Workshop on ATM Traffic Management WATM'95, IFIP, WG.6.2 Broadband Communication,
Paris, December 6-8, 1995

Adaptable Error Control for Efficient Provision of Reliable Services in ATM Networks

Georg Carle
University of Karlsruhe, Institute of Telematics,
Zirkel 2, 76128 Karlsruhe, Germany
Phone: +49 721 608-4027, Fax: +49 721 388097
e-mail: carle@telematik.informatik.uni-karlsruhe.de

Abstract

Distributed computing, distributed multimedia applications as well as advanced applications for computer-supported cooperative work depend on high performance networks providing point-to-point and point-to-multipoint communication services. ATM networks promise to provide adequate support for these applications. However, the traffic originating from these applications is highly bursty and unpredictable. While existing ATM protocol standards allow an efficient use of network resources for steady, predictable traffic, it is of high importance that ATM networks also allow for the efficient provision of services for highly bursty traffic. It is particularly difficult to provide multicast services for highly bursty sources efficiently, as there is a fundamental trade-off between throughput, cell loss rate, and use of network resources.

If non-negligible cell loss can be accepted, a large statistical multiplexing gain can be achieved at the ATM layer even for highly bursty sources. However, current adaptation layer protocols for data traffic are designed for very low cell loss probability and do not handle non-negligible cell loss well.

This paper presents a framework for error control within the ATM Adaptation Layer which allows to provide reliable point-to-point and point-to-multipoint services for different traffic management schemes even under conditions of high cell loss rates. The framework allows to select the error control scheme which provides the highest efficiency for a specific cell loss rate. It is of benefit for a variety of scenarios in combination with the ATM service categories Unspecified Bit Rate (UBR) and Available Bit Rate (ABR).

In the case of low cell loss rates, the Reliable Lightweight Multicast Protocol (RLMCP) is proposed as Service Specific Convergence Sublayer (SSCS) for AAL5, employing frame-based error control. In case of significant cell loss, large group sizes, and higher path capacities, and also for applications with stringent real-time requirements, the new adaptation layer type called the Reliable Multicast ATM Adaptation Layer (RMC-AAL) is proposed. It is the first AAL protocol offering cell-based ARQ and FEC for reliable multicasting and has a very low protocol overhead per cell. It is proposed to integrate RLMCP and RMC-AAL not only into ATM end systems, but also within ATM intermediate systems called Group Communication Servers.

1 INTRODUCTION

There exist a number of approaches in order to achieve statistical multiplexing gain for bursty sources [WoFD93]. The traffic management specification of the ATM Forum [ATMF95] presents a rate-controlled congestion control algorithm for the provision of ABR services. Alternatively, a credit-based scheme could be used [KuBC94]. Both approaches allow the sharing of bandwidth and of buffers in the switches between different virtual connections. In both approaches, the parameters of the mechanisms can be dimensioned in a way that a higher statistical multiplexing gain can be achieved, while accepting higher cell loss rates. There are advantages in such a dimensioning. It allows fast bandwidth ramp up and an increased number of admitted ABR VCs. There are other cases in which relatively high cell loss rates occur within the ATM network. One possibility is the use of the Fast Reservation Protocol with Immediate Transmission (FRP/IT) [BoTr92]. Another possibility is the assignment of different priorities to different VCs. A considerable amount of research in the area of flow and congestion control allowed to improve the understanding of the dynamics of various control strategies. There is not a single scheme which can be considered optimal. Instead, the selection of an appropriate congestion control scheme has to take into account the characterisation of the traffic sources, the strategy of resource sharing within the network, and the implementation complexity of the control algorithm, as well as costs of network resources.

In high-speed wide area ATM networks, aggressive control algorithms [KiFa95] offer the following two advantages. During transient periods, sources are not able to obtain an accurate view of the current load state. Their view of the network load is always outdated due to the propagation delay. It can not be expected that the actions performed by a source without accurate view of the current load allow to optimise the quality of service parameters (throughput, delay and loss). While conservative control algorithms may lead to long transient periods, aggressive control algorithms allow to shorten the transient period under certain conditions. Then they reach an equilibrium state faster, where the parameters can be adjusted according to optimality criteria. One example for a control algorithm that allows to adjust its aggressiveness is the concept of loss-load curves [WiCh91], which can be used to derive strategies for sources in order to maximise throughput or minimise end-to-end delay.

For multicast ABR services, different service models are possible. A conservative multicast service would limit the transmission rate according to the most congested link of a multicast tree, allowing to achieve very low cell loss rates. In contrast, an aggressive ABR multicast service may use higher transmission rates, while producing cell losses at some congested links.

