Institute of Communication Networks and Computer Engineering
University of Stuttgart
Pfaffenwaldring 47, D-70569 Stuttgart, Germany
Phone: ++49-711-685-68026, Fax: ++49-711-685-67983
Email: mail@ikr.uni-stuttgart.de, http://www.ikr.uni-stuttgart.de

Proceesings of the IEEE Conference on High Performance Switching and Routing (ATM 2000),
Heidelberg, June 2000, pp. 157-165.

A Simulation Study on Traffic Aggregation in Multi-Service Networks

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

Abstract

In this paper several different strategies for traffic aggregation in enterprise networks formed by LANs that are interconnected via a WAN are compared. We discuss on the one hand the gain in time and space requirements in the WAN dependent on the number of aggregates and on the other hand the experienced waiting time as an important measure of QoS for real-time applications. Our simulation results show that not only the need of resources diminishes if audio and video sources are grouped in one aggregate, but also the waiting time improves significantly.

1 Introduction

During the last years, several trends have been evolving that significantly influence the development of network architectures.

The first trend that can be observed is the quickly increasing number of users communicating worldwide with each other. Since the available bandwidth grows very fast – caused by the emerging WDM technology (wavelength division multiplexing) – the number of (micro-) flows transported in core networks or WANs increases approximately at the same ratio. Considering that the growth of processing capacity of network nodes cannot follow up, network architectures which are in any way aware of (micro-) flows have to support sophisticated concepts of flow aggregation in order not to struggle with scalability issues. The second trend that seriously influences the network architectures is the emerging demand for QoS support of applications (e. g. H.323 [8], SIP [7]). Therefore additional efforts in the network must be undertaken by means of aggregation strategies to provide applications with some grade of QoS. Since a considerable part of the revenue of a multi-service network is achieved by these QoS demanding applications, they play a superior role.

In order to overcome the aforementioned scalability issues, most of the currently applied network architectures provide sophisticated concepts to support traffic aggregation, e. g. the VP concept (virtual path) of ATM [6], the behaviour aggregate concept in DiffServ [1], the label stack concept in MPLS [18] and currently emerging concepts in RSVP (e. g. IP tunnels [17]). It should be pointed out that only the concepts to support traffic aggregation are standardized while the flow types that should be aggregated in order to have an ideal outcome of the aggregation are a degree of freedom.

The simulation studies on different aggregation strategies reported here are performed for a realistic traffic mix. This mix as well as the corresponding traffic models are derived from real multimedia LAN traffic.

The remainder of this paper is organized as follows: Section 2 gives an overview of traffic aggregation strategies and generally discusses their advantages and disadvantages. In section 3 the simulation scenario and the traffic types for the performance comparison of selected aggregation strategies are introduced. Finally in section 4, different traffic aggregation strategies are compared with respect to cost reduction in the WAN as well as the waiting time distribution function as an important measure of QoS.

2 Traffic Aggregation

In this paper we use the term traffic aggregation for the subsummation of different flows sharing the same identifier (may be additional) across a common path in the network. Thus if only this identifier is used to switch traffic through the network, the flows inside an aggregate are not distinguishable any more.

Some well known advantages of aggregation – especially in the case of architectures that keep flow states (like ATM, IntServ/RSVP [2], [19] and MPLS) – are reduced time and space requirements in core nodes

