

Copyright © 1992, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

**A HIGH PERFORMANCE SMDS INTERFACE
AT STS-3c RATE**

by

My T Le, Nick McKeown, and Richard Edell

Memorandum No. UCB/ERL M92/89

27 August 1992

**A HIGH PERFORMANCE SMDS INTERFACE
AT STS-3c RATE**

by

My T Le, Nick McKeown, and Richard Edell

Memorandum No. UCB/ERL M92/89

27 August 1992

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

TITLE PAGE

**A HIGH PERFORMANCE SMDS INTERFACE
AT STS-3c RATE**

by

My T Le, Nick McKeown, and Richard Edell

Memorandum No. UCB/ERL M92/89

27 August 1992

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

A High Performance SMDS Interface at STS-3c Rate


Submitted for publication June 1992

My T Le Nick McKeown Richard Edell

Abstract

This paper describes the design of a high performance interface to the Switched Multi-megabit Service (SMDS). It was implemented as a part of The BayBridge and operates at STS-3c (155Mbps) and DS3 (45Mbps) rates. The SMDS Interface consists of two circuit-boards: the Segmentation and Reassembly (SAR) board and the SMDS board.

We first give an overview of SMDS and of implementation issues. After the overview, we discuss the architecture of the SMDS interface, starting with the *Segmentation & Reassembly* board. The SMDS board is described next in three parts: *Medium Access Control*, *Physical Layer Convergence Protocol*, and *Management*. To illustrate the design, we present an example starting from the initialization process through the operational and error recovery procedures. Finally, we draw the lessons from the design of the SMDS Interface.



The
Bay
Bridge

Project Report: 7 Revision: 1.0
Department of Electrical Engineering
and Computer Sciences
University of California at Berkeley

1 Overview

1.1 What is SMDS?

Switched Multi-megabit Data Service (SMDS) is proposed by *Bell Communications Research, Inc. (Bellcore)* as a means to connect Local Area Networks (LANs) beyond the customer premises, across a metropolitan or wide area [1]. A typical topology is shown in Figure 1.




Figure 1: SMDS Topology

The principal features of SMDS are:

High Speed: DS1 (1.5Mbps) or DS3 (45Mbps).

Connectionless: transmission is on a per-frame basis. No call setup or removal is required.

Public: services are offered by the Bell Operating Companies. Each subscriber is assigned a ten-digit address following the *prefix of 1* according to CCITT Recommendation E.164 and the North American Numbering Plan (NANP).

Customer Premises Equipment (CPE) is connected to a local SMDS *Switching System (SS)* (usually located at the local telephone office) via the *Subscriber-Network Interface (SNI)*. The interface protocol across the SNI is based on the Distributed Queue Dual Bus (DQDB) Metropolitan Area Network Medium Access Control (MAC) protocol described in the IEEE 802.6 Standard.

There are two ways to connect CPEs to SMDS, *Single-CPE* and *Multiple-CPE* (see Figure 2).

Single-CPE: One CPE connected to one SS.

Multiple-CPE: A small number of CPEs connected to one SS. The CPEs and SS will operate as a DQDB network.

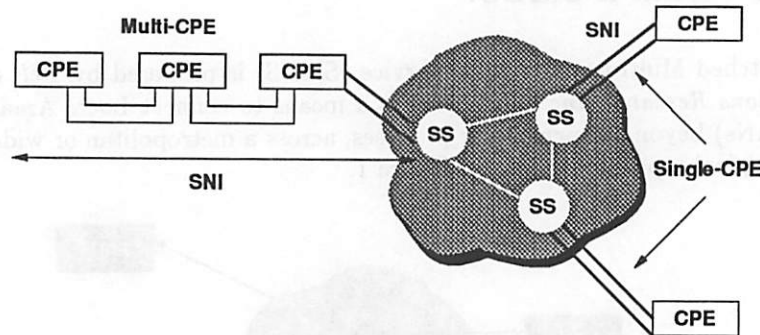


Figure 2: SMDS Configuration

1.2 Requirements to interface to SMDS

The units of transmission in SMDS are variable length messages, *L3 Protocol Data Unit (L3-PDU)* of up to 9188 bytes and fixed length cells, *L2 Protocol Data Unit (L2-PDU)* of 53 bytes. The cells are transmitted on to the medium after being arranged in a $125\mu\text{s}$ frame at the Physical Layer. The relationship among message, cells and frames is shown in Figure 3.

- Layer 3:** When SMDS is used to connect LANs across the metropolitan or wide-area, the packets from the LANs need to be *transformed* into SMDS message format. The transformation of LAN packets into SMDS messages can be implemented in two ways: *encapsulation* or *translation*. When encapsulation is being used, the entire LAN packet is placed in the payload of the SMDS message. When the LAN packet is translated, its header is stripped while its payload is placed in the SMDS message. Note that when translation is performed, the SMDS message may need to be *re-translated* back to a LAN packet once it reaches the destination CPE, unless the CPE is directly connected to the host.
- Layer 2:** After the message is constructed, it needs to be segmented into multiple cells before the data is transmitted on to the Physical Layer. Similarly, at the destination CPE, cells need to be reassembled into a message.
- Physical Layer:** Cells are placed in the $125\mu\text{s}$ frame structure before being sent out to the network. The frame structure is different depending on

