RZ 2911 (#91506) 03/10/97
Engineering Technology 12 pages

Research Report

A SCALABLE SONET/SDH ATM TRANSCEIVER WITH INTEGRATED ADD/DROP AND CROSS-CONNECT

A. Herkersdorf, W. Lemppenau *, H.R. Schindler

IBM Research Division
Zurich Research Laboratory
8803 Rüschlikon
Switzerland

* now with FH Gelsenkirchen, Abteilung Bocholt, Postfach 1441,
D - 46364 Bocholt, Germany

LIMITED DISTRIBUTION NOTICE

This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents and will be distributed outside of IBM up to one year after the date indicated at the top of this page. In view of transfer of copyright to the outside publisher, its its distribution outside of IBM prior to publication should be limited to peer communications and to specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies of the article (e.g., payment of royalties).

IBM Research Division
Almaden, T.J. Watson, Tokyo, Zurich

A SCALABLE SONET/SDH ATM TRANSCEIVER WITH INTEGRATED ADD/DROP AND CROSS-CONNECT

A. Herkersdorf, W. Lemppenau *, H.R. Schindler

IBM Research Division, Zurich Research Laboratory, 8803 Rüschlikon, Switzerland

ABSTRACT: This paper presents the architecture and functions of SMART, a **S**calable and **M**odular **AR**chitecture for SONET/SDH **T**echnology. SMART performs the transmission convergence and part of the physical media-dependent sublayer functions between B-ISDN ATM and SONET/SDH. The basic building block of the architecture is the SMART macro. Such a macro performs the mapping and extraction of an ATM cell stream into and out of STM-1/STS-3c frames as defined in the ITU-T and Bellcore standards. A unique feature of SMART, its modular scalability, allows one to interconnect N identical macros to support STM-N and STM-Nc. This support requires that certain functions handled in a single macro at STM-1 speed be performed in a distributed fashion across the N macros for STM-N and STM-Nc. Four SMART macros are integrated on one chip. A single chip can be configured to map one or multiple ATM cell streams into the following SDH signals: 4x STS-1, 4x STM-1/STS-3c, 1x STS-3, 1x STM-4 or 1x STM-4c. SMART scalability is not limited to macro interconnection within a chip, but extends to multi-chip configurations. Four SMART chips can operate in parallel to support STM-16/STS-48 (16x OC-3 or 4x OC-12c). SMART scales through the SDH hierarchy levels by operating N macros in parallel, with each macro running at a clock rate of 19.44 MHz. This is the byte clock for STM-1 operation. Consequently, the cost of a SMART-based STM-N chip or system scales linearly with the number of macros N. The low macro clock rate permits high-density CMOS implementation, which leads to a low cost per macro and, thus, to a low cost per chip as well as per system.

Other important and unique features of SMART are its on-chip support for add-drop multiplex (ADM), digital cross-connect (DCS) and automatic protection switching functions inherent to SONET/SDH, and the integrated serial line interfaces. These interfaces directly connect to electro/optical modules up to 622 Mbit/s without the need for external serializer/deserializer components. The ATM interface is both UTOPIA level 1 and 2 compliant and supports all single- and multi-PHY options.

Public domain SONET/SDH equipment must fulfill stringent reliability and availability requirements, which leads to high costs. All features of SMART described above, when applied in the private domain with its more relaxed requirements, enable the low-cost implementation of ADM and/or DCS functions. These functions, when combined with ATM layer functions (segmentation/reassemly, label swapping, etc.), can now be integrated on a single adapter card of a CPN HUB. SMART technology opens the door to an extended use of SONET/SDH in private networks and, thus, will enable a converging and homogeneous transport infrastructure for the public WAN and the private CPN world.

1 INTRODUCTION

The Synchronous Digital Hierarchy (SDH) and, its North American equivalent, the Synchronous Optical Network (SONET) are the ITU-T and Bellcore standards accepted and applied worldwide for data transmission in the wide area network (WAN) domain. SONET/SDH represents the prevailing physical layer transport infrastructure for higher-layer services in the public environment. Traditionally, the predominantly deployed services over SONET/SDH have been and still are the circuit-switched T1/E1 (1.54/2.048 Mbit/s) or T3/E3(44.74/33.36 Mbit/s) connections. The transport of cell/packet switched services, e.g. ATM (Asynchronous Transfer Mode) or IP (Internet Protocol), are attracting increasing attention.

SONET/SDH define a uniform multiplex hierarchy covering bit rates from 51 Mbit/s (OC-1/STS-1), to 155 Mbit/s (OC-3/STM-1) and up to 9.9 Gbit/s (OC-192/STM-64). The global uniformity of the SONET/SDH standard enables multi-vendor network installations and competitive pricing policies among network element equipment providers. Compared to earlier data transmission technologies (e.g., the Plesiochronous Digital Hierarchy (PDH)), SONET/SDH provides significantly enhanced support for automated network operation, administration, and maintenance (OAM). The OAM functions include frame delineation, continuous connectivity supervision, bit-error monitoring, built-in data and voice communication channels, alarm indication signaling, automatic protection switching (APS) and more. Of special importance is APS, which allows network failures to be bypassed automatically within milliseconds by connection rearrangement. A key concept of SONET/SDH is the pointer-based access to user data streams encapsulated in the SONET/SDH signal. This access method is the basis for cost-efficient implementations of add-drop multiplex (ADM) or digital cross-connect systems (DCS). See [KL90] for an introductory article on SONET/SDH. More comprehensive descriptions are found in [SR92] or in the standard literature [G.707, G.783, GR.253] itself.