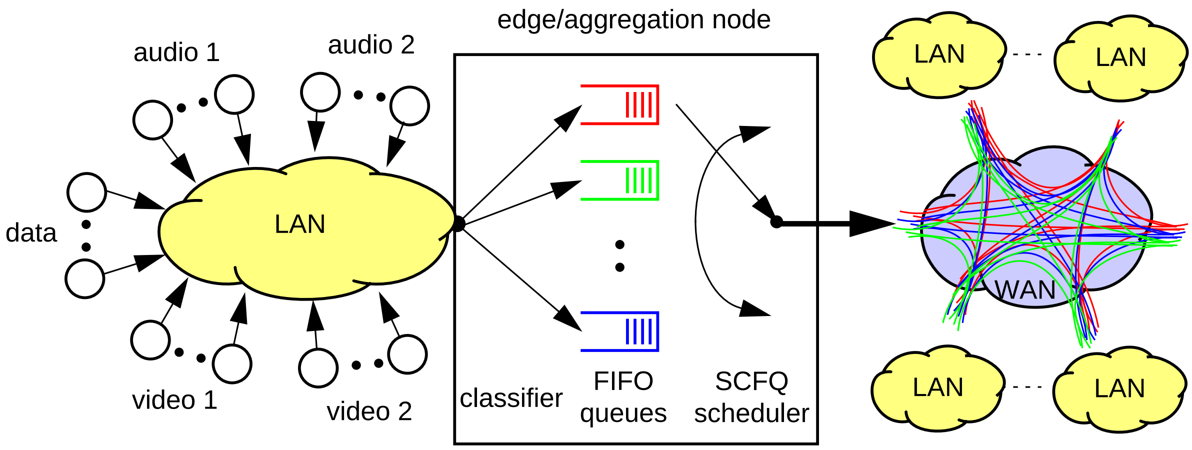

Fig. 1 Edge/aggregation node in an enterprise network

and multiplexing gain for bandwidth and buffer. A detailed discussion can be found in sections 4.1 and 4.2, respectively. On the other hand, aggregation has also some well-known disadvantages which are intensified in multi-service networks: All flows within an aggregate experience roughly the same service. Due to the lack of isolation, bursty sources can steal the bandwidth of well-behaving sources and degrade their service.

From the point of view of queuing theory, the best QoS (under the assumption that the available bandwidth is the same) is achieved if only flows with identical traffic characteristics are aggregated. As there are very few flows in a network with identical traffic characteristics, this approach only theoretically makes sense and has to be extended by some means.

An obvious and earlier applied approach in ATM is to build (service) classes of traffic with similar characteristics (behaviour aggregates). If flows of these classes are aggregated, the question immediately arises where to set limits between flows that should be aggregated and flows that should be separated from each other.

Recent publications (e. g. [13]) suggest to divide the whole traffic mix in only two aggregates, one for real-time traffic (often called stream traffic) and the other one for non-real-time traffic (often called elastic traffic). The immediately arising question is how much bandwidth has to be allocated to the real-time class in order to obtain satisfying QoS for all flows contained in it. Under the assumption that the amount of bandwidth that is used by real-time flows increases compared to the amount of bandwidth used by non-real-time flows, this question gains importance because the often proposed static priority scheduling between the two aggregates potentially entails the risk of (temporary) starvation of the non-real-time aggregate (see e. g. [15]).

3 System and Traffic Model

3.1 Simulation Scenario

Fig. 1 depicts an enterprise network which is formed by LANs that are interconnected over a WAN. Concerning the architecture of the WAN, we only assume it to be QoS-supporting and thus we expect it to add some delay, loss and jitter to the traffic stream, but not to essentially affect the characteristic of the traffic.

If the advantages and disadvantages of aggregation are to be compared in a cost function, one has to distinguish between an edge node and a core node. A widely supported aim is to push the complexity to the edge and keep the core preferably simple while retaining the QoS for the flows unchanged. Thus we evaluate an edge node and also consider the effects in the core network by means of comparing the cost of a respective aggregation strategy and the impact on the QoS behaviour. As we assume a logically fully meshed symmetric network structure it is sufficient to focus on the upstream traffic.

In our scenario all flows which belong to the same aggregate are forwarded in one unbounded FIFO queue. The classification of the packets is done according to the corresponding aggregation strategy, evaluated in section 4.2. The isolation between the aggregates is achieved with an SCFQ scheduler [4], a GPS (generalized processor sharing) approximating service discipline with low processing power requirements which is widely applied (e. g. [5]).

As we consider an enterprise network which integrates voice, video and data, the traffic mix used for the simulation consists of two different kinds of audio and video sources, respectively, corresponding to regularly used codecs and a combined source for the data traffic.