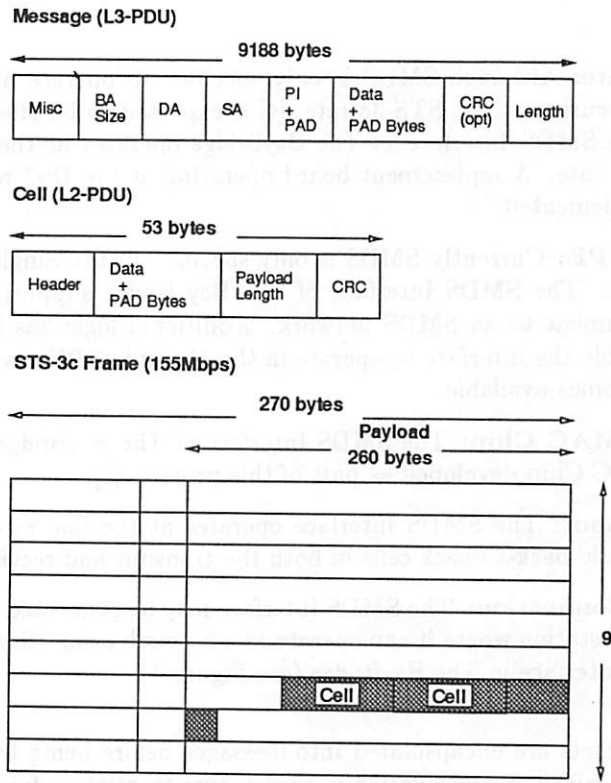


Figure 3: Message, Cell and 125µs Frame

the data rate being used (DS1 or DS3). Thus, the method of placing cells in the frame varies with each data rate.

1.3 SMDS Interface of The BayBridge

The BayBridge [2, 3] is designed to connect FDDI LANs via SMDS as an internetworking bridge and router. The main features of the SMDS Interface are:

Data Rate: Although SMDS is only specified to operate at DS1 or DS3 at the current time, STS-3c rate [4] is expected to be offered in the future. The SMDS Interface of The BayBridge operates at the SONET STS-3c line rate. A replacement board operating at the DS3 rate has also been implemented.

Single-CPE: Currently SMDS is only specified in the Single-CPE configuration. The SMDS Interface of The Bay Bridge supports Single-CPE attachment to an SMDS network. Additional logic has been designed to enable the interface to operate in the Multiple-CPE environment when it becomes available.

DQDB MAC Chip: The SMDS Interface of The BayBridge uses the DQDB MAC Chip developed as part of this project [5].

Throughput: The SMDS Interface operates at the line rates and is able to handle back-to-back cells in both the transmit and receive directions.

Direct Connection: The SMDS Interface may be connected directly to a host workstation where it can operate as a network connection instead of being an interface in The BayBridge (see Figure 4).

FDDI packets are encapsulated into messages before being transmitted to the SMDS Interface. Encapsulation is chosen over translation for its ease of implementation. The drawback of encapsulation is that some of the message payload is used to store the FDDI packet header, thus introducing another level of overhead. The encapsulation process is performed by the *Protocol Converter* of the *Bridge Board* described in [2] (see Figure 5).

The SMDS Interface is implemented to perform the following functions:

- **Segmentation & Reassembly:** The process of segmenting a message into multiple cells and the reverse process of reassembling multiple cells back to a message are performed by the *Segmentation & Reassembly Board*.

- *Medium Access Protocol*: The DQDB protocol is used as the method to gain access to transmit cells. The protocol is especially important in the Multiple-CPE configuration where a number of CPEs and the SS form a DQDB subnetwork. The DQDB protocol is implemented in the DQDB MAC Chip on the *SMDS Board*. On the transmit side, the MAC transmits cells from the SAR Board when it gains the right to transmit. On the receive side, the MAC copies cells intended for the CPE and transmits them to the SAR Board.
- *Physical Layer Convergence Protocol (PLCP)*: The PLCP logic on the *SMDS Board* is responsible for placing cells in the payload of the $125\mu\text{s}$ frame. In the reverse direction, it is also responsible for extracting cells from the frame on the receive side. Aside from placing and extracting cells, the PLCP also places and extracts the *Payload Overhead (POH) bytes*. The POH bytes are used to implement the management functions. The PLCP is connected to the Framer [6] which is responsible for turning the $125\mu\text{s}$ frame into bits before transmitting them out to the network. As the Framer and the optical components were not developed as part of The BayBridge project, they will not be discussed in this paper.
- *Management Protocol*: The Management logic on the *SMDS Board* is responsible for transmitting and receiving the POH bytes. Management information at the SMDS Interface includes a number of Error and Alarm Indications as well as the *Bus Identification Protocol* and *Message Identifier Page Allocation Protocol* required by the DQDB protocol.

The major components of the SMDS Interface of The BayBridge are shown in Figure 6.

2 Segmentation & Reassembly Board

The SAR Board operates synchronously to the SMDS line clock, segmenting outgoing messages into cells and reassembling incoming cells into messages.

