A CELL-LEVEL FORWARD ERROR CORRECTION SCHEME (FEC-SSCS) FOR ATM NETWORKS

Georg Carle[1]
Institut Eurecom
BP 193
F-06904 Sophia Antipolis Cedex,
France.
g.carle@ieee.org

Aloke Guha
Storage Technology Corporation,
2270 S. 88th St.
Loisville,
CO 80028-5289, U.S.A.
aloke_guha@stortek.com

Tim Dwight
MCI Telecommunication Corporation
R901 International Parkway, Richardson,
Texas, 75081, U.S.A.

Keiji Tsunoda,
Kumiko Kanai,
Hiroshi Esaki
Toshiba Corporation, Japan.
tsunoda@csl.rdc.toshiaba.co.jp
kumiko@isl.rdc.toshiba.co.jp
hiroshi@wide.toshiba.co.jp

ABSTRACT

This paper proposes and describes a cell-level forward error correction (FEC) scheme in the ATM Adaptation Layer (AAL) for ATM networks. An FEC scheme will benefit real-time needs of multimedia applications as well as reliable (error-free) data communications. ATM network performance degrades rapidly with increasing cell loss or the bit error rate. The achievable throughput for reliable services also deteriorates with increasing packet or frame size or with increasing number of receivers in a multicast operation. FEC-SSCS can improve end-to-end QoS in all these cases.

1. Introduction

This paper proposes and describes an AAL-level forward error correction (FEC) scheme for reliable data transmission services in ATM networks to both correct bit errors and recover from cell losses.

In ATM networks, the basic data unit is very small compared to typical application level data units (e.g., IP packets in the Internet protocol suite). One data unit at the application level corresponds to a large number, 10s to 1000s, of cells. It is expected [1] that when either the cell loss ratio (CLR) is high or the packet length is large, the peer-to-peer throughput will be severely limited.

As shown in [2] and [3], the loss rate of higher layer packets, such as TCP packets, grows linearly with the number of cells comprising a packet. In [4] it was shown that the average response time of IPX and TCP protocol also degrades rapidly with increasing cell loss. The problem grows more serious for a large scale reliable multicast service, as required in interactive conference or games over the Internet. In addition, in some environments, due to the lack of error correction capability in the lower layer services (primarily at the physical layer), it cannot be expected that the bit error ratio will be sufficiently small.

This paper therefore proposes an AAL-level FEC scheme for the SSCS termed FEC Service Specific Convergence Sublayer (FEC-SSCS) for AAL type 5. Section 2 presents the rationale for providing an AAL-level FEC scheme. Section 3 describes the FEC-SSCS specification. Section 4 provides some brief performance evaluation results. Finally, section 5 outlines our conclusions.

2. Rationale for an AAL-Level FEC Scheme

2.1 The Need for an AAL-level FEC Scheme

An AAL-level FEC scheme will be required when reliable data transmission with bounds on performance is necessary. Some important requirements for such a reliable data transmission service would include the following:

1. end-to-end data transmission throughput;
2. end-to-end data transmission latency;
3. data transmission reliability (i.e., robustness) which depends on the cell loss ratio and the bit error rate (BER);
4. data transmission costs, which will be determined by the bandwidth, the duration of call, and/or by the number of transmitted cells;
5. protocol processing costs in end systems;
6. scalability for large distances and a large number of receivers.

The transmission performance in terms of throughput and latency will degrade with increasing cell loss ratio and the BER in ATM networks. The effective cell loss ratio and BER will depend on the following:

- The physical layer that is used
  Some media may have a larger BER than the physical layers currently specified. For example, in wireless LANs, tolerance of BERs higher than those for physical layers of wired media allows a significant reduction in implementation costs.
- The class of service and associated QoS desired
  In ATM networks, QoS parameters associated with the cell loss ratio may be negotiable, and will also depend on the service class. For some service classes (e.g., UBR), a cell loss ratio sufficiently low to achieve a satisfactory application level QoS may not always be possible.
- The congestion status of the network
  During the occurrence of congestion in ATM networks, buffer overflows will result in cell losses. For certain applications, especially mission critical applications, it is

---

[1] From October 1997, Georg Carle is with GMD Fokus, Kaiser-Augusta-Allee 31, 10589 Berlin, Germany

desirable to limit the degradation of service quality even in the event of congestion.

Reducing the effective BER and cell loss ratio for the upper layer process (e.g., IP) will significantly improve the overall quality of service. This is especially true for applications that require the provision of highly reliable services. A number of previous researchers [5,6,7] have shown that the use of FEC will improve end-to-end ATM performance in terms of effective throughput and latency. In addition, it was shown [8] that an FEC scheme may be used advantageously for multiplexing VCs with different QoS requirements. If data streams with and without redundancy are multiplexed, different QoS requirements can be satisfied even for a switch that does not distinguish the data streams.

2.2 Performance Degradation due to BER

The currently defined AAL types for data transfer (i.e., AAL5 and AAL3/4) do not provide error correction. They perform error detection only, and rely on the error correction capability of the transport layer protocol (e.g., TCP). When error correction is performed by TCP or any other transport layer protocol, complete AAL PDUs are discarded if a bit error or a cell loss is detected. The error correction capability of the transport layer typically relies on sender retransmission.