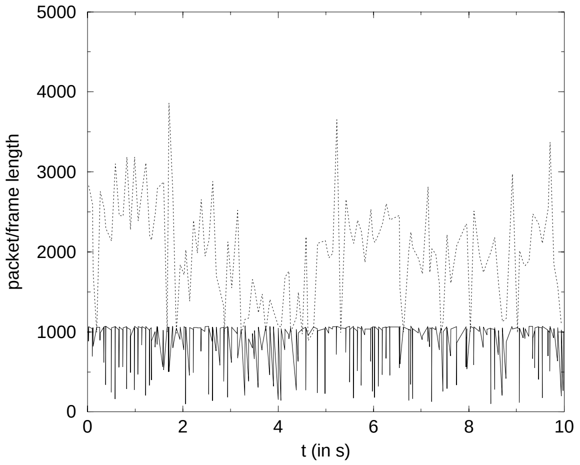

Fig. 2 Trace from Video 1

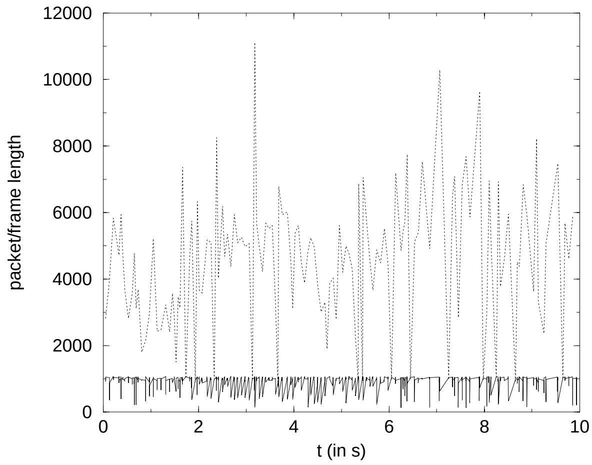

Fig. 3 Trace from Video 2

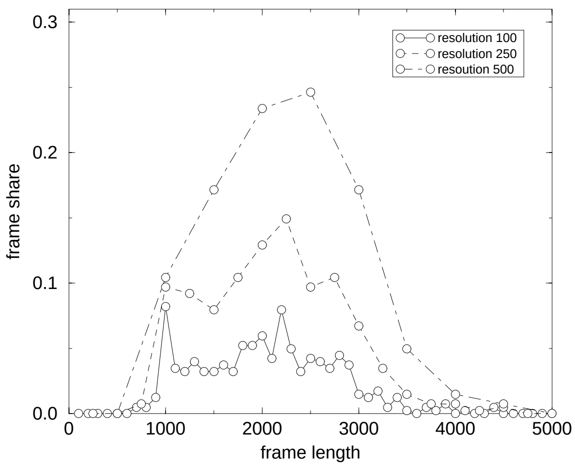

Fig. 4 Frame length probability density function from Video 1

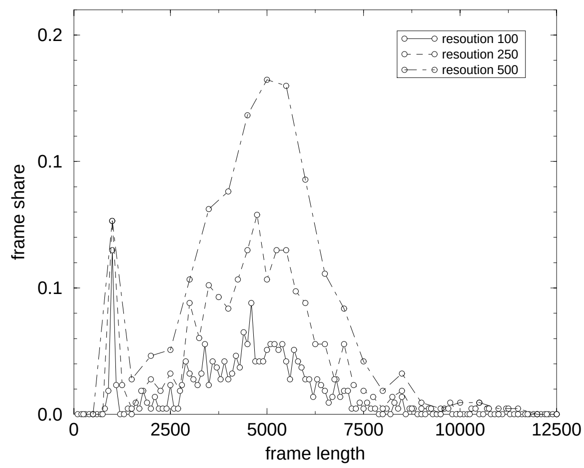

Fig. 5 Frame length probability density function from Video 2

The applied source models are described in more detail in section 3.2. For the simulation studies[1] we will restrict ourselves to a certain traffic mix, consisting of 50 audio sources (25 of each type), 10 video sources (5 of each type) and 40 data sources, leading to a bandwidth ratio of 2:3:6, which we consider reasonable for our scenario.

Among the QoS requirements for interactive multimedia communication, two important parameters are the end-to-end delay and the end-to-end loss probability. A widely accepted maximum value for the end-to-end delay in audio communication (also called one-way system delay, consisting of algorithmic delay, processing delay and communication delay) is 200 ms (e. g. [14]). If we estimate the algorithmic and processing delay of the audio codecs not to exceed 80 ms, we have a maximum of 120 ms for the communication delay. Therefore we assume a maximum delay of 20 ms in the aggregation node to be a reasonable value. Packets exceeding this maximum delay are considered as being lost. An acceptable value for the packet loss depends on the deployed codecs and individual quality demands. Ergo no commonly valid value can be given. In our simulation study we therefore chose the excess probability of the maximum delay to 10^{-3} in the aggregation node to be a good compromise keeping in mind possible additional losses in other network nodes.

1. All simulations are carried out with an object oriented simulation tool that has been developed at our institute.

3.2 Source Models

In our scenario we assume to have a certain amount of interactive audio and video traffic as well as data traffic which is further described in this section. Our main focus is on the evaluation of a realistic traffic mix,