If a single cell of an AAL3/4 or AAL5 frame is lost, typically the whole frame is discarded. By the use of a data link layer or transport layer protocol with error recovery mechanisms, a reliable service may be provided on top of an unreliable ATM service. If a reliable service in ATM networks is based on traditional transport protocols like TCP, severe performance degradations may be observed [Rom93]. For the provision of a reliable multipoint service, the probability for losses increases for a growing number of receivers. However, there are still no convincing concepts for reliable high-performance group communication in ATM networks. Therefore, the provision of reliable group communications requires the development of efficient protocols and of communication systems that achieve high performance even under conditions of high cell losses.

This paper focuses on design and assessment of error control mechanisms for correction of cell losses for point-to-point and point-to-multipoint communication. Section 2 gives an overview on related protocols for error recovery. In section 3, the proposed framework for reliable multicast communication is presented. Section 4 presents performance results of different error control schemes.

2 RELIABLE SERVICES IN ATM NETWORKS

According to the B-ISDN protocol reference model, mechanisms for error recovery may be integrated into the Service Specific Convergence Sublayer (SSCS) of the adaptation layer for provision of an assured mode service [I.363]. Up to now, only two SSCS-Protocols that offer error control mechanisms are specified by ITU. The Service Specific Connection Oriented Protocol (SSCOP) [Hen95] is subject of standardisation for an SSCS that offers assured mode service for signalling. The protocol provides end-to-end flow control and recovery of lost or corrupted data frames by selective retransmissions. However, SSCOP does not support assured mode multicast connections. For AAL1, an SSCS with FEC is proposed [I.363], based on a Reed-Solomon-Code applied on blocks of 128 cells that allows for the regeneration of up to four missing cells. Additional FEC schemes for ATM were proposed and investigated (e.g., [McA90], [Sha90], [OhKi91], and [ZhSa91]) but there are still a number of open questions concerning the combination of FEC and ARQ in ATM networks. Recently, a cell-based FEC scheme for AAL5 was proposed within the ATM Forum [CaEG95a], [CaEG95b].

While FEC allows to correct a certain number of losses originating from congestion, it may also lead to additional cell losses due to the increased load. Therefore, applying FEC is only useful in cases where the loss probability after decoding is lower than the loss rate of a data stream without FEC. As shown in [Sha90], [Bie93] and [AyGO93], an appropriately dimensioned amount of redundancy allows to reduce the loss probability after decoding.

In addition, [Bie93] showed that an FEC scheme may be used advantageously for multiplexing of VCs with different QoS requirements. If data streams with redundancy and without redundancy are multiplexed, different QoS requirements may be satisfied even for a switch that does not distinguish the data streams.

In [BaSS92] it was shown that for multiplexing of two traffic streams an increase in the burstiness of one stream negatively effects cell loss of the stream itself and the other stream. It also shows that the less bursty stream is also penalised, while the more bursty stream always benefits. It can be concluded that applying FEC for less bursty sources allows to protect less bursty sources from more bursty sources for both UBR services and for ABR services with aggressive parameters.

Transport protocols that are suitable for a connectionless network layer, like TCP, TP4 and XTP, are not very well suited to an ATM environment. The error control mechanisms of these protocols are very general and not designed for ATM cells and AAL frames. These transport protocols need to tolerate packets delivered out of sequence by the network layer. An adaptation layer protocol may benefit from the in-sequence delivery of the ATM-layer service and may use sequence number gaps for error detection. XTP offers support for reliable multicasting by a list-based algorithm and the so-called bucket algorithm. However, error control based on the bucket algorithm has significant shortcomings, as shown in [SaFd93].

TP++ [Fel93] is an example for a transport protocol that is suitable for ATM networks. It uses retransmissions in combination with FEC for error recovery (type I hybrid ARQ). However, it is only capable of unicast communication.

There are a number of additional approaches for integrating FEC into the transport layer (e.g., [ScPo94]). These schemes are capable of recovering lost packets. For transport protocols that reside above AAL5, the loss of a single cell leads to the discarding of a complete packet. In such cases, packet-based FEC within the transport layer has to recover not only the lost cell, but all cells of a corrupted frame. In ATM networks it is therefore possible to achieve a better performance with cell-based FEC schemes than with packet-based FEC schemes.

Another problem arises due to the fact that ATM signalling differs conceptually from signalling in traditional transport protocols. ATM is based on out-of-band signalling, while conventional transport protocols are based on in-band signalling. If these protocols are to be used in ATM networks, mapping of transport layer connection control to ATM signalling needs to be performed [KuSo93].