In general, the transport layer (or the network layer such as IP) does not have a cell-based error correction capability. Therefore, the complete packet must be retransmitted even if the received packet has only a single bit in error. In go-back-n schemes, not only the packet in error, but a full transmission window has to be retransmitted. For example, with a data unit of 65,535 bytes (the maximum data unit size of AAL5), the probability that the received data unit has a bit error is approximately $5*10^{-4}$, for uncorrelated errors when the bit error rate (BER) is 10^{-9}. For packets of 9,180 bytes (the default MTU size defined in [9]), the resulting packet error rate is approximately 10^{-4}. Moreover, in the case of a reliable multicast service, the packet error probability due to bit errors will linearly increase with the number of receivers. This means that for a large scale reliable multicast service, it is difficult to provide a service with satisfactory throughput and latency performance without using an FEC scheme.

2.3 Performance Degradation due to Cell Loss

In the ATM layer, the issue of cell loss due to congestion-related buffer overflows must be considered. Since the transmission service data unit, the AAL-SDU (e.g., an IP packet), will be segmented into multiple cells, a complete packet is assumed to be in error even when only one cell within the received packet is erased or missed. This packet must then be retransmitted. For the default IP MTU size of 9,180 bytes defined in [9], the corresponding AAL5-PDU consists of 192 cells. The maximum size of an AAL5 CPCS-PDU is 1,366 cells for a payload of 65,535 bytes. The approximate packet error probability due to cell loss is provided in [1]. For the example of a 64 Kbyte data unit (i.e., the maximum data unit size of AAL5), the probability that the received data unit is erroneous is about $1.3*10^{-3}$, when the cell loss ratio is 10^{-6}. For 8 Kbyte data units (i.e., the page size of modern file system), the error probability is about $2*10^{-4}$. Moreover, in the reliable multicast service, the expected packet error probability due to cell loss will linearly increase with the number of receivers.

2.4 Comparison with Other FEC Schemes

2.4.1 ITU-T SG15 Video Expert Group FEC Scheme

ITU-T SG15 video expert group is developing an AAL-level FEC scheme for the transmission of an audio-visual data streams (e.g., MPEG or H.261) over ATM networks. The target FEC scheme is for high quality real-time audio-visual signals (continuous bit streams). In addition, the FEC scheme to be developed should be independent of the AAL type and the physical media.

1. It is assumed that the ATM network provides a sufficiently small cell loss ratio (e.g., 10^{-9}) for the transmission of audio-visual signals, but can not provide a sufficiently small BER (e.g., > 10^{-7}). Audio-visual signal transmissions will use a CBR service so that the expected cell loss ratio is sufficiently small. However, it cannot always be expected that the BER is small enough, e.g., due to the poor BER of the physical media. Therefore, it is assumed that the FEC scheme corrects bit errors, although it may not correct cell losses.
2. The transmission of audio-visual signals that ITU-T SG15 is addressing imposes stricter latency requirements (i.e., delay and delay jitter), than those required for general data transmission.
3. A residual error rate can be tolerated, as dictated by human interface requirements. The error rate in the transmitted data at the receiver's application level must be sufficiently small, but this does not warrant error-free data transmission.
4. The optimization of the transmission performance of audio-visual data does not assume a transport protocol with error control mechanisms.

The following provides a comparison between our proposed FEC scheme (FEC-SSCS) and the FEC scheme developed by ITU-T SG15.

- FEC-SSCS is designed to provide efficient data communication even in the case where the ATM network can not guarantee a sufficiently low cell loss ratio. In contrast, the FEC scheme developed by ITU-T SG15 assumes that the ATM network provides a sufficiently small cell loss ratio.
- The latency requirement for the general data transmission is generally less strict than that for audio-visual signals considered by the ITU-T SG15.
- In data transmission with FEC-SSCS, we assume that the application requires an error-free data transmission.
- In order to optimize the data transmission by the transport layer, the FEC-SSCS interacts with the transport layer entity. A transport layer entity is generally required to provide an error-free data transmission.
- While a real-time audio-visual signal is a continuous bit stream, the data transmission we consider is generally not a continuous bit stream but rather an asynchronous data

stream. The data source is assumed asynchronous and can produce short (e.g., less than 50 bytes) to long (e.g., 64 Kbytes) packets.

As discussed above, the FEC scheme developed by ITU-T SG15 video expert group and the FEC-SSCS proposed in this paper are not conflicting schemes, but rather complement each other.

2.4.2 FEC Scheme at Application Level

The integration of an FEC scheme with the application level appears to be an alternative solution to obtain a high quality data transmission. In these approaches, FEC processing is performed at the application level, regardless of whether the end-station is attached to a conventional internetwork or to an ATM network. This approach appears reasonable for communication over a heterogeneous Internet environment that includes many types of data link standards including ATM.

However, even when the application applies an FEC scheme, there are still benefits for having an AAL-level FEC control capability in ATM networks for the following reasons.

