PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

All-Digital Techniques for the Optical Broadband ISDN

L. R. Linnell, H. E. Bussey, F. D. Porter and S. E. Minzer

Phil. Trans. R. Soc. Lond. A 1989 **329**, 47-56
doi: 10.1098/rsta.1989.0056

Phil. Trans. R. Soc. Lond. A **329**, 47–56 (1989)

Printed in Great Britain

All-digital techniques for the optical broadband ISDN

By L. R. Linnell, H. E. Bussey, F. D. Porter and S. E. Minzer

Bell Communications Research Inc., 435 *South Street, Morristown, New Jersey* 07960, *U.S.A.*

Advances in technology have brought down the cost of deployment for fibre-based communications systems. The extension of fibre-based systems directly into businesses and residences will have a major impact on the development of the integrated service digital network (ISDN), ushering in what is commonly referred to as broadband ISDN (BISDN). BISDN can support a wide variety of services over high-bandwidth digital interfaces ranging from very low bit-rate telementry through full-motion high-definition video.

Bellcore has been investigating two types of switching techniques for use in the BISDN: space division circuit switching and high-speed packet switching. Interest in the design of high-speed packet switches is now receiving greater emphasis because of the selection of asynchronous transfer mode (ATM) as the predominant switching and multiplexing technique for BISDN. Both types of switching fabrics can be dimensioned to operate efficiently over the proposed SONET transmission system.

This paper focuses on the ATM hardware for a second generation prototype of a local access BISDN that is being constructed to more fully understand the implications and trade-offs involved in the design of BISDN systems.

1. Broadband architecture

Figure 1 shows a target architecture for a broadband integrated service digital network (BISDN) local access network. In the distribution and drop plant between the remote electronics (RE) and the customer, single-mode fibre is dedicated to the customer, forming a star topology. Two fibres can be provided to each customer, one for upstream and one for downstream transmission, or both directions of transmission can share the same fibre by using wavelength-division multiplexing (WDM) techniques. The cost difference between these two alternatives is slight, with short loops favouring two fibres and longer loops favouring one fibre (Lu & Wolff 1988).

Single-mode fibre is also used in the feeder plant between the RE and the central office (CO), but operates at a higher rate near 2.4 Gbit s^{-1}. Although the RE-to-CO distances are longer than the customer-to-RE distances, the feeder costs are mitigated by the concentration provided by the RE and by the higher transmission rate in the feeder and represent less than 5% of the total cost (Lu & Wolff 1988).

Customers near the CO will be served directly from the CO. For customers further from the CO, the cost of dedicated optical fibre to the CO becomes increasingly significant and these customers can be served more economically through an RE (Lu & Wolff 1988). Thus, electronics costs in the RE trade-off advantageously against the cost of optical fibre in the feeder.

Several topologies are possible for interconnecting the RE to the CO. In a star topology, dedicated fibre connects each RE to the CO. In a ring topology, a fibre ring starts at the CO and passes through each RE using add/drop multiplexers to drop and insert channels. A combination star/ring is also possible for the feeder topology in which switched channels are interconnected in a star topology and broadcast channels in a bus topology. The latter

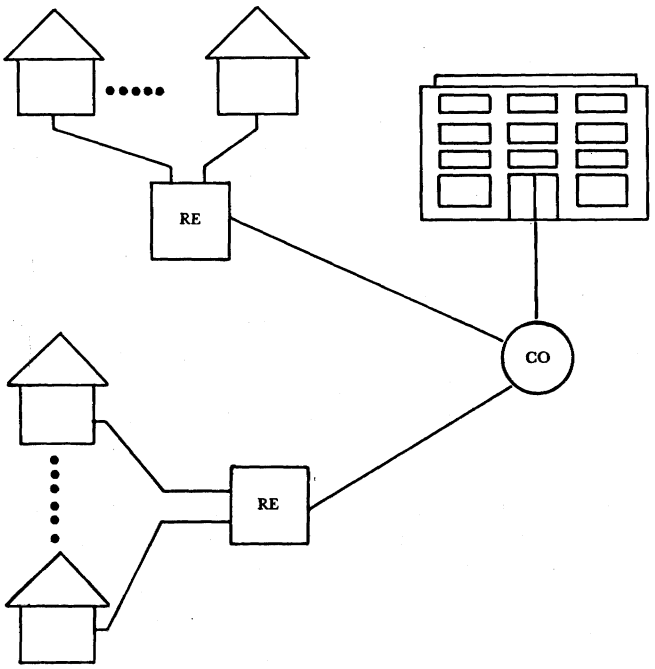

FIGURE 1. Long-term BISDN architecture.

alternative may offer some savings, but the effect on the overall cost is slight as the feeder cable represents only a small portion of the total cost.