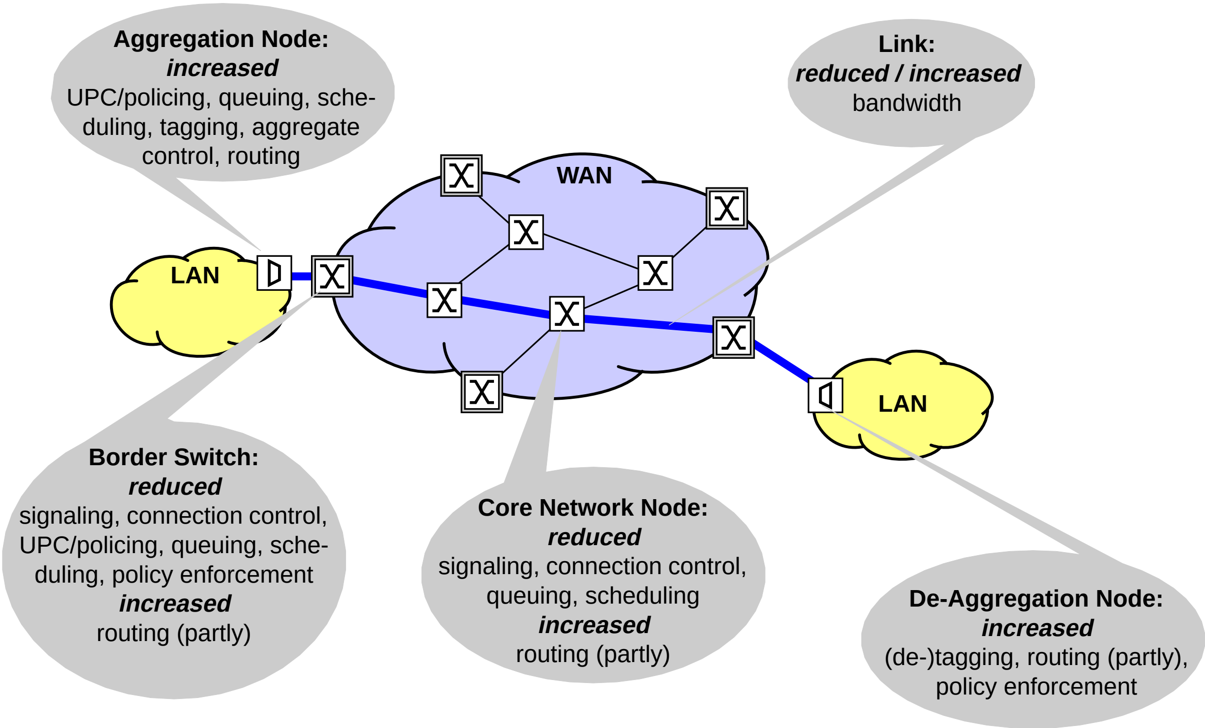

Fig. 6 Influence of aggregation on complexity

thus we use source models and parameters identified from real LAN traffic[2].

Throughout the rest of the paper we will use two different audio sources which are based on the ITU-T recommendation G.711 [9] (Audio 1, mean rate 84.2 kbit/s, frame size 20 ms) and G.723.1 [10] respectively (Audio 2, mean rate 21.845 kbit/s). Silence suppression is not standardized for Audio 1 and not applied for Audio 2. Therefore all audio traffic is modelled as constant bit rate traffic.

Our source model for interactive video is based on the H.261 standard [11] which is commonly applied in video conferencing tools. The frame length distribution[3] is taken from two characteristic settings we found during own measurements where the frame rate was sized to 15 frames/s and the resolution was chosen to be CIF (common intermediate format, for measurements see also [16]). Video 1 (mean rate 240.59 kbit/s) has a just reasonable visual quality for a video conference, whereas Video 2 (mean rate 552.98 kbit/s) offers comfortable visual quality.

2. The source parameters are taken from standards and traffic measurements in our network lab.
3. The video sources generate a frame every t ms which is – if necessary – segmented concerning the maximum transfer unit of the MAC-protocol and the codec characteristics.

Fig. 2 and **Fig. 3**, respectively, show extracts of traced packet/frame lengths of these two characteristic settings. In the depicted case, frames that are longer than 1000 Bytes are segmented in packets with a maximum length of 1000 Bytes.

From these traces we derived the frame length probability density depicted in different resolutions[4] in **Fig. 4** and **Fig. 5**, respectively. It is worth mentioning that through this simple modelling any correlations between successive frames are not considered which results in a traffic that behaves somehow different than the real traffic seen in our measurements.

To represent data traffic adequate (see [3]), we used a combination of three sub-sources (mean rate 200 kbit/s). Data 1 and Data 2 comprise WWW get requests and responses, respectively, as well as FTP and SMTP traffic. They are modelled as M/Pareto traffic without TCP ($\alpha = 1.2$, mean file size 3 KB, packet length 576 and 1500 respectively) which we believe to be a worst case estimate in our scenario. Data 3 represents the proportion of short acknowledgements in the network. We use infinite queues, again as a worst case scenario, and because finite queues are mainly reasonable in conjunction with TCP loopback. The segmentation of files to packets is done similarly to the video sources.

4. The width of the interval for which an average is obtained that corresponds to a respective value of the distribution of the frame length.