- High throughput communication

  When the application requires a high throughput communication (e.g., at OC-3 speed), it will be difficult to achieve a high throughput when the FEC processing is performed at the application level. In contrast, an AAL-level FEC scheme can provide significant benefits when high-quality high-throughput communication is required. The implementation of an FEC-SSCS scheme can be optimized for a specific network adaptor and operating system to perform the cell-level functions. It is possible to extend existing VLSI-based AAL implementations to include FEC-SSCS, thereby offering a high-performance cell-based FEC implementation to some (or all) applications. In comparison, an application-level FEC scheme is typically (i) available only for a single application, (ii) implemented in software, and (iii) operates on larger data units (i.e., packets or frames).
- Retransmission latency

  The FEC-SSCS can detect and indicate the loss of or error in the received data to the application through the transport protocol entity much faster than if the FEC processing is performed by the application. Therefore, when the application is sensitive to quality degradation due to the latency of retransmission for lost or erroneous data, an FEC scheme performed by the application level is not appropriate.
- Utilization of redundancy

  An AAL-level FEC scheme can achieve a better cell loss correction capability than an application-level FEC scheme. In most protocol stacks that reside above AAL5, the a single scheme must be used for loss detection. In contrast, an AAL-level FEC scheme can use cell sequence numbers for cell loss detection, enabling to use the redundancy completely for recovering of lost cells.
- High quality pipe among routers

  The router decides whether to use FEC-SSCS to transfer the data flow over an ATM cloud. Even when a resource reservation protocol such as RSVP is used, the router can drop packets. Therefore, to recover from such packet loss, an FEC scheme may be used by the application process. However, since routers are usually entrusted to provide a secure and high-quality ATM pipe to transfer (IP) packets, they may choose to use a resource reservation protocol. In such a case, an application-level FEC scheme and an AAL-level FEC scheme (i.e., FEC-SSCS) will co-exist.

In summary, even when the end-station uses an application-level FEC scheme, we can still obtain some significant benefits by using an AAL-level FEC scheme (i.e., FEC-SSCS).

2.5 Interaction with Frame-Based Control Schemes

FEC-SSCS has many goals. Some goals are shared by other proposed mechanisms. This section briefly explores the relationship between FEC-SSCS and a class of such mechanisms we term "frame-based" due to their objective of operating on entire frames rather than individual ATM cells.

1. Cells passing over associated connection(s) are the result of segmentation of larger "frames".
2. The boundaries of such frames can be determined by examining the cell headers.

When these conditions hold, frame-based control mechanisms have been shown [2] to improve application layer throughput by reducing the transport of "dead" cells; and by reducing the total number of frames discarded in times of congestion. The former is achieved by Partial Packet Discard (PPD) and is invoked after a switch drops a cell, while the latter is achieved by Early Packet Discard (EPD) which drop complete frames when the onset of congestion is indicated by the queue occupancy.

Both FEC-SSCS and the various frame-based control mechanisms, are or would be, optional features selected on a per-connection basis.

Frame-based control mechanisms and FEC-SSCS can constructively co-exist, but that they do have an impact on each other, and this interaction requires further study. In particular, the following issues have to be considered:

- FEC-SSCS is able to recover a CPCS-PDU if a few cells are lost (for example, due to buffer overflow caused by congestion). For connections utilizing frame-based controls, however, the likelihood is higher that either no cells are discarded, or most (PPD) or all (EPD) are discarded. FEC-SSCS is thus likely to be able to recover fewer frames on a connection utilizing frame-based controls.
- Frame-based controls may be negatively affected by the use of segmentation at the FEC-SSCS layer. For example, consider a packet split in half by FEC-SSCS. This results in 2 CPCS-PDUs, which are discernible "frames" to the frame-based control mechanism. Should such controls be activated, they may discard one of the CPCS-

PDUs but not the other. The CPCS-PDU which successfully reached the destination may nonetheless be discarded by the application (e.g., due to packet-level checksum failure). In this case, the frame-based control's objective of reducing the transport of "dead" cells was partly negated by the use of FEC-SSCS segmentation.

3. Specification of FEC-SSCS for AAL Type 5

3.1 Requirements and Goals for an FEC-SSCS

Before presenting the FEC-SSCS specification, we list the requirements and goals for an AAL-level FEC scheme.

1. FEC-SSCS should be compatible with the specification of existing AAL Type 5, i.e., be compatible with the current SAR/CPCS structure of AAL Type 5.
2. It should be possible to adjust and negotiate the parameters of the FEC algorithm. These parameters should be negotiable during a session, as well as at the connection establishment phase. Here, the parameters of the FEC processing would be the size of the appended redundant information and the maximum size of cells in the FEC frame. The possibility of parameter negotiation allows the optimization of the transmission efficiency (i.e., minimize redundant data transmission for a given effective packet error rate), and to achieve media/service independence.
3. Adjustment of actually transmitted redundant data should be allowed for the source entity. The amount of redundant data is negotiated by the FEC parameter negotiation procedure. However, the source entity can decide the amount of redundant data that is actually transmitted to the destination entity. In other words, the FEC parameter negotiation procedure only specifies the maximum length of redundant data to be transmitted.

   This operation is completely localized in the source entity, and requires no negotiation between the source and destination entities. The information that the source entity must indicate to the destination entity is the amount of redundant data actually sent to the destination entity. This information is explicitly indicated either by the LI (Length Indicator) of CPCS trailer or by the Number of Redundant Data field in the CPCS-UU of the CPCS trailer.

   Using this operation, we could optimize the data transmission efficiency by minimizing the actually transmitted redundant data length, without the FEC parameter negotiation procedure. In summary, the benefits of this operation are as follows:

   (a) Reduction of redundant data transmission overhead

   When the end-station transmits variable length packets, e.g., from 100 byte to 64 Kbyte, this feature will be very beneficial. The size of the appended redundant data is proportional to the packet size. This is especially advantageous when small packets are to be transmitted and the FEC parameters can not be renegotiated.