2.1 Segmentation

When the Bridge Board has an SMDS message to transmit over the SMDS network it hands the message to the SAR Board. Messages are segmented into 44-byte cells with 9 bytes of header and trailer information added (see Figure 7). The segmentation process is straight forward, as it involves only one message at a time.

2.2 Reassembly

While Segmentation only deals with one message at a time, the Reassembly logic must be able to handle cells from different messages arriving *interleaved*. The SMDS standard currently requires that cells be accepted when up to 16 interleaved messages arrive, but is expected to increase the requirement to 128 in the future. The SAR Board was designed to accept up to 128 interleaved messages. In addition and as an experiment, the SAR Board was designed to accept *mis-ordered* cells. Mis-ordered cells are normally rejected in SMDS, the Reassembly design of The BayBridge was implemented as an experiment of fast reassembly with cell mis-ordering. There are two issues in designing the Reassembly logic: efficient buffer management and timely cell processing. The simplest but least efficient Reassembly protocol is to always allocate a buffer equal to the length of a maximum SMDS message when a BOM or SSM cell arrives. The drawback of this protocol is the inefficient use of buffer space. While 1.2Mbytes of memory must be allocated for 128 maximum messages, a large portion of the buffer will not be used as most messages will not be of the maximum size. This means that small back-to-back messages will quickly clog the reassembly buffer.

The most complex scheme is to allocate buffers of exactly the size of the SMDS message. The drawback of this scheme is the ability to make timely cell processing decision when cells arrive every $2.8\mu s$ (at STS-3c rate).

2.2.1 Implemented Algorithm

The algorithm used in The BayBridge is based on [7] and is summarized in Figure 8. Efficient use of the cell buffer is achieved by allocating buffers cell by cell. When a BOM or SSM cell arrives, a new entry is made in the *MID Table* and a new *Message Table* entry is removed from the free list. The cell is placed in the next free location in the *Cell Buffer* and its address stored in the *Message Table*. When COM and EOM cells arrive, the MID is used to index into the *MID Table* to find the correct *Message Table* entry for this MID value. The cell is placed in the next free location in the *Cell Buffer* and its address stored in the *Message Table*. The position of the address in the *Message Table* is calculated using the Cell Sequence Number. Reassembled messages are removed from the cell buffer by indexing through the *Message Table* and reading from the *Cell Buffer*.

In this design, completed messages may be read from the cell buffer at the cell arrival rate. The advantages of this design are:

- The cell buffer is used efficiently. This becomes more important as the cell buffer size increases. This scheme is most appropriate for a switch or for a CPE serving a large number of users.
- Calculating the address of the next cell position in the cell buffer is simple — it is the next location removed from the free list. As a result new cells may be accepted continuously back-to-back, unless the Message Table or Cell Buffer is full.

However, this scheme is only worth using if the size of the Message Table is a small proportion of the size of the Cell Buffer. This may not be the case in some applications. The algorithm used in the SMDS Interface of The BayBridge is not well suited for SS implementations because of the overhead of the Message Table.

2.2.2 An alternative algorithm

We now describe an alternative reassembly design where hardware implementation is necessary (CPU is not fast enough) and where cell re-ordering is not required (see Figure 9).

- Reassembly is comprised of three components: MID Controller, FIFO Controller, and a shared RAM.
- When a cell is received, the *Cell Type* is checked. If the cell is a *SSM* or *BOM*, the FIFO Controller allocates a free FIFO to store the cell. The address of the FIFO is written in the Register File as the FIFO Write Pointer. After the first cell of the message is received, the following *COM* cells will be stored in the same FIFO. The Write Pointer is updated every time a new cell of a message is received. Once the *EOM* cell is received, the FIFO is ready to be read out.
- The Host will read from the FIFO, returning the FIFO to the *Free List*.
- The MID Controller is used to store the FIFO Number of the FIFOs that are currently being written. The *Next Expected Cell Sequence* is also kept in the MID Controller.
- To ensure that the FIFOs are efficiently used, FIFOs may be allocated in a small number of different sizes, some capable of storing maximum message size (9188 bytes) while other are capable of holding smaller messages. Upon receiving the first cell of the message, the MID Controller checks the *BAsize* to determine which FIFO is best suited to store the message.

3 SMDS Board

The SMDS Board is comprised of three logic blocks: MAC, PLCP, and Management.

3.1 Medium Access Control

The Medium Access Control (MAC) is responsible for transmitting and receiving cells (L2-PDUs). The structure of a cell is shown in Figure 10.

The MAC functions include transmitting outgoing cells and receiving incoming cells. Each MAC is connected to one bus; thus two MACs are required per CPE. An overview of the DQDB MAC chip is shown in Figure 11.

Outgoing cells:

Cells are handed to the SMDS Board by the SAR Board. When the SAR Board requests a transmission, the DQDB MAC will wait for an appropriate empty cell. When access is granted by the MAC, the SAR Board passes a cell to the SMDS board. The CRC checksum is calculated by the MAC and added to the cell trailer.

In the *Single-CPE Configuration*, the MAC at the CPE has full use of the bandwidth. Whenever an empty cell is received, the MAC can immediately use the cell, thus the DQDB protocol is not necessary. On the other hand, in the *Multiple-CPE Configuration*, the full DQDB protocol needs to be implemented as a number of CPEs will be competing to use the available bandwidth.