Today's private customer premises networks (CPN) are characterized by a diversity of physical layer architectures. Examples are Ethernet, Token Ring, FDDI, ESCON channels, and ATM. With the growing acceptance of multimedia services, a trend towards rapidly increasing network capacity requirements for CPN backbones is being observed. This evolution demands a reliable high-speed transport infrastructure. Currently, Gigabit-Ethernet and ATM switching HUBs with integrated IP support are candidates for such an infrastructure. On the other hand, in the WAN environment SONET/SDH is a proven, reliable technology for high-speed telecommunications. However, SONET/SDH network equipment, as it is available today, has been designed for the reliability, availability and serviceability (RAS) of public long-distance networks. Consequently, SONET/SDH based ADM or DCS systems are too expensive for the cost-sensitive private market. We believe that the price disadvantage of SONET/SDH can be overcome in the future. One indication is the adoption of SONET/SDH for the ATM UNI at OC-3 and above by The ATM Forum. ASIC manufacturers started to provide integrated SONET/SDH framer devices for mapping ATM cells into SONET/SDH frames. These framer devices are available today at moderate prices and

can be used for point-to-point WAN access at either OC-3 (155 Mbit/s) or OC-12c (622 Mbit/s) speeds.

This paper presents the architectural concepts of SMART, a joint development between the IBM Zurich Research Laboratory, the IBM Essonnes Development Laboratory, and the IBM T.J.Watson Research Center. SMART represents second-generation SONET/SDH framer ASIC technology with substantially enhanced flexibility and functionality. Whereas other SONET/SDH framer ASICs are either applicable only to STM-1/STS-3c [Bt8222, P5345, T1500] or to STM-1 and STM-4c for a single ATM stream [P5355, T1600], the scalability concept of SMART allows the selective mapping of up to four ATM cell streams either into four STS-1, four STM-1 frames, one STM-3, one STS-4, or one STM-4c frame with a single chip. Four SMART chips can operate in parallel to support one STM-16/OC-48 signal (16 x 155 Mbit/s or 4 x 622 Mbit/s). In addition, SMART provides on-chip support for add-drop, digital cross-connect, and automatic protection switching functions inherent to SONET/SDH. The integrated serial SONET/SDH line interfaces allow direct connectivity to electro/optical modules without the need for additional off-chip serializer/deserializer or clock recovery & synthesis components. All these features enable the low-cost integration of SONET/SDH ADM and/or DCS functions on ATM adapter cards within CPN nodes. No stand-alone, expensive, ADM/DCS systems are required to establish a SONET/SDH backbone ring in the CPN. Thus, SMART technology is an enabler for a converging transport infrastructure between the public WAN and the private CPN networking world.

2 SMART: A SCALABLE SONET/SDH ATM TRANSCEIVER

2.1 The Modular Scalability Concept

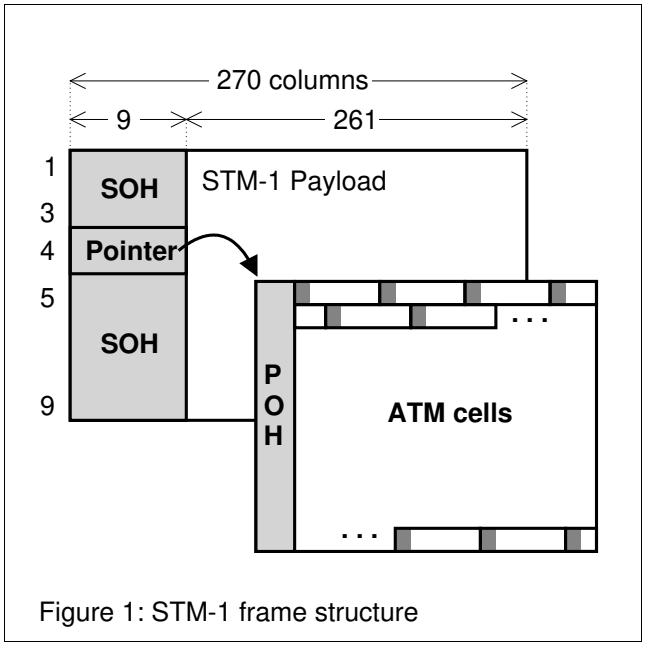

Figure 1: STM-1 frame structure

Figure 1 shows the STM-1/OC-3c frame structure. The STM-1 frame consists of 9 rows times 270 columns. Each column is one byte (8 bit) wide. Frames are periodically transmitted row-by-row as a serial data stream. The frame period is 125us, which results in a bit rate of 155.54 Mbit/s. The first 9 columns of each row contain the section overhead (SOH) and the payload pointer. The remaining 261 bytes of each row contain the STM-1 payload. Note that the STM-1 payload is also referred to as a VC-4 (virtual container of type 4) in the SDH standard literature. The STM-1 payload is partitioned into a one-column path overhead (POH) and the payload itself. The SOH and POH are the added-value of SONET/SDH compared to other transmission techniques. As the main application of

SMART is the transport of ATM traffic, Figure 1 shows ATM cells packed back-to-back in the STM-1 payload section.