   (b) Avoiding FEC parameter renegotiation overhead

   In some cases, the FEC parameter renegotiation procedure may be costly. In terms of latency and throughput, the use of this operation will reduce the frequency of FEC parameter renegotiation.

   (c) Economical implementation

   Some interface cards may support only one (or few) FEC parameter set(s). This would result in an economical (cheap) FEC-SSCS implementation. We expect that the implementation of this operation would be easier and simpler than the implementation with FEC parameter renegotiation.
4. Variable length AAL-SDUs (e.g., IP packets) should be supported
5. It should be possible to segment large AAL-SDUs into several smaller CPCS-PDUs, and to protect individual CPCS-PDUs using the proposed FEC scheme. The source FEC-SSCS entity can transfer FEC-SSCS-PDUs by pipelining. Through pipelining, the buffer resources in the source FEC-SSCS entity can be used more efficiently, and the transmission latency at the source FEC-SSCS can be reduced. Both the pipelining of AAL-SDU transmission and the pipelining of FEC frame transmission should be supported to reduce the peer-to-peer delivery latency.

   The fragmentation of an AAL-SDU into multiple FEC-SSCS-PDUs will be performed for transmission of the large AAL-SDUs. In pipelining mode, the received FEC-SSCS-PDU will be transmitted to the upper layer before the complete AAL-SDU is received by the FEC-SSCS entity. This operation mode is similar to the streaming mode defined in [10]. The benefits of such a pipelining mode are as follows:

   (a) Reduction of required receiving buffer space

   Some end-stations may not be able to allocate a sufficiently large buffer space to receive a large AAL-SDU. This may occur when the product of number of VCCs and the maximum AAL-SDU size is very large. By use of pipelining that results in AAL-SDU fragmentation, the buffer space to be allocated per VCC can be reduced.

   (b) Reduction of control latency

   Since the error in the received AAL-SDU (e.g., an IP packet) will be indicated by the pipelined FEC-SSCS entity earlier than in the non-pipelined case, the upper layer protocol entity can perform the end-to-end control (e.g., fast retransmission) earlier. This would improve the resulting end-to-end latency and throughput performance.
6. The processing costs of the FEC scheme in the lossless case (i.e., when no cell loss or bit errors occur) should be as small as possible. When there is no loss or error in the CPCS-PDUs carrying the user data, no FEC decoding should be required. Similarly, the reordering of data should be avoided in the lossless case.
7. The FEC method should operate in three modes without increased implementation complexity.

   (a) Error recovery in the case of cell loss and bit errors;

   (b) Error recovery only in the case of cell loss;

   (c) Error recovery only in the case of bit errors.

8. To support transport layer protocols like TCP with a rate control scheme that is based on the detection of losses in the network, it should be possible to indicate to the higher layer whether the transmission service data unit has experienced congestion and cell loss.

3.2 Service Provided by FEC-SSCS

The FEC-SSCS provides the capability to transfer AAL-SDUs from a source AAL-SAP to a destination AAL-SAP through the ATM network.

When a destination CPCS entity detects an error, which could be either bit error or cell loss, the FEC-SSCS entity tries to recover the original data sent from the source CPCS entity using the FEC algorithm specified.

As required, the FEC-SSCS has three operational modes associated with three different cases of error recovery:

1. SEC (Symbol Error Correction) Mode
   In the SEC mode, the FEC-SSCS entity attempts to recover only the bit errors in the symbols (data units defined by the FEC algorithm). Erased or missing symbols or cells are not recovered. When symbols are erased due to cell discard in the network, the received frame will not be recovered.
2. SLC (Symbol Loss Correction) Mode
   In the SLC mode, the FEC-SSCS entity attempts to recover only the erased or missing symbols. Bit errors are not recovered. When the recovery of erased or missing symbols fails, either the originally received data and an error indication is delivered to the upper layer entity, or nothing is delivered to the upper layer entity.
3. SEAL (Symbol Error And Loss Correction) Mode
   In the SEAL mode, the FEC-SSCS attempts to recover both the bit errors in symbols and the erased or missing cells.

The basic data format and required algorithm can be common in the three operational modes. The combination of functions for the individual modes differ.

- A CRC-10 field is used for error detection of the portions of the encoded FEC-SSCS-SDU. The CRC-10 calculation (attaching CRC-10) at the sender side can be skipped in SLC mode, while the bit error checking using CRC-10 at the receiver can be skipped in SLC mode, since bit error correction are not necessary in SLC mode.
- A Reed-Solomon (RS) error recovery algorithm (see Section 3.5) is used for error correction of an erroneous FEC-SSCS-SDU. In the RS error recovery algorithm, it is assumed that only missing symbols are subject to error recovery. The RS algorithm does not perform bit error recovery. In SEAL and SEC modes, the cells that contain symbol errors are detected by CRC-10. In the SLC mode, only the positions of the erased or missing symbols are detected by a sequence number field. Recovering both cells with symbol errors and lost cells, the complexity of the Reed-Solomon error correction algorithm can be reduced.

As a result, the core functions to be used by the three operation modes can be common. Figure 1 shows the referenced protocol structure associated with the FEC-SSCS. Table 1 shows the primitives shared between the FEC-SSCS and the upper layer entity (e.g., a TCP/IP entity).