Incoming cells:

The DQDB MAC accepts all incoming BOM (Beginning of Message) cells that match any external MAC address or cells with Message Identifiers (MIDs) corresponding to messages in progress. The MAC maintains a table of current MIDs and alerts the Host of exceptions. The MAC does not check for correct cell ordering as this is handled by the SAR Board. The MAC checks the CRC of all incoming cells and hands cells directly to the SAR Board for reassembly, indicating if the cell has passed the CRC.

3.2 Physical Layer Convergence Protocol

The Physical Layer Convergence Protocol (PLCP) is responsible for placing cells in the payload of the $125\mu\text{s}$ frame at the transmitting side once the beginning of the frame has been detected by the Framers. Conversely, it is also responsible for extracting cells from the frame payload at the receiving side. Each PLCP is connected to one direction of each bus (see Figure 12). PLCP1 is connected to the receive side of Bus A and the transmit side of Bus B. Similarly, PLCP2 is connected to the receive side of Bus B and the transmit side of Bus A. In the Single-CPE configuration, the CPE receives from Bus A and transmits to Bus B, hence only one PLCP is needed. In addition, the PLCP provides the byte clock (5.592MHz for the DS3 rate and 19.44MHz for the STS-3c rate) to the rest of the SMDS Interface. The byte clock is derived from locking the transmitting (local) clock with the incoming clock from the SS.

As the frame structures of the DS3 and STS-3c rates are different, two PLCPs were designed, one operating at 45Mbps and the other at 155Mbps.

3.2.1 DS3

As shown in Figure 13, each cell is placed in one row making the process of placing and extracting cells simple [8]. The payload is asynchronously mapped into the frame and a *nibble stuffing* technique is used to maintain the $125\mu\text{s}$ frame cycle.

The nibble stuffing protocol involves the size of the trailer. During the first frame of the stuffing cycle, 13 nibbles are transmitted. During the second frame, 14 nibbles are transmitted. During the third frame, either 13 or 14 nibbles are transmitted depending on whether the frame has been stuffed. The cycles are repeated every three frames.

3.2.2 STS-3c

As shown in Figure 14, each STS-3c row has 260 bytes of payload. This means that more than one cell can be placed per row. Cells can start at any point in the payload and can straddle across row and frame boundaries. This *floating payload* structure makes the process of extracting cell from the payload difficult.

There are two ways to indicate the beginning of a cell in the payload

Pointer: A byte in the *Payload Overhead* is used to indicate the beginning of

the cell. This *Offset Indicator* represents the number of bytes from the beginning of the row where the cell starts. After the first cell is received based on the Offset Indicator, subsequent contiguous cells can be detected as they are of the fixed 53 bytes size. If the PLCP lost track of cells, it can be re-synchronized upon receiving the Offset Indicator in the next frame.

Using Framing State Machine: Another way to recognize the beginning of a cell is to search for the correct 5-byte Cell Header. To ensure that the PLCP is detecting cells instead of mistaking a string of data for the beginning of a cell, 6 consecutive cells must be detected before the PLCP considers to be in *In-Slot-Delineation*.

Currently, the Framing State Machine is required to recognize cells while the use of the Offset Indicator is optional. The STS-3c of The BayBridge PLCP only implements the Framing State Machine.

3.3 Management

In addition to carrying cells in the payload, each row in the 125 μ s frame also carries one byte of the *Payload Overhead*. The Framer is responsible of alerting the PLCP when the first byte of the Payload Overhead is detected. The Payload Overhead is used to indicate a number of errors and alarm conditions and three management functions [9, 10].

3.3.1 Error Indicators

The Error Indicators are set when an error condition occurs which prevent operation. There are three error conditions:

Loss of Signal: An all-zeros pattern is detected on the incoming signal for a long period of time (100 μ s or longer).

Loss of Frame: A beginning of a frame is not detected after 3ms.

Loss of Pointer: A valid pointer is not detected after 8 consecutive frames. This error condition is only applicable to STS-3c rate.

3.3.2 Alarm Indicators

The Alarm Indicators are set when an condition occurs which could prevent operation. There are three alarm conditions:

Line Alarm Indication Signal: The LAIS is sent downstream to indicate that an error has been detected in the Line upstream.

Path Alarm Indication Signal: The PAIS is sent downstream to indicate that an error has been detected in the Path downstream.

Far End Receive Failure/Path Yellow: The FRF signal is used to indicate upstream that an error has been detected downstream.

3.3.3 Bus Identification

The Bus Identification Management Byte (Byte M1 in the POH) is used to indicate which Bus a node transmits on. For example, MACA receives and transmits on Bus A while MACB on Bus B. In the Single-CPE configuration, the SS needs to transmit the Bus Identification byte to ensure that the CPE is receiving on Bus A. However, the CPE does not need to send back its own Bus Identification to indicate that it transmits on Bus B. On the other hand, in the Multiple-CPE configuration, all CPEs need to transmit their Bus Identification bytes.