3 PROTOCOL ARCHITECTURE FOR RELIABLE COMMUNICATION

3.1 Frame-based ARQ

For a simple and efficient provision of reliable multicasting, the Reliable Lightweight Multicast Protocol (RLMCP) was developed as a Service Specific Convergence Sublayer for AAL5. It provides assured mode point-to-point and point-to-multipoint services using frame-based ARQ. Retransmissions may be performed in selective repeat or go-back-N mode. RLMCP has the following characteristics: Out of band signalling; no rate control (uses rate control of ATM layer); no multiplexing (uses multiplexing of ATM layer); no checksum (uses CRC-32 of AAL5); low protocol overhead (10 byte header per frame); frame size adaptable to current cell loss rate; Flag LastF for last frame of a burst (no idle messages); a lower window edge (LWE) that identifies lowest sequence number the transmitter would retransmit; and single cell acknowledgements, allowing for N:1-concentration of acknowledgement cells.

The RLMCP frame header (see Figure 1) is identical to the SSCS header of data frames of RMC-AAL explained in the next section. The protocol is based on the following data format: the first byte of the header (Discriminator/Dis) contains a cell type field (CT) which identifies an RLMCP header cell. It also contains a field (FrType) for the frame type (data frame, retransmission frame, or acknowledgement), a flag to request 'immediate acknowledgement' (I-Ack) and a flag indicating the 'last frame of burst' (LastF).

Frames carry frame sequence numbers of 24 bits, which is sufficient for high-speed WANs. Frames also carry a sequence number for the 'lower window edge' (LWE) of the transmitter, indicating the lowest sequence number a transmitter is prepared to repeat. Lost frames are detected by gaps of the frame sequence numbers, or by time-outs. The 'last frame of stream' bit allows the receivers to stop the loss detection timer. Receivers send acknowledgements periodically, after reception of a frame in which an 'immediate acknowledgement' bit is set, or after detection of a missing frame. Receivers may use cumulative positive acknowledgements, sending a lower window edge, and selective positive or negative acknowledgements, using bitmaps with a length of 32 bytes (for a sequence of up to 256 frames). For flow control, acknowledgements contain an upper window edge for the highest sequence number a receiver is prepared to receive.

3.2 Cell-based ARQ and FEC

The Reliable Multicast ATM Adaptation Layer (RMC-AAL) features cell-based ARQ and FEC for an efficient provision of reliable multicast services under conditions of higher or varying cell loss rates, and for applications with strong delay requirements. Error recovery of RMC-AAL is based on three schemes: pure ARQ, type I hybrid ARQ, and pure FEC. A fully reliable service and a service that assures delivery to a subset of K receivers are offered. Cell sequence numbers (CSN, 6 bits) are provided for detection of missing cells. Frames are identified by a frame sequence number (FSN, 24 bit) in the frame header. Like RLMCP, RMC-AAL uses the trailer of AAL5-CPCS, protecting the payload of a frame by the cyclic redundancy check CRC-32. Data frames have a protocol overhead of 10 bytes in the frame header and 8 bytes in the frame trailer. In each cell, they have an additional overhead of one byte (2 bit for cell type CT, and 6 bit cell sequence number, see Figure 2). Using cell type identifiers for RMC-AAL cells that are different from the RLMCP header cell type identifier allows to send both RMC-AAL and RLMCP frames over the same VC.

Even for high speed VCs in WANs, no large cell numbering space is required, because every cell is identified by both FSN and CSN. The alternative solution of identifying cells entirely by their cell sequence numbers leads to a significantly higher overhead per cell. For example, the protocol BLINKBLT [Gol90], which

