Towards Scaleable Error Control for Reliable Multipoint Services in ATM Networks

Georg Carle
Institute of Telematics, University of Karlsruhe, D-76128 Karlsruhe, Germany
Tel.: +49/721/608-4027, Fax: +49/721/388097; E-Mail: g.carle@ieee.org

Abstract

Advanced applications, such as distributed multimedia applications, require efficient communication subsystems providing a variety of services. Existing communication systems face increasing difficulties in fulfilling these requirements. In particular, the efficient provision of reliable group communication services in ATM-Networks remains a major unresolved issue. This paper presents a novel framework for Adaptation Layer error control mechanisms. Two adaptation layer protocols are presented that provide reliable multicast services. The first one, called RLMCP (Reliable Lightweight Multicast Protocol), is a simple and efficient adaptation layer protocol for the Service Specific Convergence Sublayer of AAL5. It uses a frame-based ARQ-scheme and is suitable for virtual connections with low cell loss rates. The second one, called RMC-AAL (Reliable Multicast ATM Adaptation Layer), features cell-based ARQ in combination with FEC. The framework permits to select the combination of error control mechanisms most suitable for the application requirements in a specific communication scenario. The achievable performance is analysed, identifying the influence of group size, cell loss rate and path capacity on throughput and delay. Guidelines are presented for selection of the error control scheme most appropriate for a specific communication scenario.

1 Introduction

Upcoming applications, for example distributed multimedia systems, computer-supported co-operative work (CSCW) applications, and virtual shared memory systems require reliable high performance multipoint communication services. Quality of service (QoS) issues of importance are not only throughput, delay, and delay jitter, but also differences of delay and reliability within the group. A key problem that must be solved to provide a reliable multipoint service is the recovery from cell losses due to congestion in the switches. The probability for cell loss may vary over a wide range, depending on the strategy for usage parameter control (UPC) and call admission control (CAC) which is applied. It is still an open question how low cell loss rates can be guaranteed for bursty multicast traffic, while using network resources efficiently. Cell losses caused by buffer overflows do not occur randomly distributed, but show a highly correlated characteristic [OhKi91]. 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 communications 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 with high cell losses.

This paper focuses on the design and assessment of error control mechanisms for the correction of cell losses in multipoint communications. Section 2 gives an overview on related protocols for error recovery. In section 3, the proposed framework for the integration of error control mechanisms into the ATM Adaptation Layer is presented. Section 4 presents a performance evaluation.

2 Protocols for Reliable Services

Adaptation Layer Protocols: 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) 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. As shown in [InMo94], it is possible to extend SSCOP to allow for partial-frame retransmissions. In [Gol90], an AAL protocol for cell-based retransmission was proposed as an extension of AAL3/4. For AAL1, an SSCS with FEC is proposed [I.363], based on a Reed-Solomon-Code applied on blocks of 128 cells that allows the regeneration of up to four missing cells. Additional FEC schemes for ATM were proposed and investigated in [Sha90] and [Bie92], but there are still a number of open questions concerning the combination of FEC and ARQ in ATM networks.

Transport Protocols: 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). Up to now, no protocol that combines ARQ and FEC was presented for multicast communication in ATM networks.

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].

Protocol Implementation: While transmission capacity was growing enormously over the last years, protocol processing and system functions in the transport component turned out to be a performance bottleneck. High performance communication subsystems, based on parallel protocol processing, and hybrid architectures with hardware components for time-critical operations [BrZ92] are required for the provision of a service with high throughput and low latency. For highest performance, complete VLSI implementations of transport subsystems are planned [ScB93]. The performance bottleneck of the transport component that can be observed for point-to-point-communication is even more crucial for reliable multipoint connections. For a growing number of receivers, processing of a growing number of control packets (known as the implosion problem), and management of a large amount of status information needs to be performed.

Selection of Protocol Mechanisms: In order to offer a wide range of services to the applications for various network parameters, several concepts of flexible communication subsystems are under development. The parallel transport system PATROCLOS [BrZ92] is a parallel implementation of a high performance transport system, offering a wide range of protocol mechanisms that may be selected according to the needs of an application. The Flexible Communication SubSystem (FCSS) [ZST93] is a configurable, function-based transport system. It allows protocol configuration and the reservation of resources in order to provide applications with a specific service quality.