Bellcore feels that the transmission system used over all of these fibre-based interfaces should have a standard optical network (SONET) frame structure†. Use of the same system for the feeder, distribution and trunking systems should significantly reduce implementation and maintenance costs. The currently proposed BISDN user-network interfaces will have rates at about 150 and 600 Mbit s^{-1} and can be based on SONET STS-3 and STS-12 signals at 155.52 and 622.08 Mbit s^{-1} respectively. BISDN interfaces need not be symmetrical, the customer selecting a rate in each direction to meet its service needs.

2. ASYNCHRONOUS TRANSFER MODE

Asynchronous transfer mode (ATM) is a multiplexing and switching technique that will enhance the flexibility of BISDN. Earlier systems, including basic and primary rate ISDN, carried services in time-division channels. In ATM, specific periodic time slots are not assigned to a channel (Beckner *et al.* 1987). Information is carried in fixed-size data units called cells, each one consisting of a header and an information field (see figure 2).

The header contains a label that uniquely identifies a channel and is used for multiplexing and routing. Signalling is carried on a designated labelled channel.

An earlier BISDN-like prototype constructed at Bellcore (Linnell 1986) was based upon a 600 Mbit s^{-1} user network interface (UNI) that was divided into four 150 Mbit s^{-1} channels. At least one of the four, 150 Mbit s^{-1} channels carried ATM cells. The remaining three channels, if

† A recommendation by the Comité Consultatif International Télégraphique et Téléphonique (CCITT) for a SONET transmission hierarchy is being standardized in the study period ending in 1988. The 1988 BISDN recommendation does not select a transmission system.

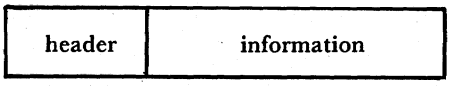

FIGURE 2. ATM cell.

active, could be used either to carry cells or as traditional time-division channels (Hayward *et al.* 1987).

More recently, a 150 Mbit s^{-1} interface has been proposed as another option complementing 600 Mbit s^{-1} (Bellcore 1987*a*). In contrast to the earlier prototype, all services over both the 150 and 600 Mbit s^{-1} interfaces are carried in ATM cells. The 150 Mbit s^{-1} STS-3 carries user information in a single 'payload module', the 'packet channel', containing a sequential stream of ATM cells. While some believe that a single payload module of close to 600 Mbit s^{-1} might be practical for an STS-12 interface, others feel that octet interleaving four STS-3 based payload modules is preferable because ATM equipment might be less expensive at the lower rate.

Because all user information is carried in cells, a single, integrated switching fabric for all services becomes fathomable. An all-purpose integrated fabric is the ideal in terms of service flexibility and might simplify operations and maintenance as well. However, because services produce different traffic characteristics and have different performance requirements, a next generation switch (NGS) (Bellcore 1987*b*) may incorporate different switch fabrics. For example, video signals may be carried over a circuit switch fabric, perhaps leaving the switched information in cell format, while lower-speed channels may be multiplexed by ATM mergers in the RE and switched by a Batcher–banyan self-routing cell switch in the CO. Although the physical implementation may incorporate multiple switching fabrics, the fabric appears to be integrated from the customer's perspective for both information transfer and signalling.

Although fixed-size cells are now preferred to variable length, agreement has not yet been reached on cell and header sizes. Several CCITT delegations favour three-octet headers plus 32 octets of information. Bellcore and the United States delegation feel that the header and information field should be larger. The three-octet header would contain a two-octet virtual channel identifier (VCI) and an error field. The VCI would not be subdivided into subfields containing values for specific ATM-related functions. Any required ATM-layer functional attributes would be obtained by table lookup, by using the VCI as key, by each processing node.

Subfields containing values for functional attributes might simplify implementation of some functions. For example, a line identifier might simplify concentration and distribution in the remote electronics, a terminal identifier might simplify media access and cell recognition. Components performing these functions would not need to access a table to determine the disposition of a cell.