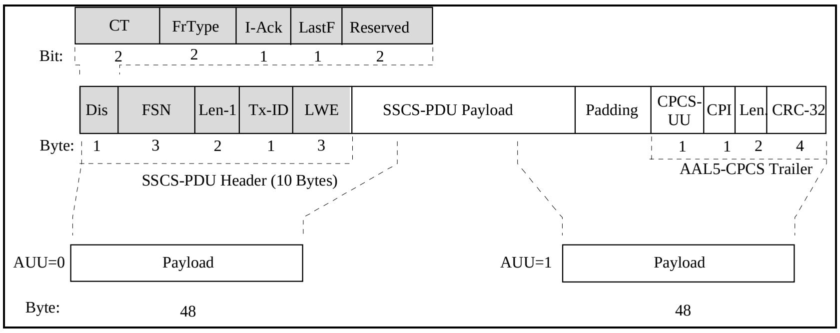

Figure 1: Data Frame for frame-based error control

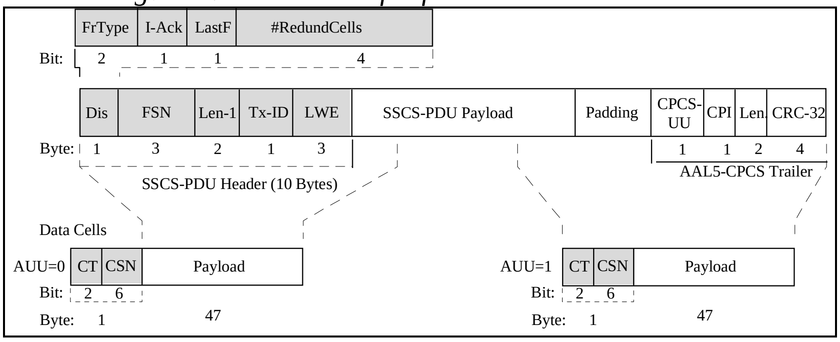

Figure 2: Data Frame for cell-based error control

also offers cell-based retransmissions, has a per-cell overhead of 4 bytes. The RMC-AAL frame header contains a transmitter identifier (Tx-id) and the length of the SSCS-PDU payload (Len-1). The frame header also contains the discriminator byte (Dis) with an identifier for the frame type (FrType), two flags (I-ACK and LastF), and the number of redundancy cells (#RedundCells) that follow the data frame. When FEC is used, h redundant cells are generated to protect the information cells of the frame. Encoding and decoding can be based on Reed-Solomon-Codes [McA90] or on simple XOR operations and matrix interleaving [Sha90]. Retransmissions may be sent by multicast or by unicast in selective repeat or go-back-N mode. It can be selected if retransmissions are frame-based (by retransmission of data frames) or cell-based (by retransmission of frame fragments). Frame fragments consist of a Fragment Header Cell, followed by a selection of original data cells of this frame. The fragment header cell contains the frame sequence number of the original frame. . A bitmap is used to indicate which cells of the original frame are retransmitted within the frame fragment. The field ‘Length of Bitmap’ indicates the valid length of the bitmap, and the field ‘Offset Bitmap’ indicates the cell number of the first bit of the bitmap.

3.3 Group Communication Server (GCS)

The application of the presented error control mechanisms is not limited to ATM end systems. The deployment of so-called Group Communication Servers with multicast error control mechanisms allows to provide reliable high-performance multipoint services for a wide range of parameters. Further improvements of performance and efficiency can be achieved by using GCSs hierarchically.

GCSs support an efficient use of network resources by performing multicast error control within the network. Allowing retransmissions originating from the server avoids unnecessary retransmissions over common branches of a multicast tree. The integration of FEC mechanisms into the GCS allows for the regeneration of lost cells and reinsertion of additional redundancy for adjusting the FEC coding scheme according to the needs of subsequent hops. By the provision of protocol processing support for multicast transmitters, GCSs also improve scalability. They also support heterogeneous multicasting by permitting to apply different protocol parameters for different branches of a multicast tree, and by conversion of different error schemes.

For groups with multiple transmitters, the GCS provides support for multiplexing of frames onto a single point-to-multipoint connection. This allows to reduce the number of required VCs significantly for large groups with many transmitters [Wei92]. Virtual LANs frequently require this multiplexing functionality. However, additionally queuing delay may be introduced by this multiplexing function. If LAN Emulation [LANE95] is used in a local ATM network, a GCS might be incorporated into a LAN Emulation Server (LES) or Broadcast and Unknown Server (BUS), thus making it possible for applications to ensure the reliable delivery of broadcast messages to all peers.

When used in combination with ABR services, a GCS may terminate an ABR control loop (corresponding to the concept of virtual sources and virtual destinations [ATMF95]), and by allowing to adapt to congestions on certain branches of a link by sending with different rates on outgoing VCs.

4 PERFORMANCE EVALUATION

It is important to know which error control scheme is best suited for a given situation. Analytical methods were applied and simulations were performed in order to evaluate the achievable performance of the proposed error control schemes. For modelling the correlation properties of lost cells, a two state Markov model (Gilbert Model) may be applied. Based on the worst case observations of [OhKi91], a probability of 0.3 was used for a cell discard following a cell discard. This is equivalent to cell losses with a mean burst length of 1.428 cells. A multicast VC was simulated, where cell losses may occur on the common link, or on the individual links. The same error model was applied to all links. Figure 6 shows the efficiencies (relation of successfully transmitted useful frames to total number of transmitted frames) mean delays of cell-based and frame-based ARQ with and without FEC for varying cell loss probability.