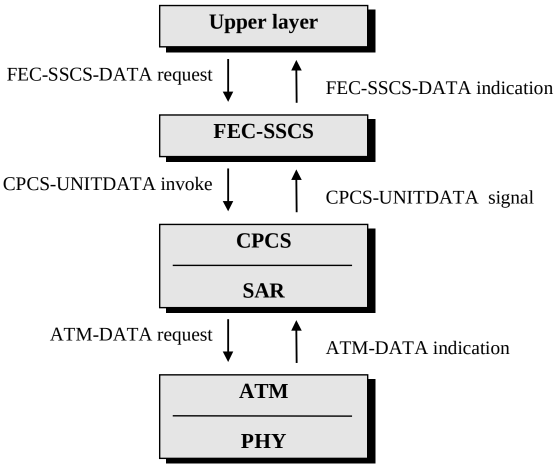

Figure 1: Protocol Reference Model for FEC-SSCS

3.3 Interaction with Management and Control Plane

The parameters of the FEC algorithm can be negotiated both during the FEC-SSCS connection set-up phase and during the duration of a connection. The negotiation during a connection must be done before the transmission of a burst using the inband signaling message generated by the upper layer. The inband signaling message is identified by the FEC-SSCS entity using the primitive FEC-SSCS-SIGI (see Table 1).

Before the signaling messages are exchanged between peer FEC-SSCS entities, the peer entities have to confirm that both entities process FEC-SSCS functionality. This is made possible by using the SSCS protocol identifier subfield in Q.2931. However, the transmission of signaling messages is not protected by the FEC algorithm.

Table 2 shows the primitives necessary for inband signaling for FEC parameter negotiation.

3.4 Structure and Coding of FEC-SSCS

The FEC-SSCS has two components. The functions of these components are described below.

- Sender side: creation of the FEC frame
  FEC-SSCS at the sender side creates an FEC frame that includes redundant information necessary to recover from errors at the destination FEC-SSCS entity. As shown in Figure 3, the FEC frame has a two dimensional matrix structure that separately interleaves the user and redundant FEC code in the matrix. A horizontal line in this interleaved frame corresponds to an FEC coded block. Each vertical line of data contains a header field to identify user and code data.
- Receiving side: error correction and error handling
  Errors in the FEC-SSCS-PDU are detected using the error detection capability of the CPCS entity. When the received FEC-SSCS-PDU contains bit errors, the FEC-SSCS entity tries to recover the error using the FEC algorithm. Otherwise, the FEC-SSCS entity transfers the received FEC-SSCS-PDU while deleting the redundant

information of the FEC-SSCS-PDU. If the FEC algorithm can not correctly recover the original FEC-SSCS-SDU sent by the source, the FEC-SSCS-SDU is either discarded or is optionally delivered to the upper layer entity with an error indication.

3.4.1 FEC-SSCS-PDU Structure and Coding

Figures 2, 3 and 4 show the proposed data structure model associated with FEC-SSCS, for user data packet transmission. For inband signaling, such as FEC parameter negotiation, the FEC code part in Figure 3 will not be appended, since the signaling message will not be protected by the FEC scheme.

The FEC-SSCS header field, shown in Figure 2, is attached to the FEC-SSCS-SDU (e.g., an IP packet). The FEC-SSCS header field includes an FEC-SSCS-UU (FEC-SSCS-User-to-User) field that is transparently transferred to the destination FEC-SSCS entity. The one bit MI (Message Identifier) field is mapped from the FEC-SSCS-SIGI that indicates the message type in the FEC-SSCS-SDU. Thus MI indicates whether the received FEC-SSCS-SDU is a user data packet or a signaling packet.

The FEC-SSCS-SDU does not have padding, length indicator, or error detection field. 48-byte alignment is provided using the functionality of the CPCS entity. One FEC-SSCS-SDU may correspond to one IP packet, and a burst of FEC-SSCS-SDUs may correspond to bulk data transmission.

An FEC-SSCS (plus the FEC-SSCS-SDU header field) usually generates multiple FEC frames. When the FEC frame is generated, the FEC code portion and the header field are appended to user data part, as shown in Figure 3.

3.4.2 Frame Format of FEC Frame

The format of an FEC frame that is actually transmitted from the source entity to the destination is shown in Figure 3.

The FEC frame has two parts: the user data part and the FEC code part. Each part has a FEC-frame-header field. The user data part contains the part or all of the FEC-SSCS-SDU with an FEC-frame-header field. The FEC code part contains an FEC redundant information field and a header field.

- User data field

  This field contains a part or all of the FEC-SSCS-SDU. The length in vertical direction **n** is determined so that the length of the vertical line in the FEC frame, {(t,1),(t,2),(t,3),...,(t,n)} is 46/a octets, where **a** is a positive integer that is usually one. In other words, the length of a vertical line in an FEC frame is usually 48 octets, equal to the payload length of an AAL5 SAR-PDU. The maximum length of a horizontal line in the user data field is specified by the MHL (see Table 2) parameter in the signaling message. This maximum length must be less than [2^{SYL}-FCPL] symbols, where SYL is the symbol size in bits and FCPL is the horizontal length of FEC code part.

- FEC code field

  This field contains the appended FEC code information to recover from bit errors and cell loss in the user data part. The data in the FEC field is calculated based on an a priori defined FEC algorithm with negotiated parameters (i.e., SYL and FCPL). The FEC algorithm used defines the code vector, {(M+1,t),(M+2,t),...,(M+S,t)}. The maximum length of an FEC code vector, expressed in **S** symbol lengths, is specified by the FCPL parameter in the signaling message.

