sequence (approximately 25 min), which are denoted as <codec>-<quality>-<frame rate>-<max. bandwidth> throughout the remainder of this paper.

In the following the measured traces are analysed according to different statistical characteristics in order to understand their impact on the network, simulated in section 4. Fig. 3.1 and Fig. 3.2 show the probability density function of the frame length for H.261 and H.263 traffic, respectively. It can be seen that the *quality* setting is dominant for its appearance while the different settings of *rate control* do not influence the graphs that much. Both figures also confirm that H.263 is intended to be used on links with smaller capacity and thus produces frames which are in average shorter than the frames generated by H.261. A comparison with the chosen MPEG traces [8] shows that the MPEG traces are also quite different with respect to the frame length distributions, but more similar to the H.26x traces without bandwidth restrictions.

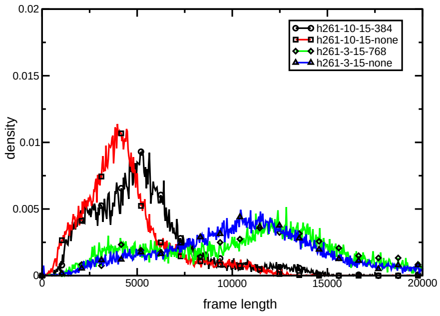

Figure 3.1: Probability density function of the frame length of H.261 frames

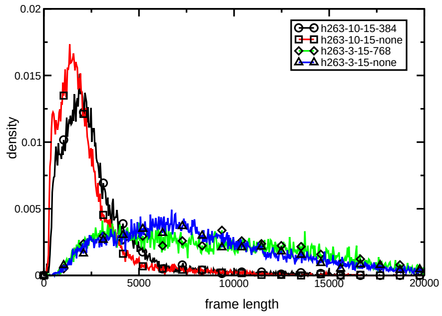

Figure 3.2: Probability density function of the frame length of H.263 frames

As both codecs use inter-picture prediction schemes it is important to have a look at the coefficient of autocorrelation between adjacent frame lengths in order to see the grade and the duration of the dependencies between successive frames. Thus Fig. 3.3 and Fig. 3.4 depict the coefficient of autocorrelation of H.261 and H.263 traffic, respectively, over the lag between two frames. It can be seen that all curves start with a high value and then drop down with different slopes that depend on the applied codec as well as on the restriction of the bandwidth, but – in case of H.261 – only hardly on the selected *quality*. In order to be able to make more concrete propositions, it is necessary to extend the series of measurements to other video material as well as to different settings. But anyhow one can conclude, that in this case the parameter *rate control* seems to be dominant while *quality* is less important. Nevertheless, it can also be stated that the codecs send – if the allowed bandwidth does not restrict them – much more (periodic) information than with restriction. This behaviour can be especially observed in a slope in Fig. 3.3 that is more flat and shows a periodic pattern. Compared to the autocorrelation of the MPEG traces, the obtained curves appear quite different and do not show the same extreme, periodic pattern.

To complete the evaluations of the correlation, we calculated the Hurst parameters from variance-time plots of all measured traces, which we found to be in a range from 0.69 to 0.77. Thus they are below those of the MPEG traces which are usually higher than 0.8. Summarizing, it can be stated, that the measured interactive video traffic has a long-term correlation between successive frames that cannot be neglected, while the traces show quite significant differences to the MPEG traces with respect to frame characteristics.

Besides the above described characteristics, the resulting frame rate is of special interest because its degradation leads to a shaky picture sequence and a change of the interarrival time of successive video frames. We found out that only H.261 traffic without bandwidth restriction is able to meet the frame rate of 15 frames/s[1]. All other settings show a significant degradation of the frame rate in case of bandwidth restriction (down to 3 frames/s). This behaviour can be explained by a few long frames – caused through scene changes or movement –

1. If the frame rate is ongoing underneath the chosen setting, either the available bandwidth or the processing power are generally too low.