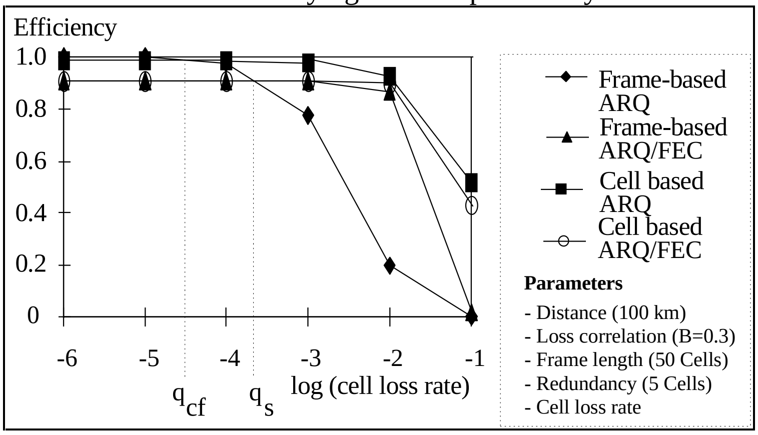

Figure 6: Simulation of efficiency for different error control schemes

For the following analysis, a system model was applied based on time slots T with T=2.831µs for the transmission of a single ATM cell. The normalised path capacity S represents the number of cells that may be stored on the links and in buffers of multiplexing equipment between the transmitter and a receiver. In the following, a multicast scenario with a common link, an ATM-switch in which copying of cells is performed, and individual links to the receivers will be investigated. For identical cell loss probabilities q on all links, the probability for successful delivery of a frame with a length of k cells to all receivers is given by

$$1 - Q = (1 - q)^{k \cdot (N+1)} \quad , \qquad (1)$$

where N denotes the number of receivers. Formula (1) uses the assumption of statistically independent cell losses. For schemes in which complete frames are discarded after the loss of a single cell, such losses lead to a higher frame loss probability Q than statistically dependent cell losses [BoLa93]. Therefore, (1) may also be used as a conservative approximation in the case of correlated cell losses.

In the next step, an approximationfor the achievable efficiency will be derived. A worst-case approximation (Weldon approximation of [Wel82]) is used to determine a proper receiver buffer size for the selective repeat protocol. Shacham [Sha87] shows that the mean receiver buffer occupancy is finite for the case when the probability to successfully transmit a frame approaches zero. The size of a finite window sufficiently large to achieve an efficiency close to the efficiency in the ideal case of infinite receiver buffers will be derived in the following. If m denotes the number of trials for successful transmission of a frame, the efficiency is in inverse proportion to the mean number of transmission trials per frame. Hence the normalised throughput efficiency η is

$$\eta = \frac{1}{\overline{m}} \quad . \qquad (2)$$

To take into account the influence of limited receiver buffers on the achievable efficiency, Weldon's approximation [Agh94] may be applied for receiver buffers which are capable of storing L path capacities (i.e., the receiver buffer holds L·S frames). It is assumed that all receiver buffers are empty at the beginning of the transmission, and that the probability of a frame being successfully received on the first L+1 transmission attempts is given by

$$P(m = i) = (1 - Q) \cdot Q^{i-1} \quad , \qquad 1 \le i \le L + 1 \quad . \qquad (3)$$

For every retransmission, S new frames have been sent before the next retransmission takes place. If more than L+1 transmission attempts are necessary, the frames that are meanwhile correctly received from the previous transmission attempts will have filled up the receiver buffers. Further frames have to be discarded due to receiver buffer overflow, so they also have to be retransmitted. Hence,

$$P(m = i + (i - (L+1)) \cdot S) = (1-Q) \cdot Q^{i-1}, \; L+2 \le i \le \infty . \tag{4}$$

The Weldon approximation gives a lower bound for the efficiency, as the maximum number of S frames is lost only if S frames have been successfully transmitted between two transmission attempts of the considered frame. While this assumption works well for low frame loss probabilities, there is a growing deviation from the real efficiency at higher frame loss probabilities. For (3) and (4), a lower bound of the efficiency is given by

$$\eta_{min} = \frac{1-Q}{1+S \cdot Q^{L+1}} = \frac{(1-q)^{k(N+1)}}{1+S \cdot (1-(1-q)^{k(N+1)})^{L+1}} \quad . \tag{5}$$