It is desirable to have the same header size on all interfaces, including customer access, feeder and internal network links. The number of simultaneous connections carried over a network node interface (NNI) might be greater than the UNI. Capacity should also be reserved for unforeseen needs.

Considering these factors, a Bellcore contribution to Committee T1 (the standards body in the United States that is responsible for BISDN) stated that 'given the large label address space that might be needed and the uncertainty of other ATM-related functions, a 48 bit or 6 byte header is not excessive' (Bellcore 1987*c*). T1 concluded that the header should be at least four octets. The United States has ruled out a two-octet header, agreeing that it must be at least four octets.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Committee T1 also agreed that the information field size should be greater than 32 octets (Bellcore 1987*d*). Some CCITT delegations prefer the smaller cell size to minimize the packetization delay incurred in collecting a full cell of digitized voice samples. By using 32-octet information fields, which incur 2 ms of delay for 64 kbit s^{-1} PCM, they hope to avoid echo when interworking with existing voice fabrics. An LSSGR (Bellcore 1987*e*) requirement limits the echo path delay for a locally switched voice call to 8 ms. Assuming ATM access is interworked to a 64 kbit s^{-1} circuit-switched fabric, this limit would be exceeded even if a cell only contains 16 octets of information (AT&T 1987). Thus, reducing the size of the information field to 32 octets might not eliminate the need for echo cancellers at interworking points. There are other applications for echo cancellers too; e.g. portable radio phones in automobiles and speaker phones. The cost of echo cancellers most likely will continue to decline.

A small cell size puts extreme pressure on maximum acceptable size of the header field, because, for small cells, larger headers significantly reduce the percentage of capacity utilization. It also severely constrains the time available for ATM protocol processing. To guarantee real-time performance, ATM requires a cell to be processed in the same amount of time it takes to be transmitted. More thorough analysis of trade-offs among delay, capacity utilization and protocol processing factors is required. Traffic patterns for services other than just voice must be considered in determining parameters for a broadband network.

3. Interface parameters

Bellcore has proposed that a 150 Mbit s^{-1} UNI to BISDN be based on a SONET STS-3. STS-3 is the lowest-rate signal that can carry simultaneously an H_4-rate service of about 135 Mbit s^{-1} (sufficient for high quality video), signalling and some narrowband services, which is considered the minimal service requirement.

The parameters of interest are

- H_4, capacity in Mbit s^{-1} for a H_4-rate service (in the range of 132–138.264 Mbit s^{-1});
- H_{22}, capacity in Mbit s^{-1} for a H_{22}-rate service (in the approximate range of 43–45 Mbit s^{-1});
- x, capacity in Mbit s^{-1} for some narrowband services;
- D, capacity in Mbit s^{-1} for user-network signalling;
- C_{info}, total information payload provided to the user by ATM in Mbit s^{-1};
- I, information payload capacity of a single 125 μs frame in octets;
- S_h, length of the header portion of a cell in octets (in range of 3–8 octets in 1988 CCITT draft recommendation);
- S_i, length of the information portion of a cell in octets (in range of 32–120 octets in 1988 CCITT draft recommendation).

The values under consideration in the United States for these fundamental parameters for the 150 Mbit s^{-1} UNI are constrained by the following set of equations:

$$H_4 + D + x = C_{\text{info}}, \tag{1}$$

$$C_{\text{info}} = 10^{-6} \times 8000 \times (8S_i)\, I/(S_h + S_i), \tag{2}$$

$$I = 2340, \tag{3}$$

$$H \geqslant 3H_{22}. \tag{4}$$

The left side of equation (1) expresses the simultaneous service mix required of a 150 Mbit s^{-1} interface and its sum gives the total ATM information payload capacity available to the user. Narrowband services, x, take up the slack left over by an H_4-rate service and signalling. Precise capacity requirements for signalling have not been established. No precise value has been selected for x.

Equation (2) computes C_{info}, the total usable capacity provided by the ATM payload module. An integral number of cells appear in each 125 μs frame. The last term on the right-hand side of equation (2) is the 'greatest integer' function.

Equation (3) indicates the information payload capacity within an STS-3; 2340 octets are available in the 9×270 SONET frame structure (9×260).

The upper bound on the value of H_4 channel rate is set by the capacity of the European CEPT 4 signal. A desirable lower bound for the United States is expressed in equation (4). It is sensible to define channel rates to allow substitution of three H_{22} services for one H_4 service. If it is necessary to lower the value for H_4 to accommodate a larger number of simultaneous services, it may also be desirable to adjust the H_{22} accordingly.