According to ITU-T [G.707] and Bellcore [GR.253], higher-order STM-N frames, with N denoting the SDH hierarchy level, are generated by one-byte interleaving (byte multiplexing) N independent STM-1 payloads into an STM-N payload. Consequently, an STM-N frame contains N x 9 SOH columns, N POH columns and N x 260 user payload columns. Defined values for N are 1, 4, 16, and 64. A single user payload with a capacity greater than the base STM-1 payload capacity is supported in concatenation

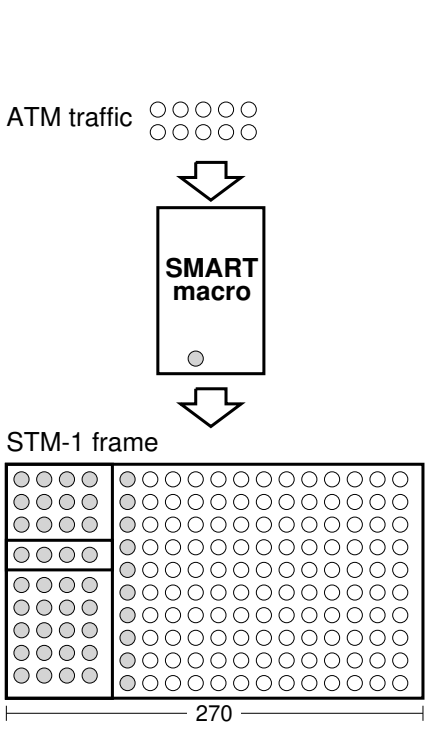

Figure 2: SMART macro

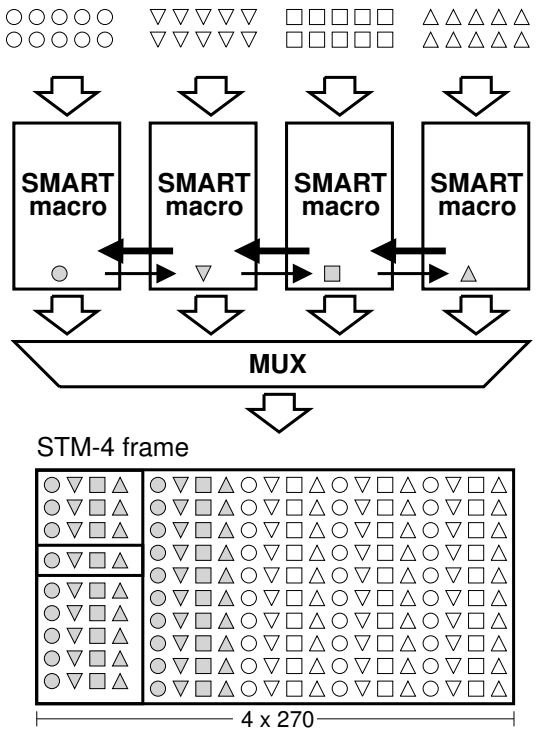

Figure 3: SMART scalability concept

mode. In this mode, N STM-1 payloads are concatenated to form a single STM-Nc payload. The STM-Nc payload contains one POH column, followed by N-1 stuffing columns and N x 260 user payload columns.

The regularity of the STM-N frame structure across all hierarchy levels and the one-byte interleaving principle have been incorporated into the SMART scalability concept: Our goal was to develop a modular building block, referred to henceforth as the SMART macro, that performs the STM-1 framing, overhead processing and transmission convergence sublayer (TCS) functions for an OC-3 ATM stream (see Figure 2). SMART scalability means that N paralleled SMART macros can interoperate to support STM-N and STM-Nc. In this mode, the STM-N framing, overhead processing, and TCS functions are performed in distributed fashion across N synchronized and interconnected macros. Figure 3 shows the example for STM-4,

where every macro contributes an OC-3 payload and a 9-bytes-per-row SOH share to the STM-N frame. The contributions per macro are illustrated with different geometric shapes, and the interoperation among macros is indicated by horizontal arrows. This distributed processing approach is highly attractive because all macros, independent of how many operate in parallel, run at the same clock rate of 19.44 MHz. This allows implementations that scale linearly in terms of cost (measured in number of 19.44 MHz clocked macros). Native SONET STS-N operation is supported by running the macros at one-third of the STM-1 clock (6.48 MHz), and by configuring the macros to process 9x90-byte frames (3-byte SOH). The multiplex stage, which merges the outputs of the N macros, performs a generic multiplex function with no SONET/SDH functionality. SMART is the only SONET/SDH framer reported in the literature that provides this flexible scalability.