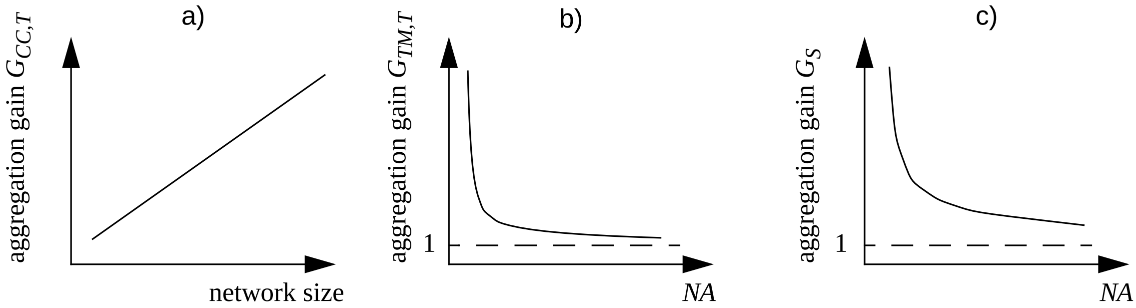

Fig. 7 Aggregation gain (qualitative)

4 Comparison of different Aggregation Strategies

Depending on the task a node has to fulfill in the network, traffic aggregation can either reduce or increase the complexity of traffic management. As mentioned earlier, aggregation increases the complexity of traffic management in the access networks in order to reduce the complexity in the WAN. **Fig. 6** gives an overview of nodes situated in different places in the network. The only functionality in the WAN with (partly) increased complexity is routing as additional information about aggregates has to be communicated in the network.

Section 4.1 discusses the qualitative progession of time and space complexity in the WAN depending on the number of aggregates whereas section 4.2 concentrates on the QoS within an aggregate and thus deals with the question weather the need of bandwidth in the WAN is increased or reduced depending on the different aggregation strategies. Section 4.3 focuses on the stability of the results obtained in Section 4.2.

4.1 Cost Aspects

In the following we present results of some related investigations that are valid for a WAN as a whole. For specific node types more concrete statements are possible but are not considered here. All results are simplified and contain only dependencies from relevant and dominant parameters and distinguish between costs due to time complexity and space complexity. Neglected are, e. g., the time complexity of policy enforcement or the space complexity of UPC, policy enforcement and shaping. To cope with situations in distinct networks a weighting of the different costs is necessary (c_i in the following).

In networks that apply connection control functions to establish, maintain and release connections there are related costs that depend directly on the mean number of nodes crossed by a connection or aggregate (*mean connection length MCL*) for time complexity, e. g. connection setup, and the maximum number of connections or aggregates (NC_{max}) that can be supported by individual network nodes for space complexity, e. g. space for connection tables. The reduction of net-

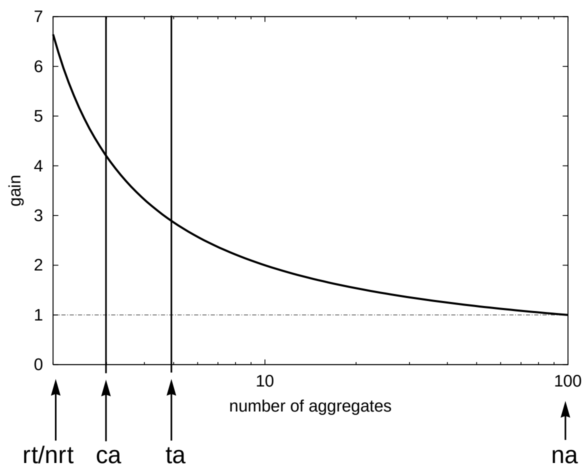

Fig. 8 Aggregation gain $G_{TM,T}$ for NC_{mean}*=100* sources

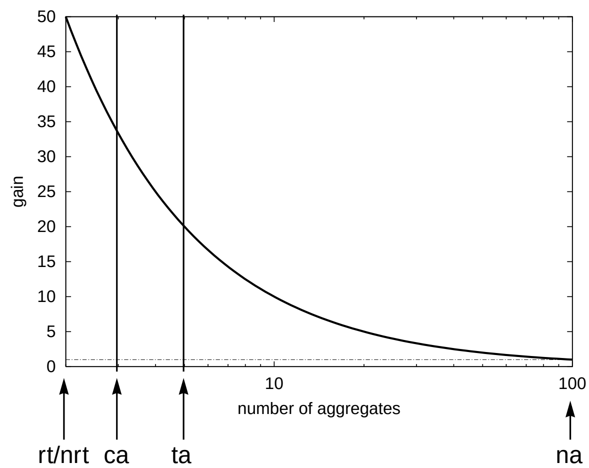

Fig. 9 Aggregation gain G_S for NC_{max}*=100* sources

work resources for connection control time due to traffic aggregation (gain) can be expressed as $G_{CC,T} = c_1 \cdot MCL$, see **Fig. 7a**. It shows that the reduction in complexity or increase in gain is most visible for large networks. For static traffic aggregation (no dynamic bandwidth allocation or dynamic setup for aggregates) there is no dependency from the number of aggregates since no connection control for aggregates is necessary. Memory resources for connection control are combined with traffic management, see below.