Taking into account the influence of protocol overhead in frames and cells, cases can be identified for which the cell-based scheme shows better performance than a frame-based scheme. In [BoLa93], a comparison of cell-based vs. frame-based selective repeat is presented for point-to-point-communication and for the assumption of unlimited receiver buffers. Here, the influence of a limited receiver buffer is also taken into account. The derivation uses the term overhead ratio as fraction of (protocol overhead + payload)/payload. In the following formulas, k_F denotes the overhead ratio of a frame-based scheme (e.g., RLMCP), in which the protocol overhead for the initial transmission attempt and the retransmissions are identical. For the cell-based scheme of RMC-AAL, the protocol overhead of the initial transmission is different from the protocol overhead of subsequent transmissions. Therefore, the overhead ratio k_{C1} is defined for the initial transmission, and the overhead ratio k_{C2} is defined for subsequent transmissions. Considering the protocol overhead of a frame-based selective repeat scheme in the derivation of (5), a lower bound for the efficiency for receiver buffers of one path capacity is given by

η_F = **Fehler!** . (6)

The overhead ratio k_F is growing for shorter frames, while the frame loss probability is decreasing for shorter frames. There exists an optimal frame size for every combination of cell loss rate and number of receivers for which the efficiency (6) can be maximised. Figure 7 shows this optimal frame size. It can be seen that significant cell loss rates, as well as larger groups, both lead to relatively short frames. A disadvantage of short frames is not only their low bandwidth efficiency, but also their increased processing costs.

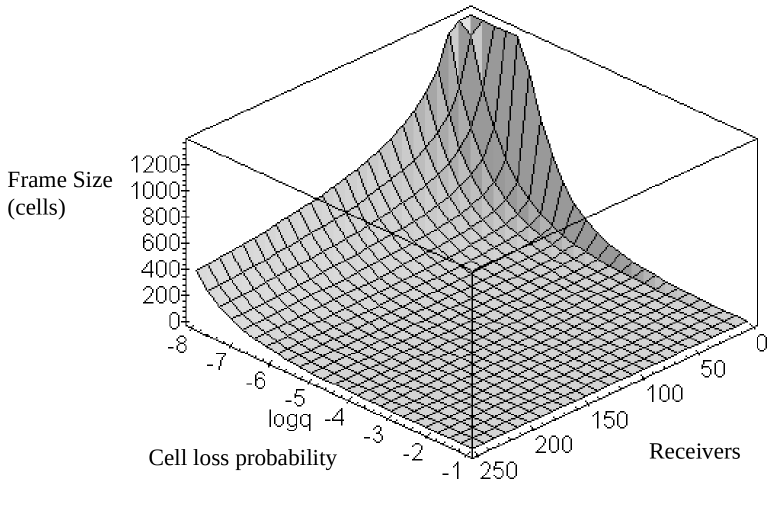

Figure 7: Optimal frame size for frame-based retransmission

For the cell-based scheme of RMC-AAL and receiver buffers of one path capacity, a lower bound for the efficiency is given by

η_C = **Fehler!** (7)

The efficiency equilibrium q_{cf} can be obtained by solving the equation $\eta_F = \eta_C$, using

$$Q = (1-q)^{kN} \tag{8}$$

for the frame loss rate. Substituting q·N by x results in a cubic equation

$$x^3 + Ax^2 + Bx + C = 0 \quad \text{with} \tag{9}$$

A = **Fehler!**; B = **Fehler!**; C = **Fehler!**. (10)

The analytical solution of a cubic equation is rather cumbersome. Solutions can be found in literature (e.g., see [Bro89]). After identifying the correct one of three possible solutions of (9), Figure 8 shows the resulting threshold for the efficiency equilibrium for a path capacity of 1750 cells (e.g., 150 Mbit/s over 1000 km). For every frame size, a line of the figure separates the region with lower cell loss rates, where a frame-based scheme results in better efficiency, from the region with higher cell loss rates, where a cell-based scheme leads to better efficiency.