3.3.4 Message Identifier Page Allocation

As indicated in the Segmentation & Reassembly section, the Message Identifier (MID) is used to transmit and receive cells in conjunction with the Source and Destination Addresses. Thus, each CPE must acquire at least one MID in order to transmit and receive in an SMDS network. The MID Page Allocation protocol has been specified as a distributed protocol where the SS continuously announces the MID value (one per frame) via the MID Page Allocation byte (byte M2 in the Payload Overhead). A brief overview of the protocol is represented in Figure 15.

In the Single-CPE configuration, there is only one CPE that needs to reserve the MID, it does not need to notify the SS which MID value it will use. This means that the CPE does not need to send out the MID Page Allocation byte. In the Multiple-CPE configuration, each CPE needs to transmit its MID Page Allocation byte to ensure that the MID value it is using is unique and has not been reserved by another node.

3.3.5 Link Status

A state machine in the Management logic is used to determine the status of the received link. Based on the current state and the incoming signals, the

incoming link is considered as *connected*, *rx_link_up* or *rx_link_down*. These *link status* indicators determine whether cells are being received. The node can then notify its upstream neighbor of the status of the link by transmitting its own Link Status on the opposite bus. The Link Status Signal is transmitted in the last 3 bits of byte G1 at DS3 rate and the last 6 bits of byte H4 at STS-3c rate.

In the Single-CPE configuration, the Link Status signals determine when cells are received from SS and when cells can be transmitted to SS. In the Multiple-CPE configuration, if the Link Status signal shows that the received link is down, it will cause the CPE to reconfigure the network to block out the faulty node.

4 Operation

In this section, the operation of a Single-CPE SMDS network at the STS-3c rate is described from the initialization process through data transmission and error recovery.

4.1 Initialization

The initialization process at the CPE can be divided into the following steps:

1. *Looking for Start of Frame:* The beginning of a frame is indicated by 3 A1 bytes (11110110).
2. *Looking for Start of Payload Overhead:* Searching for the first byte of the Payload Overhead is only necessary at STS-3 rate, as the Payload Overhead can start anywhere within the column. By recognizing J1, the first byte of POH, other POH bytes can be found in the sequence following J1 as shown in Fig 14.
3. *Checking Bus Identification:* The correct Bus Identification value (byte M1) indicates that the CPE is connected correctly to the SS and/or other CPEs. Bus Identification is important only in the Multiple-CPE configuration.
4. *Checking MID Page Allocation:* CPE monitors byte M2 in each frame to determine the MID value it will use to transmit cells. The first 511 MID values (1 to 511) are reserved for the SS while the CPEs can use the rest of the MID field (512 to 1023). The MID Page Allocation is important in the Multiple-CPE environment only where CPEs need to obtain their

unique MIDs. In this case, the CPE at the furthest distance for the Head of Bus A (HOBA) will be able to reserve the first MID while the CPE closest to HOBA will be the last to reserve its MID value (see Figure 15).

5. *Checking Link Status Signal:* CPE can begin to transmit cells to the SS only after the incoming Link Status indicates that the link is *connected* or *rx_link_up*.
6. *Looking for cells:* While the management logic processes the Payload Overhead, the PLCP logic search through the payload for cells. Only after receiving 6 consecutive cells with correct headers does the PLCP consider to be *in sync* and begins to alert the MAC of incoming cells.

After all the steps in the Initialization process have been completed, the CPE is ready to transmit cells to the SS. Note that the earliest time the CPE can begin data transmission is after receiving nearly a frame or $125\mu\text{s}$. In fact, the *SMDS Start-up Procedures* require that the CPE waits for *7 seconds* after receiving frames from the SS before transmitting busy cells.

4.1.1 Data Transmission

During the data transmission process, the CPE receives cells from the SS in the $125\mu\text{s}$ frame payload. The CPE uses the $125\mu\text{s}$ timing marks in the receiving direction to generate its own frames which carry cells back to the SS.

Cells are assumed to arrive consecutively once the *Cell Delineation* process has been completed during the Initialization process. If the cells are not received correctly, the PLCP is required to search for the correct Cell Header again (as during the Initialization process). Assuming that no serious error condition occurs, the cells will be received in the correct sequence again. The MAC transmits and receives cells as described in section 3.1.

Once the CPE has reserved its MID, it continues to check and keep the MID as MID values are broadcast on the network. This will ensure that the MID used by each CPE is unique.

4.1.2 Error Recovery

If an error or alarm signal is generated by the Framer, it means that the incoming link is no longer operating correctly. The CPE will alert the upstream node of the condition while trying to reconfigure the network to exclude the upstream

node. This is only applicable in the Multiple-CPE configuration; in the Single-CPE environment, there is little the CPE can do except to notify the SS.

Errors may also arise from the MID Page Allocation protocol. As stated in the previous section, the CPE continues to keep the MID value once it has reserved it. If the MID value has been reserved by this CPE arrives and is marked by another CPE (upstream), the CPE *lost* its MID. It needs to notify its management and to reserve another MID value. This scenario would occur in the Multiple-CPE environment but not in Single-CPE.

If the PLCP starts to receive incorrect Cell Headers, it means that it has become *out-of-sync* with the upstream node. In fact, if the PLCP receives 7 consecutive cells with incorrect headers, it will stop alerting the MAC of incoming cells. It will return to the *Cell Delineation* process to search for correct Cell Headers again. If the searching process goes on for 1ms and cells are still not detected, the PLCP will reset itself to start the initialization process again.