3 Reliable Multipoint Communication

A conceptual framework was developed that allows to select the error control mechanisms most suited for a specific environment. It describes how frame-based ARQ, cell-based ARQ, and cell-based FEC may be integrated into the adaptation layer for the efficient provision of reliable multicast services.

The framework considers a number of different group communication services: a fully reliable multicast service with assured delivery to every receiver, a K-reliable multicast service with assured delivery to at least K receivers of a group, and a real-time service in which error control is performed subject to deadlines.

As lost retransmissions contribute significantly to the QoS offered by the AAL, it is frequently of high importance to decrease the probability of lost retransmissions. The cell-based retransmission scheme and the FEC scheme of RMC-AAL explained in section 3.2 allow to decrease this probability. Additionally, the capability of ATM to offer virtual channels (VCs) with different cell loss probabilities may be used. This is in contrast to conventional networks, where initial transmissions and retransmissions will generally observe identical loss rates. The framework allows to select one of the following alternatives for reliable group communication:

- In the simplest case, a single 1:N multicast VC from the transmitter to the receivers with same QoS for all cells will be used. This requires demultiplexing above the ATM layer to distinguish ordinary transmissions from retransmissions.
- Improved performance may be achieved by an 1:N VC with different cell loss priorities according to the cell loss priority (CLP) bit in the ATM cell header. The initial transmission of a frame uses lower priority cells (CLP = 1). Retransmitted frames are sent with higher priority cells (CLP = 0). However, when ATM traffic control uses its possibility of converting high priority cells to low priority cells, the QoS of the retransmissions will not be higher than the QoS of ordinary transmissions.
- For ensuring the QoS of retransmissions, two 1:N VCs from transmitter to receivers may be used. One VC has a QoS suitable for the initial transmission. The second VC is intended for retransmissions. Its QoS will be set to low delay and low cell loss probability.
- In order to support individual retransmissions to receivers that observed losses, one of the alternatives above may be combined with a number of VCs from the transmitter to single receivers or to a subset of receivers. This allows that retransmissions lead to a reduced network load, and prevents receivers from the need to filter out unwanted retransmissions.

3.1 SSCS with 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.

The RLMCP frame header has a length of 10 bytes and 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 indicates the frame type (data frame, retransmission frame, or acknowledgement). It also contains 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. This allows to avoid unnecessary retransmission requests in K-reliable and real-time services.

Receivers send acknowledgements periodically, after reception of a frame in which an 'immediate acknowledgement' bit is set, or after detection of a missing frame. 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 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 1). 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 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. Redundancy cells use independent cell sequence numbers (RCSN), see Figure 2. 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 (see Figure 3) 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 (FSN) of the original frame. This field is called 'Start of Bitmap' (SBM). A bitmap (BM) is used to indicate which cells of the original frame are retransmitted within the frame fragment. The field 'Length of Bitmap' (LBM) indicates the valid length of the bitmap, and the field 'Offset Bitmap' (OBM) indicates the cell number of the first bit of the bitmap.

Receivers send acknowledgements periodically, after reception of a frame in which an 'immediate acknowledgement' (I-ACK) bit is set, and after detection of cell loss. The format of acknowledgements is shown in Figure 4. Acknowledgements contain a receiver identifier (Rx-id). An upper window edge (UWE) allows for window flow control. Receivers may use cumulative positive acknowledgements of frames by sending the frame sequence number of their lower window edge (LWE). Additionally, they may use bitmaps (BM) with a length of 32 bytes for negative acknowledgements of frames or individual cells. A frame sequence number (SBM), a field for the valid length of the bitmap (LBM), and a field for the offset of the bitmap identify the position of the bitmap within the window.