Figure 4 shows the functional structure of one SMART macro in transmit and receive direction. ATM cells are received from the ATM layer via the UTOPIA interface [UL1, UL2] and buffered in a four-cell-deep FIFO (AB). The UTOPIA interface (UI) is not considered to be part of the macro because it will not be replicated N times at the chip level. The ATM cell handler (CH) performs the transmission convergence sublayer (TCS) functions (see chapter 2.3). The cells pass through a second buffer stage (SB), used for clock rate decoupling and support for extended functions (see chapter 2.4), to reach the SDH processor (SP). The SP provides the STM-N frame skeleton and performs all functions related to SOH and POH processing. The SDH frames are transmitted by the Line Interface (LI). In receive direction, the corresponding functions are performed in reverse order. The chip we have developed contains four such

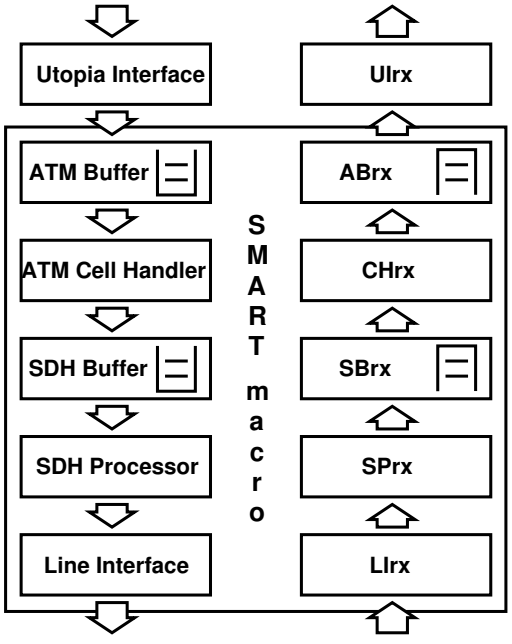

Figure 4: SMART macro structure

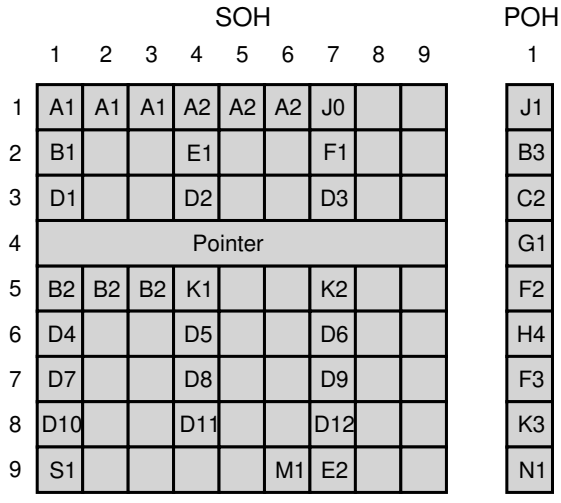

Figure 5: STM-1 SOH and POH

bidirectional macros. The following two sections discuss implementation-related aspects of the SDH Processor and the ATM Cell Handler in greater detail.

2.2 SDH Processor

The SDH processor deals with the functions related to framing, frame scrambling, section and path overhead byte generation and analysis, pointer processing, and mapping/extraction of the STM-N payloads. The data path of the SDH processor is 1 byte wide, resulting in a clock rate of 19.44 MHz. SDH processor operation is compliant to ITU-T [G.707, G.783] and Bellcore [GR.253].

2.2.1 Frame scrambling

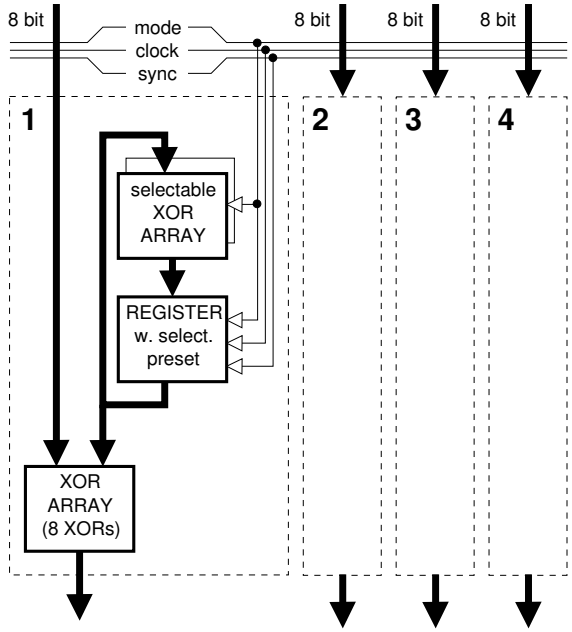

Figure 6: Multi-mode SDH frame scrambler

ITU [G.707] specifies that every SDH frame, with the exception of the first row of the SOH must be scrambled by a frame-synchronized scrambler. Frame-synchronizing means that the scrambler is set to a predefined value at the start of every frame. The generator polynomial is $1+x^6+x^7$. In the standard, the scrambler is described as a single-bit shift register with feedbacks via XORs. As shown in Figure 6, the scrambler in one SMART macro (STM-1) operates in parallel, based on byte-wide feedback to XOR arrays and proper initialization of the scrambling process. The scrambler generates a pseudo-random sequence that is independent of the data stream, which means that all feedbacks are completely within the sequence generator. This property permits an implementation of an STM-N frame scrambler on N parallel macros without any cross-coupling. The feedback pattern per macro is realized by selectable XOR arrays (controlled by mode) and selectable register preset values (controlled by mode and sync).

2.2.2 SOH Byte Processing

Figure 5 depicts the individual SOH bytes contained in the STM-1 frame of Figure 1.