5 Conclusion

In this paper, a Single-CPE SMDS has been demonstrated at STS-3c rate for direct connection or as interface to a bridge/router. While SMDS offers a comprehensive service to the users at high speed and a path towards ATM, it is not without problems. In this section, some of the drawbacks of SMDS are presented: MID Page Allocation Protocol, Redundancy in SMDS Protocol, and Cell Delineation Process.

- *MID Page Allocation Protocol*: The protocol requires that each CPE has a 10-bit counter to keep track of the 1023 MID values to be broadcast on Bus A. In addition, each M2 byte transmitted includes the two least significant bits of the current MID to ensure that all CPEs in the network are processing the same MID value. This requires a state machine to be implemented as a part of the Management logic. It would be simpler, for example, to assign each CPE a MID value before it is connected to the network - the way addresses are assigned.
- *Redundancy in SMDS Protocol*: SMDS is designed to interoperate with as many protocols as possible, as a result there is a high overhead in SMDS. There are *36 bytes* in the Header and *4 to 8 bytes* in the Trailer (optional 32-bit CRC). In addition, the LLC Header and SNAP Headers require 8 bytes which brings the Header to 48 bytes [11]. Thus, a minimum of *2 cells* are required to transmit an IP packet or encapsulated frame.

Some examples of the overheads are:

- The optional 32-bit CRC adds to the length of the header. It also adds to the complexity of the Segmentation & Reassembly design. Given the low bit error rate and other physical layer error checking mechanisms (e.g. Cell CRC), the CRC is not necessary to maintain an operational network.
 - The Payload Length in each cell is redundant, as the Length of the Message is already known.
 - The Information field is padded to be word-aligned. This means the last two bits of the *BA Size* field are always 0s. Thus, the Padding Length indication bits are redundant.
 - The *HLPI* field was defined to indicate the Higher Layer Protocol Identifier, this field could be used to indicate that the message is a *FDDI encapsulated packet*. The *HLPI* field is always set to 1 for LLC, thus the information about the message has to be stored in the LLC Header and SNAP Header which adds another 8 bytes to the Message Header.
- *Cell Delineation Process*: If the *H4 Pointer* is used to indicate the beginning of a cell, its implementation would only require a comparator and a 6-bit counter. On the other hand, to implement the Cell Delineation protocol requires a state machine, a 3-bit counter, and a 1ms timer. In addition, there is a possibility of continuing to misinterpret a 5-byte string as the beginning of a cell if it repeats every 53 bytes.

6 Acknowledgements

This work was supported in part by Pacific Bell and California State MICRO Program. We would like to thank Pacific Bell for their continued support and encouragement in this work. We are also grateful to Bellcore, Sun Microsystems, Advanced Micro Devices, TranSwitch, Integrated Device Technology, Viewlogic, Hewlett-Packard and MUSIC Semiconductors for their material and technical assistance. Much credit is also due to the following students at UC Berkeley who have contributed many hard and long hours to this project: Ting Kao, Fred Burghardt, Malik Audeh, George Kesidis, Karl Petty, Steve McCanne and Nina Taft. And finally, we thank our faculty advisors Professors Pravin Varaiya and Jean Walrand for their support.

References

- [1] Bellcore Technical Reference TR-TSV-000772, "Generic Systems Requirements in Support of Switched Multi-megabit Data Service," May 1991
- [2] Edell R., Le M., McKeown N., "The BayBridge, A High Speed Bridge/Router," in *Proceeding of IFIP Workshop on Protocols for High Speed Networks*, Stockholm, Sweden, May 1992
- [3] McKeown N., Edell R., Le M., "Architecture and Performance of The Bay-Bridge: A High Speed Bridge/Router between FDDI and SMDS," Submitted to this issue of JSAC, June 1992
- [4] Bellcore Technical Reference TR-NWT-000253, "Synchronous Optical Network (SONET) Transport Systems: common Generic Criteria," Issue 6, September 1990
- [5] Burghardt F., "DQDB MAC Chip Datasheet," *BayBridge Project Report 10, Rev 1.0*, Jan 1992
- [6] Robe T.J., Walsh K.A., "A SONET STS-3C User Network Interface Integrated Circuit," *IEEE Journal On Selected Areas In Communications; V9 N5:732-740*, June 1991
- [7] McKeown N., "Segmentation and Reassembly for Cell based Mis-ordering Networks," *Hewlett-Packard Technical Report HPL-91-176*, June 1991
- [8] Bellcore Technical Reference TR-TSV-000773, "Local Access System Generic Requirements, Objectives, and Interfaces in Support of Switched Multi-megabit Data Service," Issue 1, June 1991
- [9] European Telecommunication Standard, "Metropolitan Area Network Physical Layer Convergence Procedure for 155.520Mbps CCITT G.707-9 SHD based systems," May 1991
- [10] Institute of Electrical & Electronic Engineers, Inc., Standard 802.6, 1990
- [11] Piscitello D., Lawrence J., "The Transmission of IP Datagrams over the SMDS Service," *Draft RFC*, March 1991