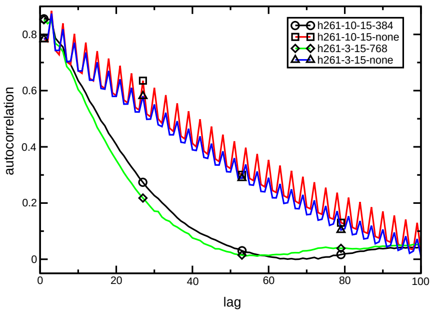

Figure 3.3: Autocorrelation of H.261 traces

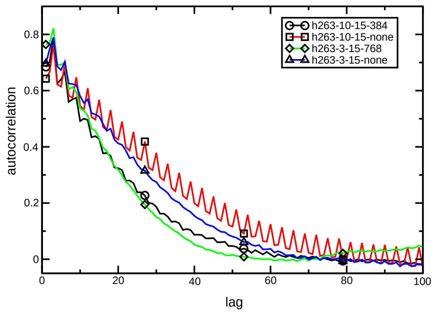

Figure 3.4: Autocorrelation of H.263 traces

that use the whole restricted bandwidth. Even the packets which are formed by partitioning of frames which exceed 1000 Bytes are not sent equally spaced over the link. Thus, it cannot be assumed that the interarrival time between successive frames or packets of the same frame is independent, or even constant.

In order to consider all above mentioned effects of the different settings, we decided to directly use the packet traces for our simulation study, discussed in the next section. In difference to the presented and analysed (video) frame traces, the packet traces are already segmented in packets with a maximum length of around 1000 Bytes.

To cover a broad range of traffic originated by video conferences while still keeping the number of different sources small to be able to interpret the obtained results, we decided to use the traces h261-10-15-none (mean rate 526.4 kbps) and h263-3-15-none (mean rate 901.2 kbps). We did not use the low quality of H.263 because the visual quality was not good enough to be realistic in an enterprise environment. The traces without bandwidth restrictions were chosen to represent an upper bound concerning the coefficient of autocorrelation.

4 Simulation Study

One of the unsolved questions related to traffic aggregation is, how many classes of traffic are really needed and moreover, how much bandwidth has to be allocated to each of these classes in order to obtain a sufficient QoS for the contained traffic. Working on that problem, we carried out simulation studies to compare different aggregation strategies. As an important measure of QoS not only for multimedia traffic, we used the complementary waiting time distribution function (waiting time CDF).

As our main focus is on a network with not too much complexity, we compared the allocation of weights according to the mean bandwidth and according to the effective bandwidth[2] of the respective traffic types (determined by simulations as the probability of 10^{-3} to exceed a maximum delay of 20 ms in the edge node).

2. The effective bandwidth does not stand in opposition to the postulated simplicity of the network, because it is sufficient to use two static curves „number of sources“ versus „factor for rate reservation“ – one for interactive video and one for streaming video – to determine the weight that has to be allocated to the scheduler if the traffic pattern changes. For audio and data traffic, we suggest to allocate mean rates.

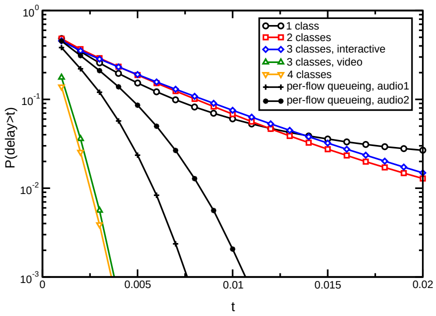

Figure 4.1: Waiting time CDF with mean rate allocation for audio

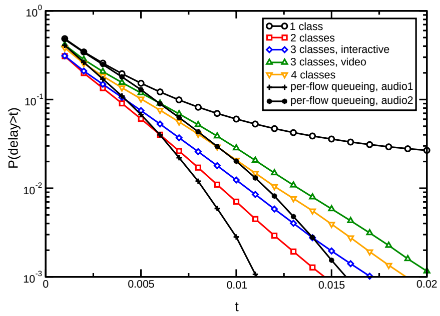

Figure 4.2: Waiting time CDF with effective bandwidth rate allocation

It is appropriate to mention here that the weights allocated to an SCFQ scheduler are of relative nature as the link bandwidth is divided proportionally to the weights of the actually active sources. This is especially important for the bandwidth that is not allocated at all or temporarily not needed because it can be used by other aggregates proportionally to their weights. Furthermore it should be explained that all GPS emulating schedulers have an implicit coupling between reserved bandwidth of a class and its experienced delay because the time between two consecutive grants is proportional to the reserved share of the bandwidth.