4. ATM HEADER TYPES SUPPORTED

To investigate the implications of some design decisions associated with ATM header functionality, a second generation of ATM prototyping hardware is being designed. It incorporates a custom, single chip (Chao *et al.* 1988) interface to a high-speed, framed, serial bit stream.

The hardware described in this paper supports two distinct header formats: partitioned and unpartitioned. The partitioned formal results in a long header (6 octets), while the unpartitioned format results in a short header (4 octets). The cell body, which makes up the remainder of the cell, may be 44, 66, 88, 132 or 264 octets.

4.1. *Partitioned (long) header*

The partitioned (or long) format (figure 3*a*) explicitly includes fields for a variety of functions that might be performed by the ATM layer of the network, such as processing according to traffic type or priority, or identifying the physical address of the source by use of a line equipment number (LEN). The fields in the header are defined as follows.

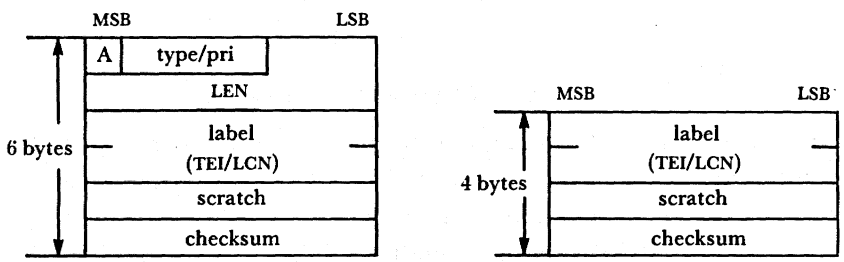

FIGURE 3. (*a*) Long and (*b*) short header format.

A is the activity bit, is 1 if the position is occupied by an active cell and 0 otherwise.

Type/Pri is a four-bit field that tells to which class of traffic the cell belongs. This could act as a priority field when multiple cells are competing for the same resource, or it could be used

to separate various types of traffic (e.g. voice, video, data). This field might be treated as if it is part of the VCN.

The line equipment number (LEN) is an 11-bit field that identifies a customer line within an RE. A cell multiplexer in an RE could fill in this field so that upstream cell-processing equipment could differentiate between cells from different customers.

VCN is a 16-bit field that identifies a logical stream of cells on a customer's line. This could be treated as one 16-bit name space, or it could be divided into two subfields: the terminal endpoint identifier (TEI) and the logical channel number (LCN).

Scratch is a byte included as a scratch pad area that can be used in a variety of ways, depending on its context. For example, customer premises equipment (CPE) may use it for media access control to a multiaccess ATM link, or a self-routing switching fabric could use it for insertion of switch port numbers.

Error. While error correction schemes are also being considered, the prototype uses this field to detect errors in the cell header. Cells with erroneous checksums are discarded, decreasing the probability of misdelivery of cells. Thus the checksum should use a robust error detection method, such as a cyclic redundancy code (CRC). To simplify implementation, however, we use a simple exclusive–or checksum.

4.2. *Unpartitioned (short) header*

The short format is shown in figure 3*b*. Information such as whether a cell position (slot) is occupied, the cell priority, etc. is implicitly given by the label, and obtained by table lookup. The short header is a subset of the long header, consisting of the label, the scratch field and the checksum (over just the label and scratch fields). As no activity bit is included, a special logical channel number (e.g. *label* = 0) identifies an empty slot.

4.3. *Implementation*

The second-generation research prototype supports both the long and short header formats, enabling research on the merits of each approach. In the long format, only the activity bit, 3 bits of priority, four bits of LEN, and 8 bits of label are examined. This limits the number of lines and logical channels that can be supported, but for a prototype this is acceptable. In use, either the appropriate field is accessed directly (as in a very simple cell multiplexer, which might examine only the occupancy field), or all 16 bits are used by the hardware to perform table lookups, exactly as the 16 bit label of the short header is used.

5. Second generation ATM hardware

A broadband ATM system can be configured using a small number of primitive functions.

Selection (1 input, 1 output). Packets are selected based on some criteria, e.g. whether or not they belong to some set of logical channels. The remaining cells are deleted.