For all networks, necessary traffic management resources can be reduced by traffic aggregation, e. g. for usage parameter control, policy enforcement (like filtering) or scheduling. The effects are summarized in the formula $G_{TM,T} = c_2 \cdot \log NC_{mean} / \log NA$ for time complexity, see **Fig. 7b**, where NC_{mean} depicts the mean number of connections active in core nodes and NA the number of aggregates, and in $G_S = c_3 \cdot NC_{max} / NA$ for space complexity (covers connection control as well), see **Fig. 7c**. With increasing aggregation ratio (decreasing number of aggregates) the aggregation gain can be very high.

To get an impression of the order of magnitude of the aggregation gain, **Fig. 8** depicts the ratio of the gain $G_{TM,T}$ of an aggregate of arbitrary size over the aggregation gain $G_{TM,T}$ in the case of per flow queueing. Comparably, **Fig. 9** depicts the spacial ratio of gain G_S. Thus these two figurs show the gain in time and space complexity that can be obtained through aggregation. It can be seen that an increase in the number of aggregates to get additional isolation of traffic types has a strong impact if the number of aggregates is small whereas the impact reduces if the number of aggregates is already large.

4.2 Statistical Service Guarantees

In order to achieve a certain per (micro-) flow QoS, the factor of overbooking (effective bandwidth) of each source type has to be determined depending on the number of (micro-) flows. For our scenario this factor is calculated as the fraction of the required bandwidth over the cumulated mean value of all sources to obtain a probability of 10^{-3} that the delay exceeds 20 ms. Depicted over the number of aggregated sources, this factor represents the economy of scale of the respective source type which mainly depends on the burstiness of a source. The more similar these curves are, the higher is the influence of economy of scale in an aggregate compared to the influence of giving-up the isolation between the source types.

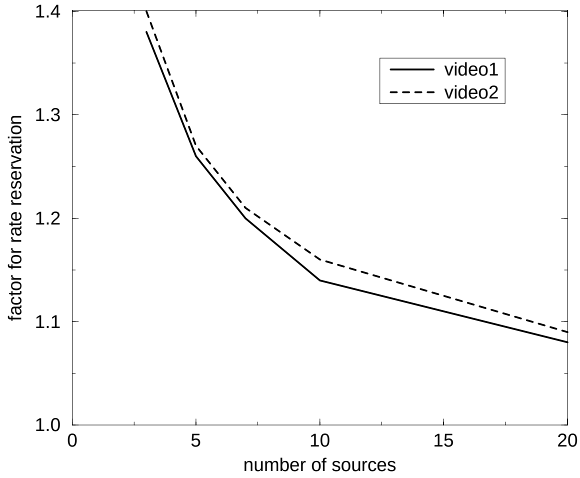

Fig. 10 Multiplex gain of different video sources

For both of the introduced audio source types, no overbooking is needed as they are CBR sources. The need for over-allocated bandwidth for the introduced video sources is determined by simulations which result in the rather similar curves depicted in **Fig. 10**. Thus – although the two video source types are quite different – it is obvious that it is better to aggregate them in order to efficiently take advantage of the available bandwidth. Nevertheless, we obtain an overbooking factor of about 1,26 that has to be considered for five video sources in order to meet our QoS requirements. For the introduced data traffic, we recommend to reserve a certain amount of bandwidth on which data sources with feedback mechanisms adapt their rate.

In the following simulations five aggregation strategies are evaluated. For ‘type aggregation’ (*ta*) each of the five introduced traffic types is queued separately, whereas for ‘class aggregation’ (*ca*) the number of aggregates is reduced to three by queueing the audio traffic, the video traffic and the data traffic in one aggregate each. ‘Audio video conference’ (*avc*) is similar to *ca* despite the fact that in the video queue for each video connection one audio connection is queued[5]. ‘Real-time/non-real-time’ (*rt*/*nrt*) combines all audio and video sources in one single queue and all data sources in another queue. These strategies are compared to ‘full aggregation’ (*FIFO*) where all sources are stored in the same queue and ‘per-flow queueing’ (no aggregation, *na*) where for every source a separate queue is provided.

5. This aggregation strategy considers the fact that in an audio video conference a video signal and an audio signal belong together and have to be resynchronized after transmission.