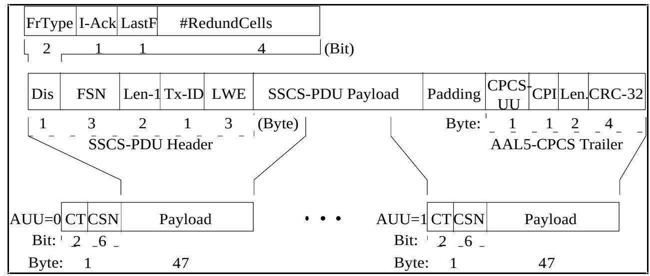

Figure 1: RMC-AAL Data Frame

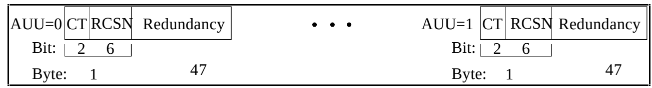

Figure 2: RMC-AAL Redundancy Frame

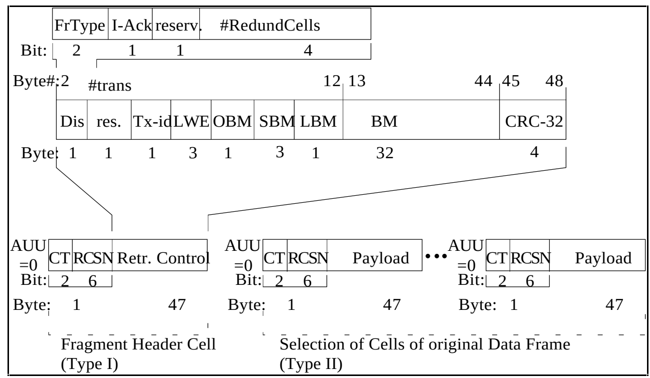

Figure 3: RMC-AAL Frame Fragment for Retransmissions

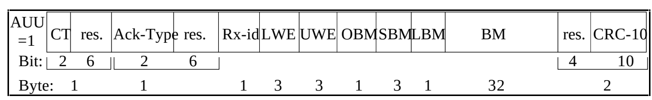

Figure 4: RMC-AAL Acknowledgement cell

4 Performance Evaluation

4.1 Assessment of Error Control Schemes

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. The following four alternatives were compared: frame-based ARQ, frame-based ARQ combined with FEC, cell-based ARQ, cell-based ARQ combined with FEC.

For the following analysis, a system model was applied based on time slots T for the transmission of a single ATM cell. For a 155Mbit/s SDH link with a transfer rate of 149,76Mbit/s at the ATM layer, the mean cell interarrival time is T=2.831µs. The normalized path capacity S represents the number cells that may be stored on the links and in buffers of multiplexing equipment between the transmitter and a receiver. With RTT denoting the mean round trip time, S may be expressed as

$$S = \frac{RTT}{T} \quad . \tag{1}$$

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. The path capacity of the common link will be called S_C, and the path capacity of the individual link to receiver number i will be called S_i. Similarly, the cell loss probability of the common link will be called q_C, and the cell loss probability of the individual links will be called q_i. Without FEC, the probability for successful delivery of a frame with a length of k cells to all receivers is given by

$$1 - Q = (1 - q_C)^k \cdot \prod_{i=1}^{N} (1 - q_i)^k \quad . \tag{2}$$

with Q denoting the frame loss probability. For identical cell loss probabilities on all links, (2) can be simplified to

$$1 - Q = (1 - q)^{k \cdot (N+1)} \quad . \tag{3}$$

where *q* denotes the cell loss probability and *N* denotes the number of receivers. Formula (2) 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, (2) may also be used as a conservative approximation in the case of correlated cell losses.

In the next step, an upper bound and a lower bound for 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

$$\eta = \frac{1}{\overline{m}} \quad . \tag{4}$$

4.1.1 Infinite Receiver Buffer

For infinite receiver buffers, the probability that the receivers get a frame in the m = i-th transmission trial is geometrically distributed [Sab82]:

$$P(m = i) = Q \cdot (1 - Q)^{i-1}, \; i \geq 1 \quad . \tag{5}$$

Then, a lower bound for the mean of the random variable m is given by

$$\overline{m}_{min} = \sum_{i=1}^{\infty} i \cdot P(m=i) = (1\text{-}Q) \sum_{i=1}^{\infty} i \cdot Q^{i-1} = 1\text{-}Q \cdot \frac{1}{(1\text{-}Q)^2} = \frac{1}{(1\text{-}Q)} . \tag{6}$$