Framing: The A1 and A2 byte are used for frame delineation. The frame alignment word of an STM-N frame consists of 3xN A1 bytes followed by 3xN A2 bytes. Multiplexing the three A1 and three A2 bytes independently generated by each of the N paralleled macros results in a correct STM-N frame alignment word. No macro interoperation is required. In receive direction, the LIM module, described in the chapter entitled "External Interfaces", performs byte and frame delineation.

Bit interleaved parity (BIP) processing: The B1 byte contains an 8-bit BIP checksum for error monitoring purposes. The B1 checksum is computed over all bytes of an STM-N frame. Transmitting equipment inserts the B1 code of the current frame in the B1 byte location of the next frame to be transmitted. Receiving equipment performs B1 calculation over the current frame and compares this B1 code with the B1 byte received in the next frame. Mismatches between computed and received B1 codes are counted.

To compute the B1 checksum of an entire STM-N frame, each byte of the frame has to be included in the computation. However, in a structure in which SOH processing is split into N paralleled macros, each macro sees only the N-th part of a frame. The solution is to perform B1 calculation and checking as a two-stage process: First, B1 is calculated locally within each macro resulting in N intermediate checksums. Second, the same B1 calculation is performed over the N intermediate checksums to obtain the final B1 checksum. This approach is valid because parity calculation is basically a modulo-2 addition operation, which is associative by nature. Furthermore, this approach is scalable because the time interval between the end of the current frame (the availability of the N intermediate checksums) and the transmission of the B1 checksum in the next frame (beginning of row 2 of the next frame) allows sequential adding of the intermediate B1 checksums to compute the final B1 value. In the case of STM-16, each of the 16 stage-two operations may take up to 800 ns (125us / (9 rows x16 modules)), which can be well handled by today's CMOS technology even when crossing chip boundaries.

The Nx3 B2 bytes contain a BIP-Nx24 checksum covering the entire STM-N frame except for the first three rows of the SOH. In principle, BIP-Nx24 computation is defined analogously to BIP-8 computation. The difference is that the B2 checksum is Nx24 bits wide and, hence, the computation is based on Nx24 bit (Nx3 byte) quantities of the STM-N frame. The advantage of a parallel macro structure is that the BIP-Nx24 checksum is equivalent to N one-byte interleaved BIP-24 checksums. Consequently, every SMART macros computes a BIP-24 checksum based on the data that is processed by the macros, and the N checksums are one-byte interleaved like the rest of the STM-N frame without the need to exchange B2 checksums among macros.

Remote error indication (M1): The M1 byte conveys the error count detected by the received B2 checksum back to the remote transmit side. As the B2 checksum is distributed over N macros, a dedicated macro interconnection is used in SMART to assemble the B2 error count in macro #1 from where it is inserted into the M1 byte of the current frame. Note that macro #1 is the leftmost macro in Figure 3. Receiver equipment counts the number of block errors extracted from the M1 bytes.

The processing of the following SOH bytes can be accomplished within each macro without the need for inter-macro communication. In the case of STM-N or STM-Nc, the bytes appear only once, namely at the location corresponding to macro #1, in the SOH. In particular, these functions are:

Section trace (J0): A 16-byte section access point identifier (SAPI) that can be stored in the transmit macro is repetitively transmitted in the J0 byte. The receiving equipment validates its continued connection to the intended transmitter by comparing the received J0 bytes against the expected SAPI (stored in the receiver macro).

Orderwire (E1,E2) and user channel (F1): Two orderwire bytes, each byte representing a 64 kbit/s voice channel, and one data byte channel may be used for special maintenance purposes. These bytes can be accessed by the microprocessor both in transmit and receive direction.

Data Communication channels (D1-D3, D4-D12): Two data communication channels (DCC) are defined and supported with dedicated serial interfaces (See chapter "External Interfaces").

Automatic protection switching (K1, K2): The K1 and K2 byte are allocated for automatic protection switching (APS) signaling and remote defect indication (RDI), see chapter 2.4 "Extended Functions". Detecting an RDI in the K2 byte results in a microprocessor interrupt.

Synchronization status (S1): The S1 byte contains the synchronization level (i.e., the clock quality) of the transmit equipment. The SPrx checks the received status pattern for consistency. Modifications in the synchronization status result in a microprocessor interrupt.

Bytes reserved for national use or future standardization (i.e., the blank byte locations in Figure 5), are filled with a default byte pattern. No dedicated support is provided for these bytes. However, they can be accessed by the microprocessor.

2.2.3 Pointer processing

Pointer processing is a key architectural feature of SONET/SDH. It allows one to decouple the STM-1 payloads from the STM-N frame phase, and to compensate for minor clock frequency variations between transmit and receive equipment of different network nodes. A detailed description of the pointer mechanism is beyond the scope of this paper, but can be found in [G.707, GR.253].

SMART fully supports SONET/SDH pointer processing. Every STM-1 payload has its own pointer. As every STM-1 payload is processed by a separate SMART macro, pointer processing can be accomplished entirely within the macro. There is no need for inter-macro communication, the exception being payload concatenation, STM-4c operation. An STM-4c payload is a single payload for which a single pointer is sufficient; it is provided by macro #1. The pointer bytes of the other macros are filled with concatenation information, a fix fill pattern defined in the standards.