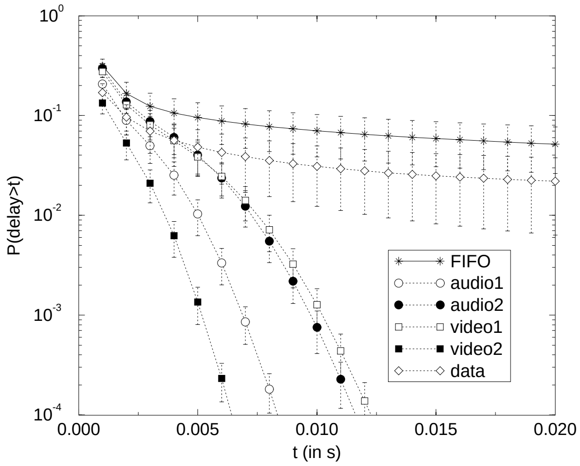

Fig. 11 Complementary waiting time distribution function of ‘full aggregation’ and ‘no aggregation’

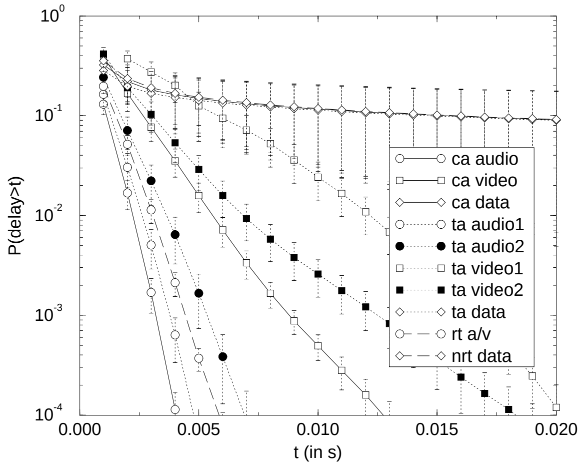

Fig. 12 Complementary waiting time distribution function of different aggregation strategies

Fig. 11 and **Fig. 12** show the complementary distribution function of the waiting time of the aforementioned traffic aggregates. **Fig. 11** depicts the cases ‘no aggregation’ and ‘full aggregation’ as reference values for the aggregation strategies ‘type aggregation’, ‘class aggregation’ and *rt*/*nrt* which are depicted in **Fig. 12**. We want to point out that there is an implicit coupling between bandwidth and delay[6] that affects the progression of all curves. As flows receive consecutive grants proportional to their allocated weights, smaller flows are discriminated which results in longer waiting times compared to larger flows. This can be observed e. g. for the different aggregation strategies for audio.

We also want to stress that in **Fig. 11** the complementary distribution function in case of ‘full aggregation’ is worse than all distribution functions in case of ‘no aggregation’ because the conservation law is not kept as the scheduler considers the lengths of packets [12].

The contemplation of the statistical delay bounds of the different aggregation strategies compared to the per-flow queueing shows that the need of bandwidth in the WAN increases in case of type aggregation. If class aggregation is applied, the required bandwidth in the core is about the same whereas *rt/nrt* aggregation saves bandwidth.

From a traffic engineering point of view, it can be concluded that the type aggregation strategy only leads to disadvantages for an enterprise network of the assumed size. The class aggregation strategy yields a good QoS that is approximately comparable to the QoS of per flow queued sources. The *rt*/*nrt* aggregation strategy yields the best results and is thus the preferred option in our scenario. However, if further services like non-interactive multimedia (streaming audio and video), multimedia database applications or games are added, the result for real-time traffic aggregation may change significantly.

6. This holds especially for all GPS emulating schedulers, e.g. SCFQ.

4.3 Stability of the Service Guarantees

As the considered scenario is static with respect to the traffic mix, we carried out further simulations to show that the obtained results are stable with respect to changes to the carried traffic or variation of the number of *nrt* sources.

Fig. 13 depicts the variation of the carried traffic. In this scenario, the number of sources is kept constant while the overall load is altered simply by changing the link bandwidth. It can be seen that the service experienced by the *rt* traffic decreases rapidly with increasing load but the statistical delay bound for delay and loss is still met. Admittedly the service experienced by data sources is unacceptably bad. Here, only a rate regulation (e. g. in TCP) leads to acceptable results with respect to experienced delays.

Fig. 14 shows another change to the scenario. Here, the number of *nrt* sources and the link bandwidth are scaled in such way that the load and the number of *rt* sources are kept constant. It can be seen that the behaviour of the real-time traffic is almost independent of its fraction on the link bandwidth.