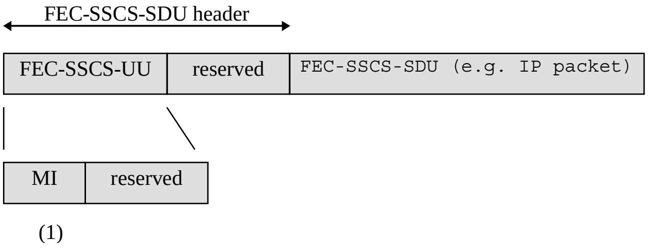

Figure 2: Frame Format of FEC-SSCS

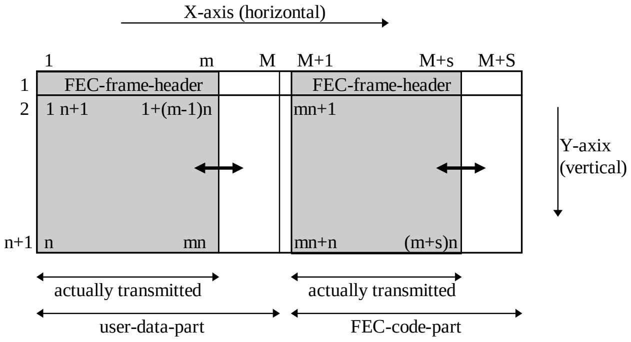

Figure 3: Frame Format of FEC frame

- FEC-frame-header field

  This field is used to identify the position of erased or erroneous vertical lines and to detect bit errors in vertical lines. The length of the FEC-frame-header is 2 octets in each vertical line. The detailed format and coding rule is specified in Section 3.4.3.

The order of writing and reading are the same to avoid the latency in reordering of symbols at the destination FEC-SSCS. Each horizontal line (i.e., the horizontal vector) in the user data part corresponds to the user data protected by the FEC code that is located at the same vertical position in the FEC matrix. **m** symbols {(k,1),(k,2),...,(k,M)} of user data are protected by the *s* symbols {(k,M+1),(k,M+2), ...,(k,M+s)} ($s \le S$) of FEC code.

3.4.3 Format and Coding of FEC-Frame-Header Field

The function of the FEC-frame-header field is as followed.

1. Identification of the positions of erroneous or erased vertical lines

   In order to recover the erroneous or erased symbols in the user data part, the FEC algorithm requires determination of the exact position of vertical lines with bit errors or erased symbols.

2. Error detection in the vertical lines

   In order to improve the error correction capability of the FEC algorithm, the frame-level FEC algorithm requires the detection of vertical lines with bit errors by additional per-cell redundancy in the FEC-frame-header. In SLC mode, in which bit errors are not corrected, this function is not required.

In order to provide the above functions, the following fields and associated coding rules are proposed. All fields exist in each vertical line. The length of the FEC-frame-header field in the vertical line is 2 octets.

- Sequence Number (SN) field

  The SN field is used to identify the position of the vertical line in the horizontal direction (between 0 and m+s-1). For a SN field with length **q**, the maximum number of vertical lines in the FEC code part denoted by FCPL is determined by 2q ≥ FCPL.
- User/FEC (U/F) field

  The U/F field is used to identify which part, i.e., user data or FEC code part, the received vertical line belongs to. The one bit U/F takes two values. "U" represents the user data part, and "F" represents the FEC code part.
- Parity (P) field

  The P field is used to perform the parity check on the FEC-frame-header field excluding the CRC field. The purpose of the P bit is to detect a single bit error in the FEC-frame-header.
- CRC field

  A CRC field of ten bits is used to detect bit errors in a vertical line.

3.4.4 FEC Frame Mapping to CPCS-PDU

In order to avoid large padding fields in the CPCS-SDU for the transmission of FEC code fields, the following FEC frame transmission method is used. If the FEC frame (aligned to 48 octets) is transmitted as a single CPCS-SDU, each CPCS-SDU uses a 40 octets padding field.

3.5 FEC Algorithm

In the FEC frame specified in Section 3.4.2, one Reed-Solomon block (RS-block) corresponds to one horizontal line. Therefore, the FEC frame contains *n* RS-blocks.

The FEC uses a Reed-Solomon (M, M+S) code that is able to correct up to [S/2] erroneous symbols or **a** erased (missing) symbols. With **m** being the number of actually transmitted vertical lines for the user data portion, and **s** being the number of vertical lines for the FEC code portion in the encoded FEC frame, **m** and **s** are limited by m ≤ M and s ≤ S. **m+s** vertical lines are actually transmitted from the source entity to the destination entity.

The RS(M,M+S) code can correctly recover data when the following equation is satisfied.

$$a + 2b \leq S \qquad (1)$$

where **a** is the number of erased (missing) symbols, and **b** is the number of symbols in error, whose position in the FEC frame is not known. Since the FEC frame uses CRC for bit error detection in each vertical line, the actual error correction capability of the FEC algorithm based on RS(M,M+S) is increased. Since the redundancy is used only to correct erased (missing) symbols, symbols in a vertical line are assumed to be erased when bit errors occurred in the corresponding to the line. Then, the actual error correction capability of the FEC algorithm is given by the following equation.

$$a + b \leq S \qquad (2)$$

If the total number of vertical lines with bit errors is greater than **s** (s ≤ S), then the RS code can be used for bit error detection in every row of symbols. In this case, the FEC frame can be recovered as long as not more than s/2 symbols per row are in error. Also, the receiving entity of FEC-SSCS entity can use the error correction capability of CRC-10 to correct a single bit error per vertical line.