Besides an aggregation strategy with *2 classes*, we compare two different possibilities of having *3 classes*, and *4 classes* with the reference curves of total aggregation (*1 class*) and *per-flow queueing*.

In *2 classes* – the often postulated subdivision in a real-time (stream) and a non-real-time (elastic) class (e. g. [7]) – all audio and video traffic is multiplexed in one queue and isolated from data traffic which forms the second class. *3 classes interactive* considers the fact, that the streaming video has not so stringent requirements concerning delay, but on the other hand should be isolated from data traffic as that is based on UDP. Therefore audio and interactive video traffic form one class, streaming video traffic the second class and data traffic the third one. *3 classes video* is an approach to isolate the small audio traffic (first class) from the video traffic (second class) while still being able to obtain a multiplexing gain among all video sources. Finally, *4 classes* is the extreme in our scenario, because one class is used for each of the four different traffic types.

In our simulation scenario[3] we chose a combination of 25 audio sources (G.711, G.723) and 5 video sources (H.261, H.263, MPEG) of each type and 56 combined data sources. This leads to a bandwidth ratio of *audio : interactive video : streaming video : data* of *1 : 3 : 1 : 5*, which we believe to be a reasonable value in the chosen scenario. The sum of the mean bandwidth of all sources is 23.3 Mbps, while the sum of the effective bandwidth is

3. We used a simulation program, modeled and built from components of a simulation library developed at our institute. The playout of each trace was delayed compared to each other during the transient simulation phase in order to avoid synchronization effects on the packet level and to take into account that only one trace is used for five interactive sources, respectively.

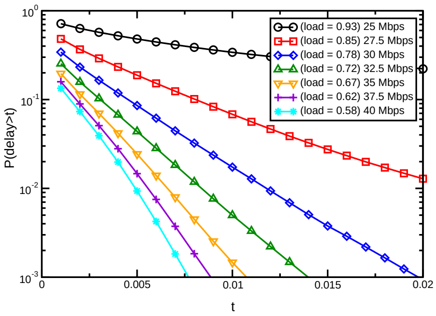

Figure 4.3: Waiting time CDF with mean rate allocation for stream class

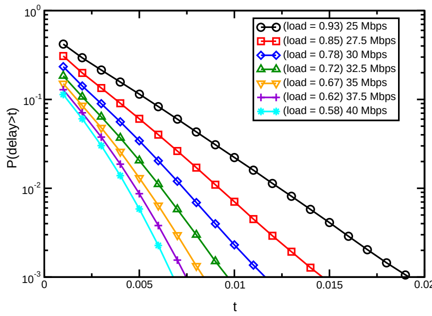

Figure 4.4: Waiting time CDF with effective bandwidth rate allocation

quite higher at 39.9 Mbps. The link bandwidth was chosen to be 27.5 Mbps which equals a carried traffic of 0.85.

Fig. 4.1 and Fig. 4.2 show the complementary waiting time distribution functions of audio traffic, which is certainly most critical because of its stringent requirements in terms of delay and jitter, and its relative small bandwidth. It can be seen, that for mean rate allocation a sufficient service for audio is only possible when kept in a class for itself (*3 classes video* and *4 classes*). This is due to the fact, that in the other cases the video traffic dominates the aggregate and therefore the performance. When allocating according to the effective bandwidth the curves show a rather similar behaviour (with *2 classes* being the best) and the audio traffic keeps the statistical delay bound even for only two classes. It should be mentioned here, that this allocation strategy leads to a higher bandwidth reservation of especially the video traffic, which is substracted from the data traffic and therefore worsens its service as expected.

Figures of interactive and streaming video sources are omitted here due to space reasons. But it can be said that if the reservation is carried out according to the effective bandwidth the curves are similar to the presented curves in Fig. 4.1 while the curves for mean rate allocation are somehow worse. Thus complexity and performance suggest to use only two classes *stream* for real-time traffic and *elastic* for non-real-time traffic.