This graph also shows that the influence of economy of scale of the data sources only slowly (with increasing number of sources) starts to overcome the influence of given-up isolation. This is an explanation why

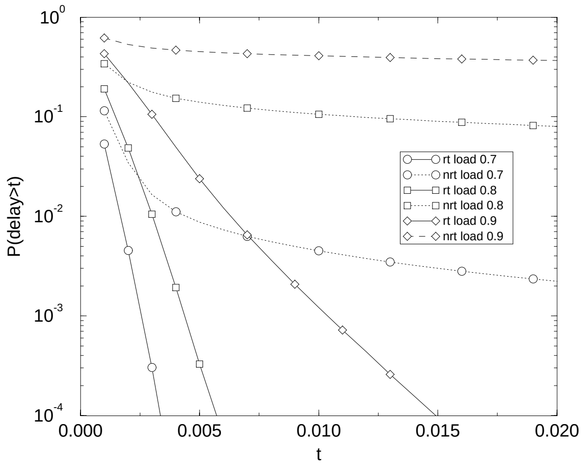

Fig. 13 Variation of the load in case of *rt/nrt* aggregation with a constant number of sources

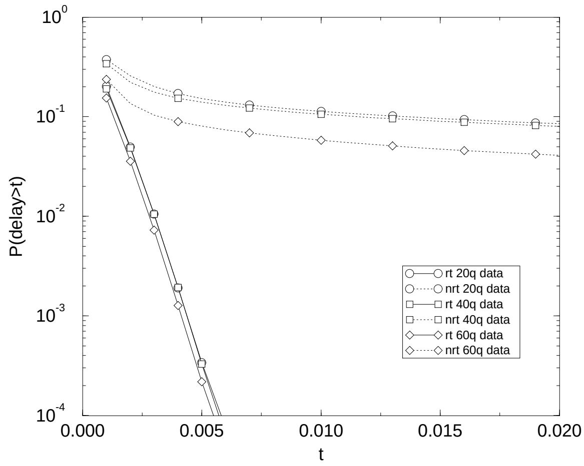

Fig. 14 Scaling of the number of *nrt* sources and the link bandwidth (constant load) with constant number of *rt* sources

the curves of the different aggregation strategies depicted in **Fig. 12** are much worse than the per-flow queueing case in **Fig. 11**.

Thus, from **Fig. 13** and **Fig. 14**, we can conclude that the obtained results for the behaviour of the real-time traffic are stable with respect to the overall load situations and the background traffic and not dependent on our chosen traffic scenario.

5 Conclusion

In this paper we presented approaches to evaluate different aggregation strategies with respect to costs in the WAN and waiting time in the aggregation node. We showed the qualitative progession of the gain concerning traffic management in time and space complexity that is obtained through aggregation. Afterwards, this aggregation gain of an arbitrary aggregation strategy compared to the per flow queueing strategy was presented. We showed that this gain rapidly decreases for a small number of aggregates.

Concerning the QoS requirements, we chose the waiting time as an important measure for real-time flows. In our simulations, we compared different aggregation strategies with each other as well as with FIFO and per-flow queueing. Since not only the aggregation gain but also the waiting time distribution shows a better behaviour than for all other aggregation strategies, the often suggested subdivision of the traffic in a real-time part and a non-real-time part is also our preferred solution in this scenario. However we want to point out that under different traffic conditions the experienced QoS may suggest another aggregation strategy. In this case the aggregation gain and the QoS have to be balanced in order to satisfy not only the scalability requirements but also the QoS requirements of individual flows.

Acknowledgement

We would like to thank our colleague Stefan Bodamer for his kind support and 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.: 'Messung und Modellierung von Internetverkehr', *Tutorium MMB'99*.

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

[5] 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.

[6] Händel, R.; Huber, M.; Schröder, S.: *ATM Networks: Concepts, Protocols, Applications*, Addison-Wesley, 1998.

[7] Handley, M.; Schulzrinne, H.; Schooler, E.; Rosenberg, J.: *SIP: Session Initiation Protocol*, IETF RFC 2543, March 1999.

[8] ITU-TSS: *Packet-based multimedia communication systems*, ITU-T Recommendation H.323, Feb. 1998.

[9] ITU-T: *Pulse Code Modulation (PCM) of voice frequencies*, ITU-T Recommendation G.711.

[10] ITU-T: *Dual rate speech coder for multimedia communications transmiting at 5.3 and 6.3 kbit/s*, ITU-T Recommendation G.723.1, March 1996.

[11] ITU-T: *Video codec for audivisual services at p x 64kbit/s*, ITU-T Recommendation H.261, March 1993.

[12] Kleinrock, L.: 'Queueing Systems, vol. II: Computer Applications', John Wiley & Sons, New York/Chichester/Brisbane/Toronto, 1976.

[13] 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.

[14] Perkins, M.; Dvorak, C.; Lerich, B.; Zebarth, J.: 'Speech Transmission Performance Planning in Hybrid IP/SCN Networks.' IEEE Communications Magazine, Vol. 37, No. 7, July 1999, pp. 126-131.

[15] Ramakrishnan, K.; Hjalmtysson, G.; Van der Merwe, J.: 'The Role of Signalling in Quality of Service Enabled Networks.' *IEEE Communications Magazine*, Vol. 37, No. 6, June 1999, pp. 124-132.

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

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

[18] 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.

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