RS codes to be used are built over a Galois Field $GF(2^N)$ [11,12], and are based on a generator polynomial G(X) given by

$$G(X) = \prod_{i=0}^{S-1} (X - \alpha^{i+k}) \qquad (3)$$

where α is a root of the primitive polynomial in $GF(2^N)$ and **k** is the base exponent of the generator polynomial.

The symbol length is determined by (i) feasibility of implementation, and (ii) the allowable maximum size of the encoded block (M + S). For a feasible implementation, the symbol length should be 4, 8, 16, 32 or 64 bits. The allowable size of the FEC frame (i.e. user data part plus FEC code part) is given by Table 3. For the current IP (IPv4), 4-bit symbols or 8-bit symbols would be preferable. In a native ATM environment, 8-bit or 16-bit symbols would be preferable. 16-bit symbols allow to cover the maximum SDU size in AAL Type 5 in a single FEC frame, while 8 bit symbols can cover the default MTU for AAL5 in a single FEC frame.

4. Performance Evaluation

4.1 Point-to-Point Communication

This section covers the influence of FEC-SSCS onto the packet error probability. Applying FEC-SSCS results in a significant reduction of packet loss. Figure 3 shows the improvement in the packet error probability for uncorrelated cell losses. The evaluation is based on IP packets that are segmented into **M** ATM cells. The ATM cells are considered erased or in error at each data link segment with probability β. It is further assumed that VCs consist of **d** data link segments. Without FEC-SSCS, an IP packet is corrupted whenever at least one cell contains a bit error or is lost. Applying FEC-SSCS with a single redundancy cell per frame, the IP packet is segmented into one or more FEC-SSCS-SDUs, and each FEC frame is segmented into **f** ATM cells (**f-1** cells for the user data part, and one cell for the FEC code part). Now, the IP packet is corrupted when at least a single FEC frame in the IP packet is corrupted. The evaluation of figure 3 is based on **f** = 10 and **d** = 5.

Applying FEC-SSCS substantially improves the IP packet error probability. For an IP packet size of 1330 cells, the IP packet error probability without FEC will be large (e.g., 0.6 with $\beta = 10^{-6}$). On the other hand, the IP packet error probability is substantially improved by the use of FEC scheme, e.g., $2.99*10^{-8}$ with $\beta = 10^{-6}$. In this example, the IP packet error probability is improved by five to six order of magnitude.

4.2 Point-to-Multipoint Communication

In this section, the packet error probability for point-to-multipoint communication is evaluated. With the evaluation model in the previous section, IP multicast packets are transferred from one sender to **N** receivers.

Figure 4 shows the improvement in the packet error probability. Applying FEC-SSCS again results in a significant improvement. For example, the IP packet error probability without FEC-SSCS is larger than 0.5 for 10^3 receivers with IP packets of 160 cells. With FEC-SSCS, the IP packet error probability is less than 10^{-5} for 10^3 receivers, and less than 10^{-2} for 10^6 receivers.

5. Conclusion

This paper discusses and proposes a cell-level forward error correction (FEC) scheme for the Adaptation Layer of ATM networks.

The proposed AAL-level FEC scheme can improve the end-to-end QoS both for point-to-point and point-to-multipoint communications. Furthermore, the FEC-SSCS scheme is very beneficial in cases where retransmissions by a higher layer protocol leads to low QoS and an in efficient use of network resources, or where delay requirements preclude the use of retransmission. By evaluation of the cell level and packet level error probability, a substantial improvement through the use of the FEC-SSCS scheme was shown. The cell loss probability can be improved by one or more orders of magnitude, even when the cell losses are correlated. The packet error probability can be further improved, in some cases even to five or six orders of magnitude.

REFERENCES

[1] H. Esaki: "Reliable IP Multicast Communication Over ATM Networks Using Forward Error Correction Policy", To appear in IEICE Transactions on Communications.

[2] A. Romanow and S. Floyd: "Dynamics of TCP Traffic over ATM Networks," IEEE JSAC, Vol. 13, No. 4, May 1995, pp. 633-641.

[3] T. Chen, L. Jones, S. Liu, and V. K. Samalm: "Effect of ATM Cell Loss on TCP Packet Loss"; ATM Forum technical contribution 94-0914, 1994.

[4] H. Li, K.-Y. Siu, H.-Y. Tzeng: "IPX and TCP Performance over ATM Networks with Cell Loss", ATM Forum technical contribution 95-0151, Feb. 1995.

[5] N. Shacham, P. McKenny, "Packet Recovery in high-speed networks using FEC coding," in Proceedings of IEEE INFOCOM '90, San Francisco, CA, June 1990.

[6] H. Ohta and T. Kitami, "A Cell Loss Recovery Method Using FEC in ATM Networks," IEEE JSAC, vol. 9, pp. 1471-1483, Dec. 1991.

[7] E. Ayanoglu, R.D. Gitlin, and N.C. Oguz, "Performance Improvement in Broadband Networks Using Forward Error Correction," Journal of High Speed Networks, vol. 2, 1993, pp. 287-304.