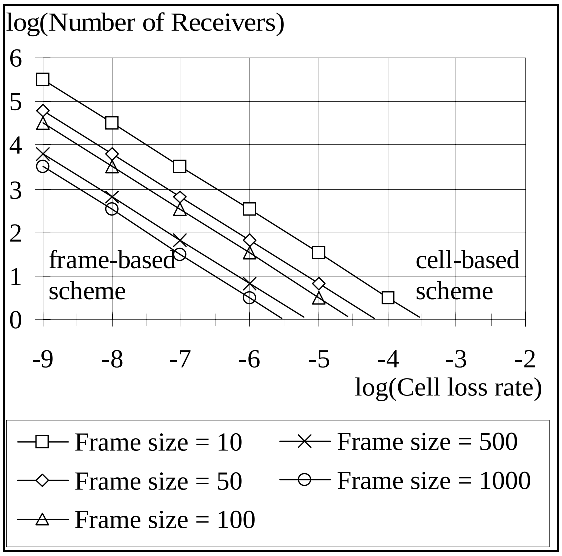

Figure 8: Frame-based vs. cell-based retransmissions

5 CONCLUSION

Current AAL protocols for data traffic are designed for very low cell loss probability and do not handle non-negligible cell loss well. For steady, predictable traffic, they allow for an efficient use of network resources. However, a large portion of the traffic to be transported over ATM networks is highly bursty and unpredictable. If non-negligible cell loss can be accepted by the AAL, a large statistical multiplexing gain can be achieved at the ATM layer even for highly bursty sources.

An adaptable framework for error control was presented in this paper which provides reliable point-to-point and point-to-multipoint services for a wide range of cell loss rates. For low cell loss rates, a frame-based end-to-end error control is most appropriate. In this case, the Reliable Lightweight Multicast Protocol (RLMCP) employing frame-based error control is proposed as SSCS for AAL5. In case of significant cell loss, large group sizes, and higher path capacities, and also for applications with stringent real-time requirements, the new adaptation layer type called the Reliable Multicast ATM Adaptation Layer (RMC-AAL) is proposed. It is the first AAL protocol offering cell-based ARQ and FEC for reliable multicasting and has a very low protocol overhead per cell. It is proposed to integrate RLMCP and RMC-AAL not only into ATM end systems, but also within ATM intermediate systems called Group Communication Servers.

Results of the performance evaluation are useful for selection of the error control scheme, and for dimensioning of frame sizes. It is shown how the number of receivers affects the appropriate frame size. A cell loss threshold is derived for which a cell-based scheme outperforms a frame-based scheme.

References

[Agh94] Aghadavoodi, J.: „Multicopy ARQ Strategies for Heterogeneous Networks,“ 4th IFIP International Workshop on Protocols for High Speed Networks, Vancouver, B.C, Canada, Aug. 1994

[ATMF95] Sathaye, S. (Ed.): „Draft ATM-Forum Traffic Management Specification 4.0,“ ATM Forum Document 95-0013, 1995

[BaSS92] Bae, J.; Suda, T.; Simha, R.: „Analysis fo Individual Packet Loss in a Finite Buffer Queue with Heterogeneous Markov Modulated Arrival Processes: A Study of Traffic Burstiness and Priority Packet Discarding,“ IEEE Infocom 1992, pp. 219-231

[Bie93] Biersack, E.: „Performance Evaluation of Forward Error Correction in an ATM Environment,“ IEEE Journal on Selected Areas in Communication, Volume 11, Number 4, pp. 631-640, May 1993

[BoLa93] Bondi, A.; Lai, W.-S.: „The influence of cell loss patterns and overheads on retransmission choices in broadband ISDN,“ Computer Networks and ISDN Systems 26, pp. 585-598, 1994

[BoTr92] Boyer, P.; Tranchier, D.: „A reservation principle with applications to the ATM traffic control,“ Computer Networks and ISDN Systems 24 (1992), pp. 321-334

[Bro89] Bronstein, I., Semendjajew, K.: „Taschenbuch der Mathematik,“ (in German) Verlag Harri Deutsch, Thun und Frankfurt/Main 1989.

[CaEG95a] Carle, G.; Esaki, H.; Guha, A.; Tsunoda, K.; Kanai, K.: „Necessity of an FEC Scheme for ATM Networks,“ Contribution ATMF/95-0325; ATM Forum Technical Committee 'Service Aspects and Applications SA&A', April 10-14, 1995, Denver, U.S.A.

[CaEG95b] Carle, G.; Esaki, H.; Guha, A.; Tsunoda, K.; Kanai, K.: „Proposal for Specification of FEC-SSCS for AAL Type 5,“ Contribution ATMF/95-0326; ATM Forum Technical Committee 'Service Aspects and Applications SA&A', April 10-14, 1995, Denver, U.S.A.