Merging (*N* inputs, 1 output). Packets from multiple links are merged onto a single link. This function requires buffering. Note that the average traffic on the input links must not exceed the capacity of the output link for an extended length of time, or buffers will overflow and cells will be lost.

Switching (*N* inputs, *M* outputs). For every logical channel at a switch input, the cells associated with the logical channel are switched to one of the output links. Like merging,

switching requires buffering. This definition of switching does *not* include broadcasting or header translation.

Translation (1 input, 1 output). The header of each cell, or portions of it, is translated into a new header.

Terminal Interface (1 input, 1 output). At the input, cells on the link are made available to the terminal equipment. At the output, the terminal equipment puts cells onto the link. A terminal may be part of the network, e.g. an interface to call processing, or part of the CPE. This set of primitives leads to the following observations:

selection can be viewed as a subset of translation, because a cell header can be translated in a way such that the cell becomes 'inactive';

merging can be implemented with examination of only occupancy information in the header. In contrast, selection, switching and translation are functions whose actions are based on the contents of the cell headers;

only translators *modify* cell headers.

Based upon the above considerations, we have partitioned our implementation into the following modules:

cell translator-filter module (CXF) examines the input value of a cell's VCN for selection and possible translation;

cell multiplexer (CM) merges cells from several input ATM payload modules into a single output payload module;

cell interface (CI) is a terminal interface primitive for data;

cell video codec (CV) is a terminal interface primitive for video;

cell switch (CS) switches cells, based on a switch output port number derived from the VCN, from one of a set of input streams to one of a set of output streams.

5.1. *The cell translator and filter*

Figure 4 shows a block diagram of the CXF. The header of each incoming cell is examined for cell occupancy. If it is occupied, the header is latched by the controller, which then uses the VCN to index into the translation table. The translation table entry includes a new VCN value for the output link. A switch output port number is placed in the 'scratch' field of the header for self-routing through the CS, and a counter field is incremented to keep track of the number of cells transported over the virtual circuit. While the header is being processed, the rest of the cell is buffered in a FIFO. By the time the last byte of the cell has been loaded into the FIFO, the translation is complete and the controller has latched the new header. This new header is appended to the information field and the entire cell is inserted in the next slot for transmission. A cell must be completely processed and its buffer available for the next cell within the time required to transmit a cell so that CXF can handle cells back-to-back without loss.

The cell filter is a special case of translation, in which the input header value translates to an inoperative VCN value, causing the cell to be discarded. In this case, the cell body is flushed from the FIFO while the empty header and a dummy (zeroed) body are inserted into the next output slot.

A VME bus interface is used by the control processor to initialize the board, maintain the translations, and access the traffic counters. The cell translator–filters are packaged four to a board (all boards for this prototype are double-Eurocards).

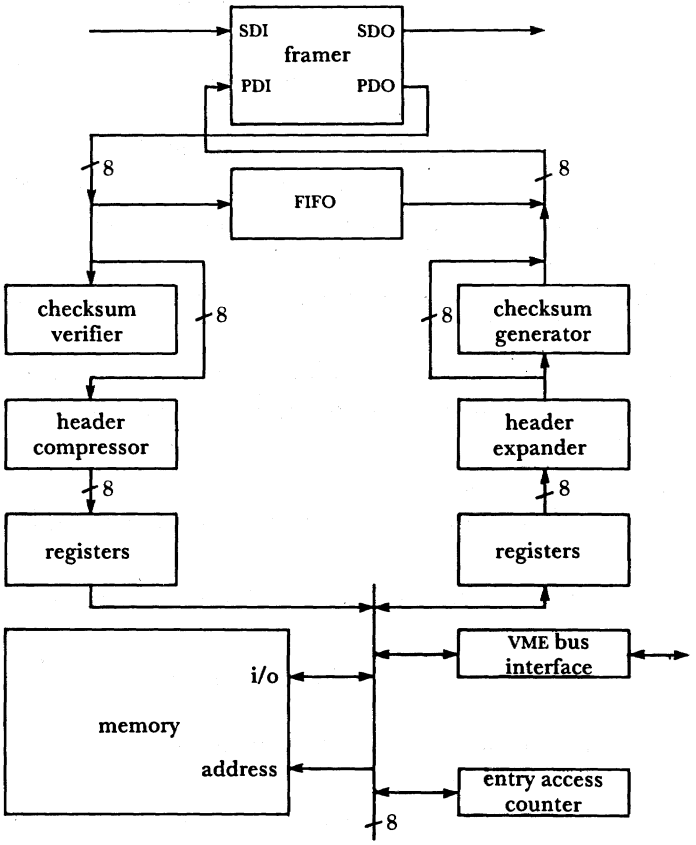

FIGURE 4. The cell translator–filter module. SDI is 150 Mbit s^{-1} serial data input; SDO is 150 Mbit s^{-1} serial data output; PDI is parallel data input; PDO is parallel data output.