2.2.4 POH Byte Processing

BIP processing (B3): The B3 byte contains a BIP-8 checksum covering all bytes of an STM-1 or STM-Nc payload. Except for the STM-Nc case, the STM-1 payloads are processed within a single macro. In the case of STM-Nc, the payload is distributed over N macros. Hence, B3 computation requires an interconnecting scheme that is similar to the one used for B1 computation.

The following POH bytes are processed within a macro without the need for inter-macro communication.

Path trace (J1): In analogy to the section trace byte J0, the J1 byte is used to transmit repeatedly a 16-byte or 64-byte access point identifier. SMART supports both path trace formats.

Signal label (C2): The C2 byte identifies the composition of the payload. The SPrx equipment compares the expected signal label with the received C2 byte. A signal label mismatch results in a microprocessor interrupt.

Path status (G1): The G1 byte conveys a 4-bit error count (B3 checksum errors) and a 1-bit RDI back to the remote path originator.

For the user channel bytes F2/ F3, the position indicator H4, the APS channel K3, and the network operator byte N1, no dedicated processing is provided. All bytes are accessible by the microprocessor.

2.3 ATM Cell Handler

The ATM Cell Handler supports all TCS functions for SONET/SDH as defined by ITU-T [I.432]. One SMART chip contains four ATM Cell Handlers to support four 155 Mbit/s data streams (8-bit data path) or one 622 Mbit/s concatenated data stream (32-bit data path). As an option, the cell handlers can be switched into a transparent data path mode. This option permits the use of external non-ATM frame payload mappers/ demappers (e.g., for IP, FDDI, T3/E3, etc.) via a dedicated CS interface, see chapter 2.5.

A four-cell-deep ATM Buffer (AB FIFO in Figure 4) provides clock separation between the ATM and the PHY layer. Whenever the buffer does not hold a complete user cell from the ATM layer, idle or unassigned cells are automatically inserted into the transmit path to guarantee a continuous cell stream towards the SDH Processor. On the receive side, idle/unassigned cells are either discarded or may be forwarded to the ATM layer. In transmit direction, online calculation of the header error control (HEC) byte, and, in receive direction, single-bit error forward correction and multi-bit error detection are performed. Self-synchronized ATM cell payload scrambling and descrambling [LEE94] is provided.

In the case of 32-bit cell handler operation, the 5-byte cell header is extended to 8 bytes. This results in a 56-byte cell, in which both header and payload are word (32-bit) aligned. Of course, these three extra bytes have to be removed before the cells are transferred to the SDH Processor for back-to-back alignment in the SDH frame. Correspondingly, three bytes are inserted (and removed by the UI) in the reverse direction. The removal/insertion process is performed with two 4x4 byte-wide switching matrices. These matrices are well justified because the word alignment of the cells simplifies the HEC calculation and cell scrambling significantly. Furthermore, as will be shown in chapter 2.4 “Extended Functions”, the matrix serves a second purpose, i.e., APS and loopback.

The cell handler supports a variety of test and diagnostic functions: Header configuration of idle/unassigned cells, cell sequence payload bit error monitoring and activation/deactivation of HEC calculation, HEC error forward correction and cell

payload scrambling/descrambling. A unique feature is HEC corruption performed on cells with selectable VPI/VCIs.

2.4 Extended Functions

2.4.1 Integrated Serial SDH Interfaces

An important feature of SMART is its on-chip serializer/deserializer functions for the SDH line interfaces. The serializer consists of an eigth-to-one bit multiplexer and a PLL-based clock generation circuit. The deserializer provides PLL-based clock extraction from the bit-serial data stream, byte delineation, one-to-eigth bit demultiplexing and, finally, SDH frame delineation. Three of the four serial interfaces support 155 Mbit/s, and one supports both 155 Mbit/s and 622 Mbit/s. As the serial interfaces can be connected directly to optical/electrical modules, no additional external serializer/deserializer components are required, which leads to considerable cost savings for system integration.

2.4.2Add/Drop Multiplex and Digital Cross-Connect

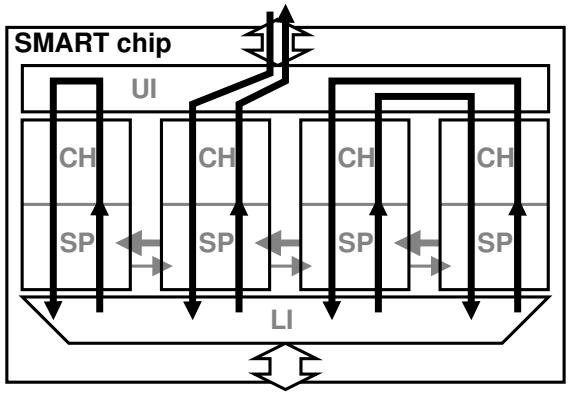

Figure 7: Add/Drop and Cross-Connect on a chip

So far, the description of SMART has been concentrated on its use for bidirectional SONET/SDH path termination for ATM traffic. However, SMART can also be configured to operate in various multiplex section transit modes, such as: native SONET/SDH multiplexing (4x STM-1 to/from 1x STM-4), SONET to SDH conversion (and vice versa), add/drop and digital cross-connect. ADM and DCS are important and frequently used functions to realize SONET/SDH-based ring or meshed networks. With ADM, a fraction (e.g., STM-1) of a high capacity data stream (e.g., STM-16) can be extracted economically and a new one can be inserted while the bulk traffic is bypassed. With DCS, input ports of a network node can be arbitrarily assigned to the output ports.