Using (3) and (4), the maximum achievable efficiency is

$$\eta_{max} = (1 - Q) = (1 - q)^{k(N+1)} \quad . \tag{7}$$

4.1.2 Finite Receiver Buffer

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

$$P(m = i) = (1 - Q) \cdot Q^{i-1} \quad , \qquad 1 \leq i \leq L + 1 \quad . \tag{8}$$

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,

$$P(m = i + (i - (L + 1)) \cdot S) = (1 - Q) \cdot Q^{i-1}, \; L + 2 \leq i \leq \infty \, . \tag{9}$$

However, the Weldon approximation is a pessimistic approximation: 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 (8) and (9), an upper bound for the mean of the random variable m can be derived (see [Wel82]) to

$$\overline{m}_{max} = \frac{1 + S \cdot Q^{L+1}}{1 - Q} \quad , \tag{10}$$

and a lower bound of the efficiency is given by

$$\eta_{min} = \frac{1 - Q}{1 + S \cdot Q^{L+1}} = \frac{(1 - q)^{k(N+1)}}{1 + S \cdot (1 - (1 - q)^{k(N+1)})^{L+1}} \quad . \tag{11}$$

Hence, the achievable efficiency is bounded by

$$\eta_{min} \leq \eta \leq \eta_{max} \quad . \tag{12}$$

4.1.3 Evaluation of buffer size and frame size

Figure 5 shows the analytical results and the results of a simulation for a multicast scenario with sixteen receivers and a frame-based ARQ scheme. The analysis is based on Weldon's approximation and receiver buffers of size S, of size 5·S and of unlimited size. Furthermore, a frame size of 50 cells and a path capacity of 8 frames (approximately 50 km for a 150 Mbit/s link) were assumed. For larger receiver buffer sizes, a higher efficiency may be achieved in particular for larger cell loss rates.

The following procedure is proposed to determine a sufficiently large receiver buffer size and an appropriate frame

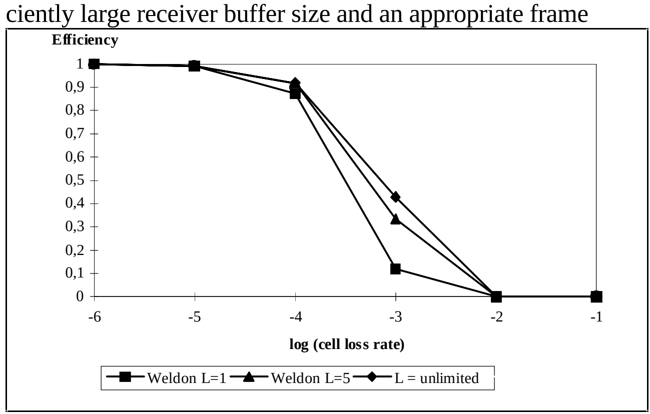

Figure 5: Efficiency for different buffer sizes

size k in respect to the cell loss probability q and the group size N.

Step 1: Define a desired efficiency η.

Step 2: Determine the frame size k for a given group size N and given cell loss rate q:

$$\eta_{max} = (1-q)^{k(N+1)} = \sum_{i=0}^{k(N+1)} (-1)^i \binom{k\cdot(N+1)}{i} q^i \qquad (13)$$

$$= \eta \le 1 - k\cdot(N+1)\cdot q + 0(q)$$

with 0(q) being the Landau-symbol, indicating that 0(q) will be infinitely small for q close to zero. Therefore, (13) can be approximated for small q by

$$\eta \le 1 - k\cdot(N+1)\cdot q \quad . \qquad (14)$$

Hence, the maximum frame size is given by

$$k \le \frac{1-\eta}{(N+1)\cdot q} \quad . \qquad (15)$$

Formula (15) shows that the frame size decreases with increasing group size N. This indicates that a frame-based error control scheme is not suitable for large groups. In this case, a cell-based ARQ scheme or a hybrid ARQ/FEC scheme is more suitable. If the frame size k is fixed, the maximum permissible cell loss rate is given by

$$q \le \frac{1-\eta}{(N+1)\cdot k} \quad . \qquad (16)$$

Step 3: Determine a proper receiver buffer size L. For the calculation, it will be demanded that the difference between the maximum and the minimum achievable efficiency for a given frame loss rate Q is smaller than $\Delta\eta$. For the ideal case η_{max}, the frame loss rate is given by (7).