5.2. *The cell multiplexer*

This module (figure 5) is a board that merges cells from four input links onto a single output link. On each input line, active cells are detected by examination of the headers and buffered in a separate FIFO for each input, so that cell ordering is preserved. The output server uses one of two service disciplines for transferring cells from input link FIFOs to the output: round-robin or spatial priority, in which lower-numbered inputs have priority over higher-numbered inputs. Larger multiplexers can be constructed from multiple 4:1 modules, by cascading them in a tree topology.

5.3. *The cell switch*

The CS module routes cells between up to 32 inputs and outputs. Incoming cells are routed to the output port indicated in the 'scratch' field in the header. To minimize the design effort, much of the design of the packet switch of the first-generation prototype will be used, including the ring reservation method of output port arbitration. The cell switch comprises multiple boards: the Batcher–banyan fabric board, the cell switch interface (CSI) boards (one per populated port), the timing and synchronization board, and the ring head-end board. The two latter boards may be combined.

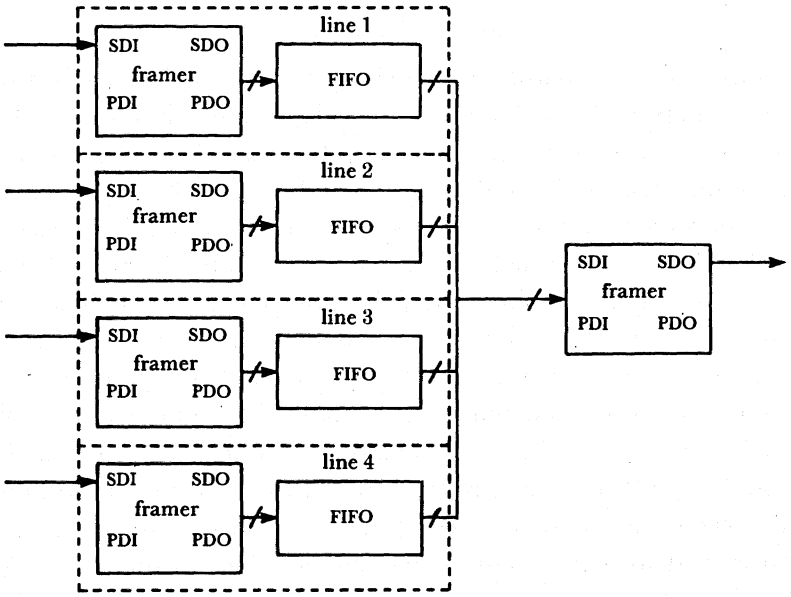

FIGURE 5. The cell multiplexer module.

5.4. *The cell interface*

The CI provides the terminal interface function. While many terminal interfaces are possible, a buffer board, similar to the one in the first-generation packet interface, seems to be the most reasonable choice. The CI acts as a memory-mapped peripheral device. The CI is connected to a VME bus to allow processor access. As figure 6*a* shows, incoming cells are detected and filtered according to their headers. Cells to be received by a CI are stored in a receive FIFO. The CI can generate VME bus interrupts following cell arrivals, or it can also be operated in a polled mode. Outgoing cells are written to an output FIFO and transmitted in the next available slot. To provide the capability to either delete or overwrite incoming cells, the input and output sides of the CI can be connected with several timing signals that indicate and control the appropriate action. The delay through the CI is only several bytes.

CIs can be connected in a variety of topologies (figure 6*b*), including a point-to-point, a daisy chain and a dual-bus configuration. These configurations provide for a wide variety of terminal equipment interconnection topologies.