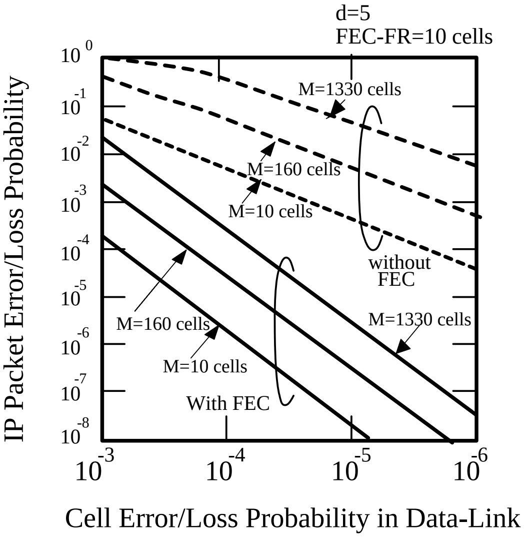

Figure 3: Packet Error Probability for Point-to-Point Communications

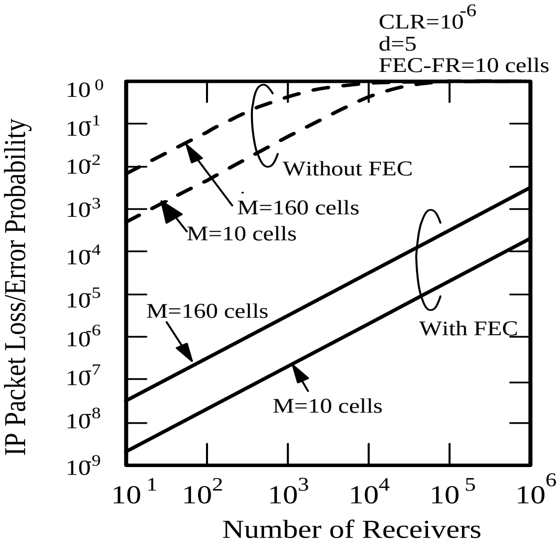

Figure 4: Packet Error Probabilty for Point-to-Multipoint Communications

[8] E.W. Biersack: "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.

[9] R. Atkinson: "Default IP MTU for use over ATM AAL5", RFC1626, May, 1994.

[10] ITU-T Recommendation I.363: "B-ISDN ATM Adaptation Layer (AAL) Specification", March, 1993.

[11] S. Lin, D.J. Costllo: "Error Control Coding ; Fundamentals and Applications", Prentice Hall, 1983.

[12] A.J. McAuley: "Reliable Broadband Communication Using a Burst Erasure Correcting Code", ACM SIGCOMM90, Sept. 1990.

Parameter	Type	Comments
Interface Data (ID)	request indication	whole or partial FEC-SSCS-SDU
More (FEC-SSCS-M)	request indication	FEC-SSCS-M=0: end of FEC-SSCS-SDU FEC-SSCS-M=1: not end of FEC-SSCS-SDU
Burst indication (FEC-SSCS-BI)	request indication	FEC-SSCS-BI=0: end of burst FEC-SSCS-BI=1: not end of burst
FEC-SSCS-Loss Priority (FEC-SSCS-LP)	request indication	- invoke: FEC-SSCS-LP=0 (high priority) - signal: Mapped from the ATM layer's congestion indication parameter
FEC-SSCS-Congestion indication (FEC-SSCS-CI)	request indication	Mapped to from the ATM layer's congestion indication parameter FEC-SSCS-CI=1: congestion experienced FEC-SSCS-CI=0: no congestion experienced
FEC-SSCS Cell Loss indication (FEC-SSCS-CLI)	request indication	
FEC-SSCS Bit Error indication (FEC-SSCS-BEI)	request indication	
FEC-SSCS Signaling indication (FEC-SSCS-SIGI)	request indication	FEC-SSCS-SIGI=1: signaling data FEC-SSCS-SIGI=0: not signaling data
FEC-SSCS-User-to-User indication (FEC-SSCS-UU)	request indication	Transparently transported by the FEC-SSCS
FEC-SSCS Reception Status (FEC-SSCS-RS)	request indication	Indication of corrupted FEC-SSCS-SDU

Table 1: Primitives and parameters between FEC-SSCS and the upper layer

Parameter	Type	Comments
Symbol Length indication (SYL)	request indication	8xn bits (n>0) are specified
Vertical Length indication (VL)	request indication	Usually indicating 46 byte length
Maximum Length of Horizontal Line (MHL)	request indication	Must be less than [2^(8n)-FCPL]
FEC-Code-Part Length indication (FCPL)	request indication	2^(SNL) must be larger than the maximum length of FEC-code-part. (*)
Selection of Operation Mode (SOM)	request indication	SOM=1 : Cell Loss is not recovered SOM=2 : Bit Error is not recovered SOM=3 : Both are recovered
Negotiation Status indication (NS)	request indication	NS=0h : hello polling NS=1h : hello positive acknowledgment NS=2h : hello negative acknowledgment NS=3h : polling parameters NS=4h : polling positive acknowledgment NS=5h : polling negative acknowledgment NS=6h : confirmation

(*): SNL is the length of sequence number field in the FEC-frame-header field.

Table 2: Primitives and parameters for FEC parameter negotiation

Symbol length	Max. of M+S	Max of User data + FEC code part
4 bits	15	690 Byte
8 bits	255	1.73 Kbyte
16 bits	65,535	3.015 Mbyte
32 bits	$4.30\text{x}10^{9}$	197.57 Gbyte
64 bits	$1.85\text{x}10^{19}$	$8.51\text{x}10^{20}$ Byte

Table 3: Symbol length and Maximum FEC Frame Size