In STM-4c mode, the 16-bit data provided via the UTOPIA interface has to be demultiplexed and word-aligned to 32-bit format. This requires a byte-wide 2x4 switching matrix in transmit direction, and a 4x2 matrix in receive direction. If the two matrices are replaced by a single 6x6 matrix, it can alternatively be used to perform SDH multiplex, SONET to SDH conversion, ADM and DCS functions.

For example, configuring the receiver sections of the four macros of a SMART chip for STM-1 as well as the transmit sections for STM-4 operation, and folding the received STM-1 payloads within each macro back to the transmit direction by means of the 6x6 switch matrix results in a 4x STM-1 to STM-4 SDH multiplexer. The inverse operation works as well. SONET to SDH interworking is achieved by configuring the macros in one direction for SONET (Bellcore) and in the opposite direction for SDH (ITU). As the

ATM Cell Handler operation can be disabled (transparent data path mode), the multiplexers are not restricted to ATM traffic. This means that SMART provides a smooth migration path between emerging ATM and traditional (circuit switched) services.

As shown in Figure 7, DCS operation is realized by configuring the 6x6 matrix such that the STM-1 payload received from one macro is transmitted via a different macro. For ADM, the matrix is set such that some of the STM-1 payloads are terminated at the UTOPIA or CS interface, while the others are retransmitted i.e., bypassed. Note that a single SMART chip provides ADM and DCS functions for up to four STM-1/OC-3's out of one STM-4/OC-12 without the need for any external components.

2.4.3 Automatic Protection Switching

The protection switching operation, based on the K1/K2 byte signaling protocol, is defined in ITU-T [G.783]. All time-critical portions of the signaling protocol are handled in the SDH Processor (SP). For the switching of the STM-1 channels we have generalized the two 4x4 matrices used for byte removal/insertion in the ATM Cell Handler (CH) to one 8x8 matrix and extended it to 9x9. As shown in Figure 8, this now permits 1+1 and 1:N APS within one SMART chip (1<=N<=4), and even across multiple chips (4<N<=14).

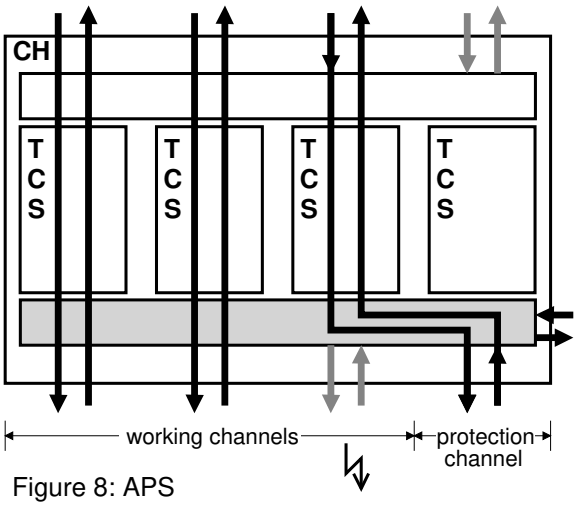

Figure 8: APS

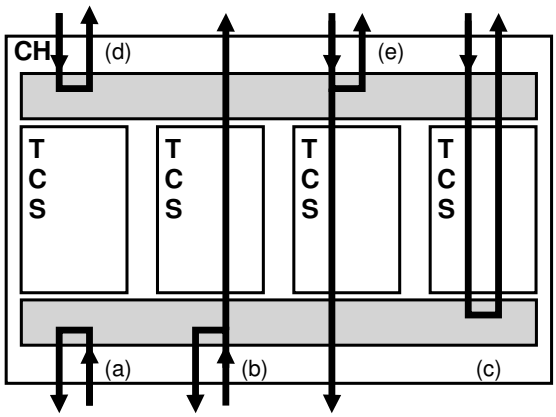

Figure 9: Loopbacks

The example in Figure 8 assumes a line failure on working channel three, and shows the switch-over to the protection channel.

2.4.4 Loopbacks

As shown in Figure 9, the matrix described above in the ATM Cell Handler can also be used for loopback purposes, for example, a line loopback without (a) or with (b) copy-on-the-fly, or a local loopback traversing the TCS functions (c). For fault isolation purposes, a second loopback stage (8x8 matrix) was inserted at the top of the TCS functions. It can be used for local loopbacks without (d) or with (e) copy-on-the-fly.