To get an idea of the stability of the simulation results, Fig. 4.3 and Fig. 4.4 show the complementary waiting time distribution function for the real-time class (*2 classes* aggregation strategy) with mean rate allocation as well as rate allocation according to the effective bandwidth, depending on the link bandwidth (and the carried traffic, respectively). While it is possible for increasing link bandwidth (30 Mbps/carried traffic of 0.78) to keep the statistical delay bound with mean rate allocation as well (overprovisioning), this approach is not viable for a pretty high load on the link. Remarkable is the fast deterioration of the service quality, when the carried traffic exceeds 0.8. For the effective bandwidth allocation a highly loaded link is feasible while still having some reserve for bandwidth fluctuation without worsening the real-time traffic too much.

5 CONCLUSIONS

In the first part of this paper, we presented the evaluation of measured interactive video traffic. Our focus was on the impact of changes in the parameter settings (accessible via the user interface) on the traffic characteristics. We discovered that not only the characteristics of successive frames are heavily dependent on these settings, but also the temporal behaviour of the codecs.

In the second part we used the evaluated traces, together with MPEG traces, audio and data sources in order to obtain measures of the QoS of multimedia traffic. Therefore we compared the results gained from different aggregation and bandwidth allocation strategies. Our simulation results confirm that – in our scenario – two classes provide sufficient QoS support for all applications if the respective rates are allocated according to the effective bandwidth. Furthermore we showed that this statement also holds true if the load on the link is increased.

ACKNOWLEDGEMENT

We would like to thank Thomas Schmidt for his help to carry out the measurements and Stefan Bodamer for the implementation of the evaluation tool as well as many valuable discussions.

REFERENCES

[1] Blake, S.; Black, D.; Carlson, M.; Davies, E.; Wang, Z.; Weiss, W.: *An Architecture for Differentiated Services*, IETF RFC 2475, Dec. 1998.

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

[3] Charzinski, J.: 'Internet Traffic Measurement and Modelling', *Tutorium ICATM'2000*.

[4] Dolzer, K.; Payer, W.; Eberspächer, M.: 'A Simulation Study on Traffic Aggregation in Multi-Service Networks', in *Proceedings of the ICATM'2000*, June 2000.

[5] Golestani, S.: 'A Self-Clocked Fair Queuing Scheme for Broadband Applications', In *Proceedings of the IEEE INFOCOM 1994*, April 1994, pp. 636–46.

[6] Guerin, R.; Li, L.; Nadas, S.; Pan, P.; Peris, V.: 'The Cost of QoS Support in Edge Devices – An Experimental Study', In *Proceedings of INFOCOM'99*, March 1999.

[7] Lindberger, K.: 'Balancing Quality of Service, Pricing and Utilisation in Multiservice Networks with Stream and Elastic Traffic', In *Proceedings of the International Teletraffic Congress 16*, June 1999, pp. 1127-36.

[8] Rose, O.: 'Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems.' *Proceedings of the 20th Annual Conference on Local Computer Networks*, 1995, pp. 197-406. Traces available from ftp-info3.informatik.uni-würzburg.de/pub/MPEG.

[9] Ryu, B.: 'Modeling and Simulation of Broadband Satellite Networks – Part II: Traffic Modeling.' *IEEE Communications Magazine*, July 1999, pp. 48-56.

[10] Terzis, A.; Krawczyk, J.; Wroclawski, J.; Zhang, L.: *RSVP Operation over IP Tunnels*, Internet-Draft, May 1999, work in progress.

[11] vic: available from www-mice.cs.ucl.ac.uk/multimedia/software/vic/.

[12] Viswanathan, A.; Feldman, N.; Wang, Z.; Callon, R.: 'Evolution of Multiprotocol Label Switching.' *IEEE Communiacations Magazine*, Vol. 36, No. 5, Mai 1998, pp. 165-173.

[13] White, P.P.: 'RSVP and Integrated Services in the Internet: A Tutorial.' *IEEE Communications Magazine*, Vol. 35, No. 5, May 1997.