[Fel93] Feldmaier, D.: „An Overview of the TP++ Transport Protocol Project,“ Chapter 8 in ‘High Performance Networks - Frontiers and Experience,’ Ahmed Tantawy (Ed.), Kluwer Academic Publishers, 1993

[Gol90] Goldstein, F.: „Compatibility of BLINKBLT with the ATM Adaptation Layer,“ ANSI Technical Subcommitee T1S1.5/90-009, Raleigh, NC, U.S.A., Feb. 1990

[Hen95] Henderson, T.: „Design Principles and Performance Analysis of SSCOP: A new ATM Adaptation Layer Protocol,“ Computer Communication Review, Vol. 25, Nr. 2, Apil 1995, pp. 47-59

[I.363] ITU-TS Draft Recommendation I.363: „BISDN ATM Adaptation Layer (AAL) Specification,“ Geneva, 1993

[KiFa95] Kim, H.; Farber, D.: „The Failure of Conservative Congestion Control in High-Speed Networks,“ Proceedings of Second Gigabit Networking Workshop (GBN'95), Boston, MA, U.S.A., April 1995

[KuBC94] Kung, H.T.; Blackwell, T.; Chapman, A.: „Credit-Based Flow Control for ATM Networks: Credit Update Protocol, Adaptive Credit Allocation, and Statistical Multiplexing,“ Proceedings of ACM SIGCOMM 94, London, England, August 1994, pp. 101-114

[KuSo93] Kure, O.; Sorteberg, I.: „XTP over ATM,“ Computer Networks and ISDN Systems, Vol. 26, pp. 253-262, 1994

[LANE95] ATM Forum: „LAN Emulation Over ATM: Draft Specification,“ LAN Emulation Sub-working Group, ATM Forum Technical Committee, August 1995

[McA90] McAuley, A.: „Reliable Broadband Communication Using a Burst Erasure Correcting Code,“ ACM SIGCOMM '90, Philadelphia, PA, U.S.A., Sep. 1990

[OhKi91] Ohta, H.; Kitami, T.: „A Cell Loss Recovery Method Using FEC in ATM Networks,“ IEEE Journal on Selected Areas in Communications, Vol.9, No.9, Dec. 1991, pp.1471-1483

[Rom93] Romanov, A.: „Some Results on the Performance of TCP over ATM,“ Second IEEE Workshop on the Architecture and Implementation of High Performance Communication Subsystems HPCS’93, Williamsburg, VA, U.S.A., Sep. 1993

[SaFd93] Santoso, H.; Fdida, S.: „Transport Layer Multicast: An Enhancement for XTP Bucket Error Control,“ in: Danthine, A.; Spaniol, O. (Eds.): High Performance Networking, IV, IFIP, North-Holland, 1993

[ScPo94] Schatzmaier, R; Popescu-Zeletin; R..: „Providing Support for Data Transfer in a New Networking Environment,“ in D. Hutchison, A. Danthine, H. Leopold, G. Colson (Eds.): Multimedia Transport and Teleservices; International COST 237 Workshop, Vienna, Austria, November 13-15, 1994, Lecture Notes in Computer Science 882, Springer-Verlag, pp. 241-255

[Sha87] Shacham, N.: „Queuing Analysis of a Selective-Repeat ARQ-Receiver“, in Proceedings IEEE INFOCOM '87, San Francisco, California, February 1987, pp. 521-526

[Sha90] Shacham, N.; McKenny, P.: „Packet recovery in high-speed networks using coding,“ in Proceedings of IEEE INFOCOM '90, San Francisco, CA, pp. 124-131, June 1990

[Wei92] Wei, L.; Liaw, F.; Estrin, D.; Romanow, A.; Lyon, T.: „Analysis of a Resequencer Model for Multicast over ATM Networks,“ Third International Workshop on Network and Operating Systems Support for Digital Audio and Video, San Diego, CA, U.S.A., Nov. 1992

[Wel82] Weldon, E.: „An Improved Selective-Repeat ARQ Strategy,“ IEEE Trans. Com., vol. COM-30, pp. 480-486, Mar. 1982

[WiCh91] Williamson, C.; Cheriton, D.: „Loss-Load Curves: Support for Rate-Based Congestion Control in High-Speed Datagram Networks“, Proceedings of ACM SIGCOMM 1991, Sep., 1991, Zurich, Switzerland, pp. 17-28

[WoFD93] Worster, T.; Fischer, W.; Davis, S.P.: Resource Allocation for Packet Data Traffic on ATM :Problems and Solutions“, in Gerner, N.; Hegering, H.-G.; Swoboda, J. (Ed.): Proceedings of „Kommunikation in Verteilten Systemen“; Munich, Germany, March 1993 Springer-Verlag, pp. 100-113

[ZhSa91] Zhang, L.; Sarkies, K.: Modelling a virtual path and its application for forward error recovery schemes in ATM networks, in Proc. of SICON'91, Singapore, Sep. 1991, pp. 259-264