$$\eta_{max} - \eta_{min} \le \Delta\eta \qquad (17)$$

$$\Rightarrow (1-Q) - \frac{1-Q}{1+S\cdot Q^{L+1}} \le \Delta\eta \quad . \qquad (18)$$

With Q=Q(N,k) according to (2), this allows to determine the required receiver buffer size to

$$L \ge \frac{1}{\log Q(N,k)} \cdot \log\left(\frac{\frac{\Delta\eta}{1-Q(N,k)}}{1-\frac{\Delta\eta}{1-Q(N,k)}} \right) - 1 \quad . \qquad (19)$$

4.1.4 Hybrid ARQ schemes

Generally, the frame-based schemes are appropriate for small group sizes, because the achievable efficiency decreases for both the number of receivers and the frame sizes. For large group sizes, the cell-based schemes show better performance properties. Therefore, the additional implementation complexity for RMC-AAL may be justified for large group sizes and significant cell loss.

The achievable performance of ARQ/FEC schemes can be evaluated as follows. First, the probability that m = i transmission attempts are necessary for successful delivery of a frame will be evaluated. In the frame-based hybrid ARQ/FEC scheme, k information cells are protected by h redundancy cells. The frame loss probability decreases from the frame loss probability of a pure ARQ scheme Q(k,k) to Q(n,k) for the hybrid scheme (with the length of a frame with redundancy being n = k + h cells):

$$P(m = i\cdot\frac{n}{k}) = (1-Q(n,k))^{N+1}\cdot(1-(1-Q(n,k))^{N-1})^{i-1} \qquad (20)$$

Hence, the mean value $\overline{m}$ can be derived to

$$\overline{m} = \sum_{i=1}^{\infty} i\cdot\frac{n}{k}\cdot P(m = i\cdot\frac{n}{k}) \qquad (21)$$

$$= \frac{n}{k}\cdot\frac{1}{(1-Q(n,k))^{N+1}} \quad , \qquad (22)$$

and the achievable efficiency is given by

$$\eta_{ARQ/FEC} = \frac{1}{\overline{m}} = \frac{k}{n}\cdot(1-Q(n,k))^{N+1} = \frac{k}{n} + 0(Q) \approx \frac{k}{n} \qquad (23)$$

Under this assumption, a formula for the general efficiency equilibrium of frame-based ARQ schemes and FEC schemes (see Figure 6) may be derived. The threshold for the cell loss probability is given by

$$\eta_{ARQ}(q_s) \approx \eta_{ARQ/FEC}(q_s) \quad , \qquad (24)$$

$$1 - k\cdot(N+1)\cdot q_s = \frac{k}{n} \quad , \qquad (25)$$

$$q_s = \frac{1-k/n}{k\cdot(N+1)} = \frac{h/n}{(n-h)\cdot(N+1)} \quad . \qquad (26)$$

Therefore, the following criterion is proposed to determine the appropriate error control mechanism for frame-based schemes:

$q \le q_s$ => frame-based ARQ,

$q > q_s$ => frame-based ARQ/FEC .