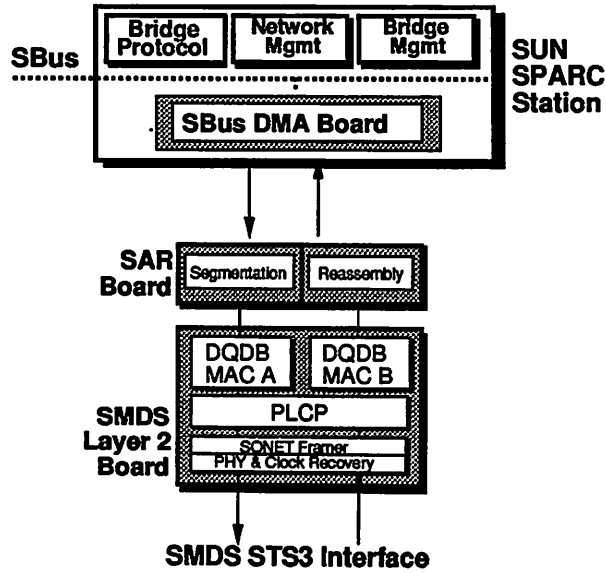


Figure 4: SMDS Interface as a direct connection

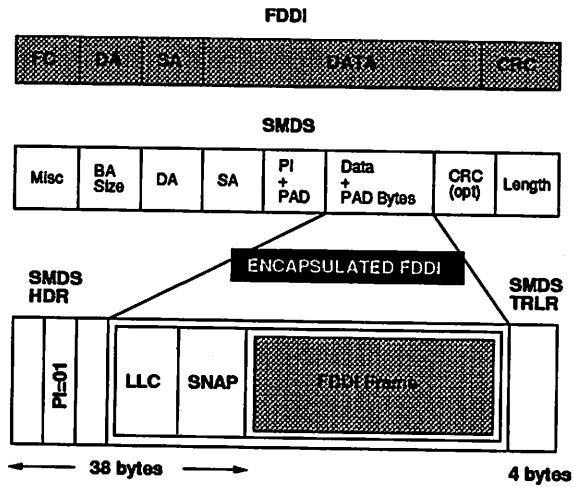


Figure 5: Encapsulating FDDI PDU to L3-PDU

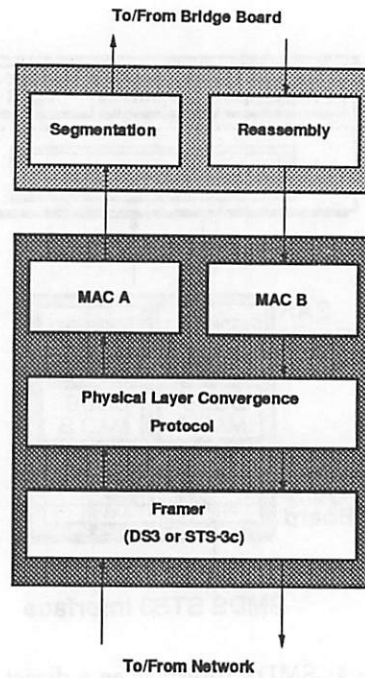


Figure 6: SMDS Interface of The BayBridge

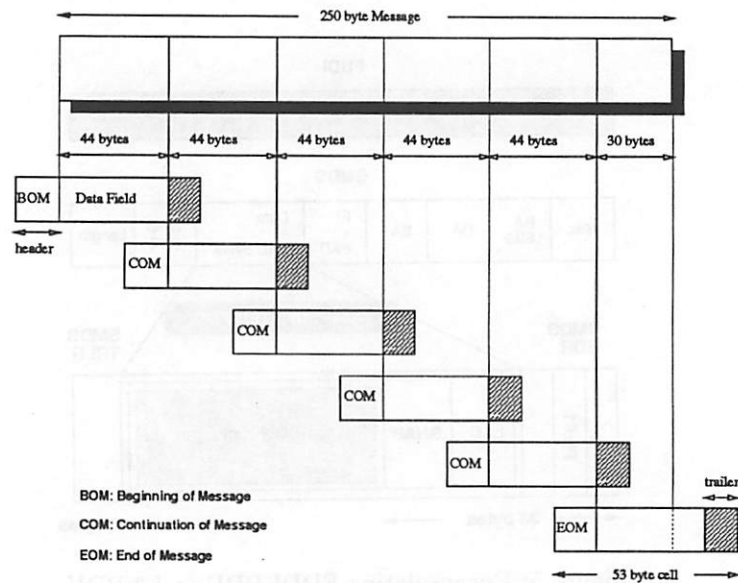


Figure 7: Segmentating a 250-byte message

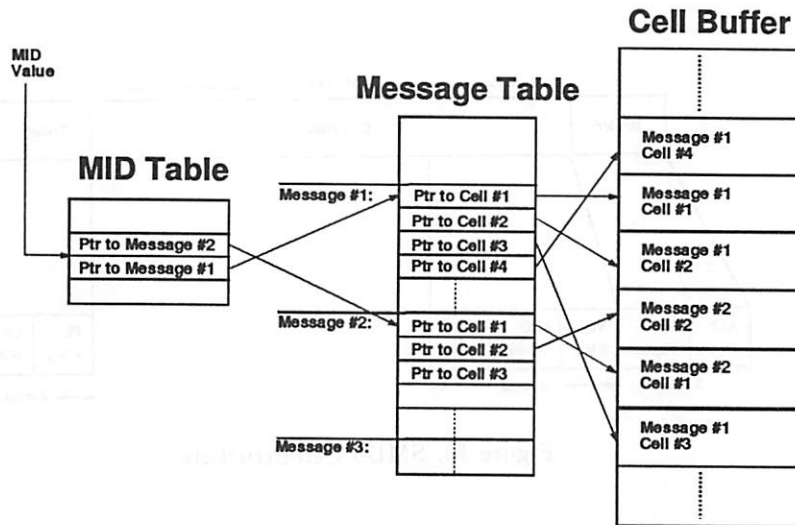


Figure 8: Reassembly for interleaved messages with misordered cells. The MID Table, Message Table and Cell Buffer are implemented as separate RAMs.

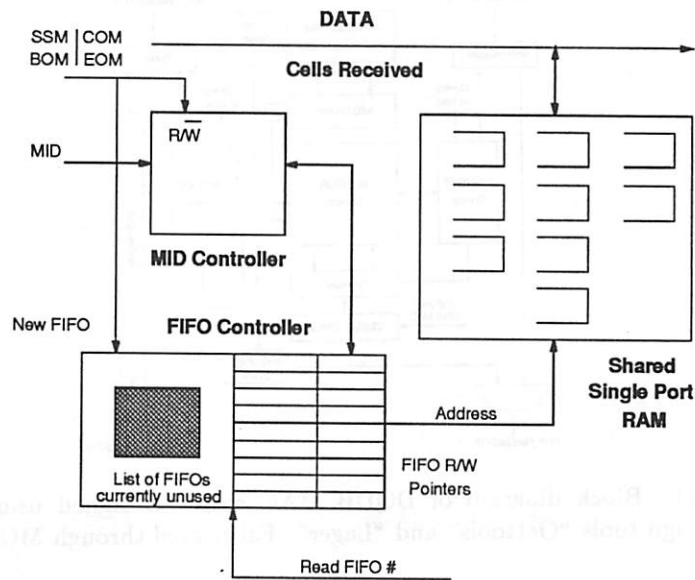