2.5 SMART External Interfaces

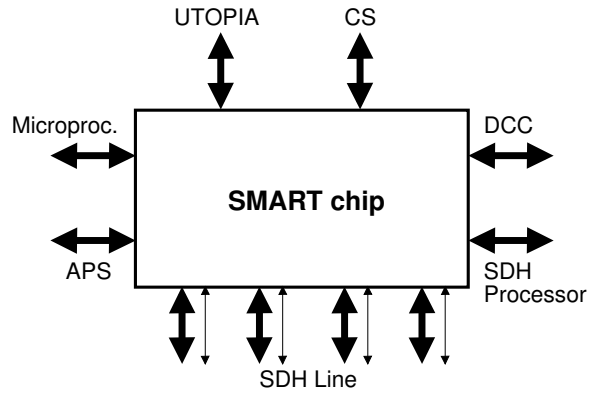

Figure 10: External Interfaces

The SMART chip communicates with its environment via the dedicated interfaces shown in Figure 10: The chip is configured and controlled via the microprocessor interface. This interface is also used to retrieve status information from the chip. External ATM layer devices are connected to the UTOPIA interface. The 16-bit UTOPIA interface can be configured to operate as one Utopia-II interface (supporting all single- and multi-PHY as well as status indication options up to 622 Mbit/s) or two independent 8-bit UTOPIA-I interfaces for 155 Mbit/s each. The circuit switched (CS) interface can be used for transparent data path channels. The CS interface supports one 622 Mbit/s or up to four 155 Mbit/s transparent channels. SMART can handle any combination of ATM and transparent channels. The APS interface interconnects multiple SMART chips for automatic protection switching. The SDH Processor interface exchanges the control signals to support STM-N processing distributed across N parallel macros. One DCC interface per SMART macro can be connected directly to external HDLC controllers for using the SONET/SDH inherent data communication channels. In addition to the serial line interfaces (introduced in chapter 2.4), SMART provides four 8-bit parallel SDH line interfaces. Again, three of the four interfaces support 155 Mbit/s, whereas one is designed for 155 Mbit/s and 622 Mbit/s. The parallel interfaces are useful for STM-16 operation, where the serial data stream is generated with a commercially available 2.5 Gbit/s serializer/deserializer component [VT95].

3 CONCLUSIONS AND OUTLOOK

We have developed a SMART chip which demonstrates that the regular multiplexing principle inherent to the SONET/SDH hierarchy can effectively be translated into a VLSI architecture. As a result, the SMART chip shows unmatched scalability and extended functions. A single chip can map one 622 Mbit/s or multiple 155 Mbit/s ATM cell streams into 4x STS-1, 4x STM-1/STS-3c, 1x STS-3, 1x STM-4 or 1x STM-4c. Four SMART chips support STM-16/STS-48. A SMART chip contains four basic macros, each running at a modest clock rate of 19.44 MHz. Applied to today's CMOS technology, this leads to low power dissipation and low cost. Unique features of SMART are its on-chip support for ADM, DCS and APS functions, and the integrated serial line interfaces.

All these features enable the low-cost integration of SONET/SDH ADM and/or DCS functions on ATM adapter cards within CPN nodes. No stand-alone, expensive, ADM/DCS systems are required to establish a SONET/SDH backbone ring in the CPN.

Thus, SMART technology is an enabler of a homogeneous transport infrastructure permitting seamless interconnection of public WANs and private CPN networks.

3.1 Acknowledgments

We would like to acknowledge the continuous support of and the many contributions to the design of the SMART chip by our colleagues from the Essonnes laboratory, J.-C. Le Garrec, V. Guerre, D. Louis, J.-P. Nuez, P. Mone, D. Plassat, A. Steimle, G. Taroni, from the T.J. Watson Research Center in Yorktown, J. Ewen, B. Parker, D. Pearson, A. Widmer, D. Zumbo, and from the IBM Zurich Research Laboratory, P. Buchmann, R. Clauberg, T. Engbersen, C. Landry, R. Luijten, and D. Webb.

References

[B8222] Brooktree, Bt8222, ATM Receiver/Transmitter with UTOPIA Interface

[G.707] ITU-T, 'Network Node Interface for the Synchronous Digital Hierarchy', (07/95)

[G.782] ITU-T, 'Types and Characteristics of Synchronous Digital Hierarchy (SDH) Equipment', (01/94)

[G.783] ITU-T, 'Characteristics of Synchronous Digital Hierarchy (SDH) Equipment Functional Blocks', (01/94)

[GR.253] Bellcore, GR-253-CORE, 'Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria', (12/94)

[I.432] ITU-T, 'B-ISDN UNI - Physical Layer Specification', (03/93)

[KL90] M.J. Klein, 'The Synchronous Digital Hierarchy: Principles, variants and applications', Philips Telecommunication Review, Vol.48, No.4, Dec 1990

[LEE94] B.G. Lee, S.C. Kim, "Scrambling Techniques for Digital Transmission", Springer-Verlag, London, 1994

[P5345] PMC-Sierra, PM5345, S/UNI-155 Saturn User Network Interface

[P5355] PMC-Sierra, PM5355, S/UNI-622 Saturn User Network Interface

[SR92] M. Sexton, A. Reid, 'Transmission Networking: SONET and the Synchronous Digital Hierarchy', Artech House, Norwood, MA, 1992

[T1500] Texas Instruments, TNETA1500, SONET/SDH ATM Receiver/Transmitter

[T1600] Texas Instruments, TNETA1600, SONET/SDH ATM Receiver/Transmitter

[UL1] ATM Forum, 'UTOPIA - An ATM PHY Interface Specification', Level 1, Version 2.01, (03/94)

[UL2] ATM Forum, 'UTOPIA - An ATM PHY Interface Specification', Level 2, Version 0.95, (06/95)