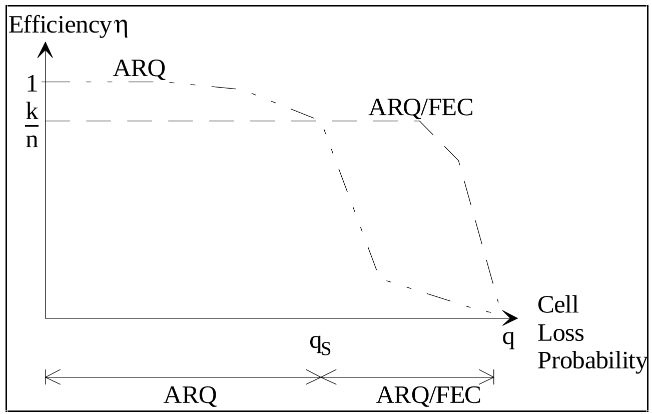

Figure 6: Threshold q_S for efficiency equilibrium

An interpolation of the simulation results for the frame-based schemes in a point-to-point scenario, in a multicast scenario with 4 receivers and in a multicast scenario with 16 receivers was performed. Table 1 shows the cell loss probability thresholds which were obtained. It can be seen that for point-to-point-communication, FEC is benefitial only for relatively high cell loss rates. With an increasing number of receivers, FEC is benefitial also for lower cell loss rates.

Number of receivers	Threshold (analysis)	Threshold (simulation)
1	1.82E-3	2.51E-3
4	3.64E-4	3.98E-4
16	1.07E-4	1.00 E-4

Table 1: Cell loss probability thresholds

5 Conclusion

It was pointed out that a large number of alternatives exist for the provision of a reliable multicast service in ATM networks. Existing approaches have significant shortcomings, however there does not exist a single approach best suited for all scenarios that need to be considered. A new framework was presented which has the potential to fulfil many requirements. For small groups and low cell loss rates, a frame-based end-to-end error control is most appropriate. In this case, RLMCP or another lightweight protocol for reliable multicast can be used 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.

Results of the performance evaluation allow the dimensioning of buffers and frame sizes. It is shown how the number of receivers affects the appropriate frame size and the cell loss threshold for which a cell-based scheme outperforms a frame-based scheme. Thresholds are also given to identify cases in which an ARQ/FEC allows higher efficiency than a pure ARQ scheme.

Acknowledgements

The author would like to thank Martina Zitterbart, Torsten Braun, Stefan Dresler, Jochen Schiller, Burkhard Stiller and the other members of the high-performance networking group for valuable discussions. Special thanks are to Axel Westenweller for implementing the simulation model. The author is grateful for provision of the simulation tool BONeS Designer by the Alta Group of Cadence Design Systems. The support by the Graduiertenkolleg "Controllability of Complex Systems" (DFG Vo287/5-2) is also gratefully acknowledged.

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

[Bie92] Biersack, E. W.: „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

[BrZ92] Braun, T.; Zitterbart, M.: „Parallel Transport System Design,“ 4th IFIP Conference on High Performance Networking HPN’92, Liège, Belgium, Dec. 1992

[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

[GoJa84] Gopal, I., Jaffe, J.: „Point-to-Multipoint Communication Over Broadcast Links,“ IEEE Trans. Commun., Vol. Com-32, No. 9, pp. 1034-1044, Sep. 1984

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

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

[InMo94] Inoue, I., Morita, N.: „Partial-Frame Retransmission Scheme for Data Communication Error Recovery in B-ISDN“, Proceedings of 4th IFIP Workshop on Protocols for High-Speed Networks, pp. 328-348, Vancouver, Canada, Aug. 1994

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

[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, pp.1471-1483, Dec. 1991

[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

[Sab82] Sabnani, K.: „Multidestination Protocols for Satellite Broadcast Channels,“ Ph. D. Dissertation, Columbia University, NY, U.S.A., 1982

[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

[ScB93] Schiller, J.; Braun, T.: „VLSI-Implementation Architecture for Parallel Transport Protocols,“ IEEE Workshop on VLSI in Communications, Stanford Sierra Camp, Lake Tahoe, CA, U.S.A., Sep. 1993

[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

[Wat92] Waters, A. G.: „Multicast Provision for High Speed Networks,“ 4th IFIP Conference on High Performance Networking HPN’92, Liège, Belgium, Dec. 1992

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

[ZST93] Zitterbart, M.; Stiller, B.; Tantawy, A.: „A Model for Flexible High Performance Communication Subsystems,“ IEEE Journal on Selected Areas in Communications, Vol. 11, No. 4, pp. 507-517, May 1993