Figure 9: Proposed Reassembly design

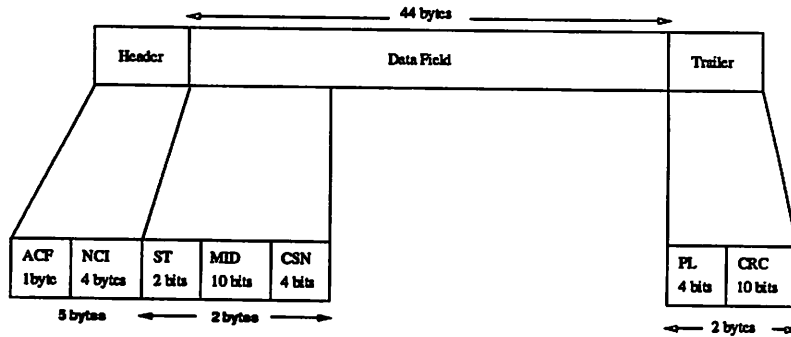


Figure 10: SMDS Cell Structure




Figure 11: Block diagram of DQDB MAC chip. Designed using Berkeley VLSI design tools "Octtools" and "Lager". Fabricated through MOSIS in 2µm CMOS.

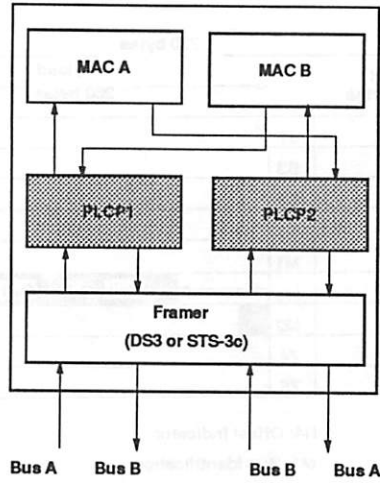


Figure 12: PLCP Connections




Figure 13: DS3 Frame


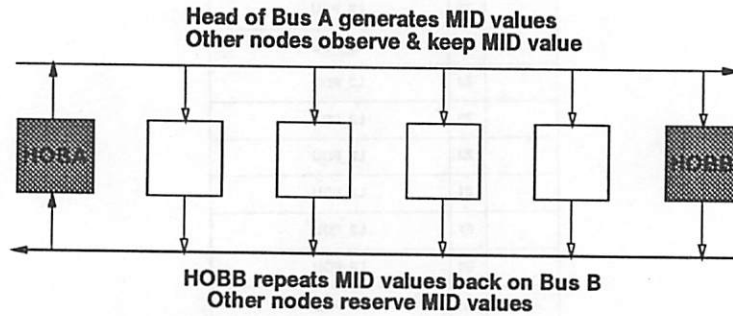


Figure 14: STS3 Frame



MID is a 10-bit field
There are 1023 MID values
1 to 511: reserved for SS
512 to 1023: to be used by CPEs

Figure 15: MID Page Allocation Protocol