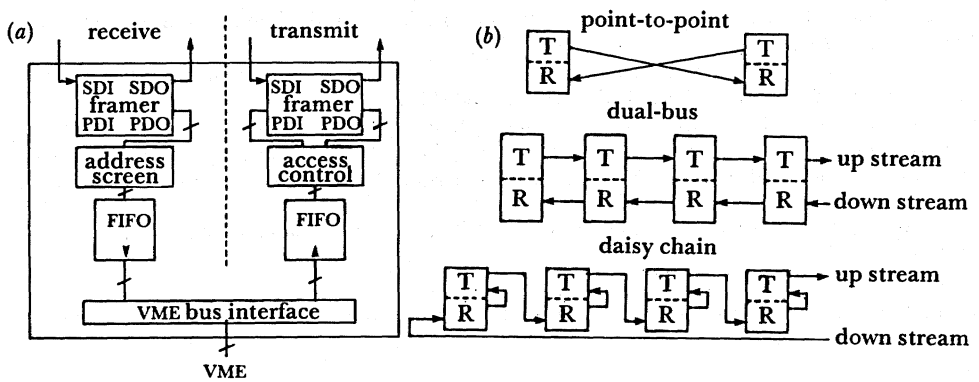

FIGURE 6. (*a*) Design and (*b*) configurations of the cell interface.

5.5. *A cell video codec*

This module consists of two boards: the coder board and the decoder board. These boards are being prototyped to provide our laboratory with high-bandwidth terminal equipment and to investigate the problems of recovering the transmitter's sampling clock at the receiver, when video is transmitted in ATM cells (Chao & Johnston 1988). As such, no particular attention was given to coding, compression, or signal quality issues. The coder board digitizes incoming audio and video channels. The audio channel is digitized using a 16-bit linear analogue–digital (A–D) converter clocked at a 64 kHz sampling rate. The video channel is digitized at 7.5 MHz† by using an 8-bit linear A–D flash converter. The PCM data from the audio and video channels are clocked into separate FIFOs. As soon as a FIFO contains a complete cell, the control circuitry that allows the cell to be inserted into an empty slot is enabled. A control lead completely inhibits cell generation, allowing the coder to be turned off when no video call is active.

The decoder board detects incoming cells addressed to it and then routes the cells to either the audio or video FIFOs, depending on the cell header information. The average filling level of the video FIFO is used to regulate an oscillator that recreates the coder's 7.5 MHz clock, which is used for the video D–A converter. The 64 kHz audio PCM clock is derived from the network clock.

6. Conclusions

We have described cell-handling primitives that can be used in the architecture of cell-switched broadband networks. These components support both long and short headers and various data lengths so that alternative header functionality may be investigated. We also have partitioned the primitives into circuit modules, and shown how the modules may be interconnected in a variety of ways for construction of broadband network prototypes.

References

AT&T 1987 Delay associated with an ATM access to a synchronous local exchange. T1D1.1/87–285.

Beckner, M. W., Lee, T. T. & Minzer, S. E. 1987 A protocol and prototype for broadband subscriber access to ISDNs. In *Proc. Int. Switching Symp. 1987, Phoenix, Arizona*, pp. 462–469.

Bellcore 1987*a* Preliminary special report on broadband ISDN access. **SR-TSY-000857** (1).

Bellcore 1987*b* Next generation switch symposium. **SR-TSY-000755** (1).

Bellcore 1987*c* Some considerations in selecting a cell header size. T1D1.1/87–492.

Bellcore 1987*d* Some considerations in selecting a cell information field size. T1D1.1/87–493.

Bellcore 1987*e* LATA switching systems generic requirements. **TR-TSY-000064**.

Chao, H. I. & Johnston, C. 1988 A packet video system using the dynamic time-division multiplexing technique. In *Proc. Globecom. 1988*.

Chao, H. J., Robe, T. J. & Smoot, L. S. 1988 A 140 Mbits s^{-1} CMOS LSI framer chip for a broadband LSDN local access system. *IEEE J. solid St. Circts* **JSSC-23** (1), 133–141.

Hayward, G. A., Linnell, L. R., Mahoney, D. D. & Smoot, L. S. 1987 A broadband ISDN local access system using emerging-technology components. In *Proc. Int. Switching Symp. 1987, Phoenix, Arizona*, pp. 597–601.

Linnell, L. 1986 A wide-band local access system using emerging-technology components. *IEEE J. selected Areas Commun.* **SAC-4**, 612–618.

Lu, K. W. & Wolff, R. S. 1988 Cost analysis for switched star broadband access. In *Proc. Globecom. 1988*, pp. 48.4.1–48.4.7

† This sampling rate allows two video coders to share the same access link, but results in low-definition video, as no data-compression is used.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES