Flexible aggregation of channel bandwidth in primary rate ISDN

J W Burren
Rutherford Appleton Laboratory
Chilton, Didcot, Oxon, England, OX11 0QX

Abstract

This paper describes an algorithm that allows a number of 64 Kbps ISDN circuits running between the same pair of subscribers to be aggregated into a single wideband circuit. The algorithm permits flexible aggregation in the sense that the bandwidth of the wideband circuit may be changed by the addition or removal of 64 Kbps circuits without disturbing the flow of data on the wideband channel. The paper also describes an implementation of this algorithm in an apparatus that interfaces a LAN to a primary rate ISDN interface. This interface is built using an array of transputers, and can achieve high utilization of the bandwidth available on the primary rate channel.

1 Introduction

In the coming years the technology used for wide-area digital communication will be based on the Integrated Services Digital Network (ISDN) recommendations of the CCITT [1]. ISDN networks currently being implemented in Europe and North America provide 64 Kbps calls between subscribers. Digital voice and any form of digital data may be carried on these calls. Domestic and small business customers are supplied through **basic rate** interfaces providing network access at 144 Kbps. The interface is configured as two 64 Kbps channels, the B-channels, and a 16 Kbps signalling channel, the D-channel. Larger subscribers in Europe are provided with **primary rate** access at 2.048 Mbps configured as 30 B-channels and a 64 Kbps D-channel, with the remaining bandwidth used for synchronization. Subscriber equipment operates a packet protocol on the signalling channel to originate and accept circuit-switched calls on the B-channels. In North America, primary rate access is at 1.544 Mbps, configured as 23 B-channels and a 64 Kbps D-channel.

The 64 Kbps B-channel is the standard for subscriber calls and is the fundamental unit on which the elements of the circuit switched network operate. While sufficient for voice and terminal traffic, 64 Kbps does not provide sufficient bandwidth to support the transfer of all types of data, particularly when digital images or large files of computer data are involved. This limited unit of bandwidth is also not appropriate for circuits carrying packet multiplexed traffic, as in the case of circuits joining widely separated Local Area Networks (LANs). The ISDN recommendations also describe H-channels which can be assembled from multiple 64 Kbps channels. However, many aspects of the implementation of these channels are still under study and it seems unlikely that these will be widely available for a number of years.

The problem of constructing equipment functioning outside the network that can aggregate B-channels to form channels of larger bandwidth is non-trivial, since calls, that is B-channels, between the same end-points do not in general follow the same path through the network. Thus the network transmission delays on different B-channels between the same end-points may be unequal and if data were fed down these B-channels as if they together constituted a single higher bandwidth channel, this data would be scrambled. However, the transit delay associated with a B-channel once established is constant and if the relative delays associated with the B-channels that are to be aggregated are measured, suitable apparatus at the receiving end can perform an unscrambling operation. To make this measurement marker signals must be transmitted from each end in a predefined sequence; these signals may be recognized at the receiving end and their relative positions in the data stream noted. The receiving end may then calculate displacements that restore the marker signals to their correct sequence and use these to rearrange all

the received data to produce a correctly-sequenced data stream.

This paper describes an apparatus called a **ramp** that implements this aggregating operation for primary rate interfaces. The ramp allows a set of n B-channels calling a single destination, where $1 \leq n \leq 30$, to be grouped together to form a single, wider channel of $n \times 64$ Kbps. The data from packets to be transmitted over this wide channel is spread byte-by-byte across its constituent channels, the ramps at each end ensuring that any scrambling of the data in the network is removed. A ramp can handle as many of these wide channels as can be constructed with the 30 B-channels available from the primary rate interface. The ramp also handles these channels in a flexible fashion which allows their bandwidth to be varied dynamically by adding or subtracting B-channels without disturbing the flow of data.

In Section 2 of this paper, the method used to aggregate channels is described. An implementation using Inmos Transputers is outlined in Section 3 and finally in Section 4 some performance measurements are given. (*Note* The algorithm described in Section 2 is the subject of U.S. patent number 4805167 and patents in the U.K. and Europe are pending.)

2 Variable rate channels for ISDN

The B-channel on ISDN, which is designed to carry digital voice, is a bidirectional channel carrying 64 000 bps in each direction and delivering octets of data to each end every 125 μsec, that is 8000 times per second. A primary rate bearer channel operating at 2.048 Mbps is time-division multiplexed into 32 sub-channels. Time-division frames contain 32 octets of data, one octet (byte) for each sub-channel. In this paper, the 32 positions will be called slots and will be numbered from 0 to 31. In each frame, slot 0 is used to identify the start of frame and slot 16 is reserved for the signalling channel, the D-channel. The remaining 30 slots may be used as B-channels to carry voice or data. The primary rate channel carries 8000 frames per second and a 64 Kbps signal is sent through the network as a series of octets inserted into the same slot of successive frames. Primary rate channels in the USA also employ 8000 frames per second, but the frames consist of 24 slots plus a single frame start bit.

Data in a particular slot passes from its source to its destination through the network via one or more switching centres. At each of these centres the data will be extracted from its incoming slot and placed in another slot in a frame travelling along the route to the destination. Octets placed in the same slot in successive frames travel through the network in exactly the same fashion and arrive at the destination as octets in a given slot in successive frames. However, the slot position within the carrier frame at the destination will not in general be the same as the slot position at the source. Thus, if octets carried on a B-channel between source A and destination B leave A in slot s_A in A's transmission frames and arrive in B's reception frames in slot s_B, then in general $s_A \neq s_B$. However, since B-channels are duplex, octets transmitted by B in slot s_B will arrive at A in slot s_A in A's reception frames. If a number of B-channels are to be established between A and B, each channel has to be established independently and will in general be subject to different delays in each exchange and may in fact pass through different exchanges in crossing the network. These differences in paths give rise to two types of effect, rearrangement within a frame and skew between slots. Thus if slots i and j, with $i < j$, are used at A for two B-channels to B, with corresponding slots l and m at B, it is possible that $m < l$, that is that the channels are rearranged within the frame. It is also possible for the network to cause skew in the alignment of octets within frames. Two octets that leave A in the same frame in slots i and j may arrive at B in slots m and l in different frames. To aggregate a number of B-channels into a single, wideband channel, ramps must be able to remove rearrangement and skew on reception.

In this paper, the wideband channels made up of a number of aggregated B-channels will be called U-channels. To allow a partner ramp to make the necessary measurements to determine any rearrangement or skew, ramps transmit a marker signal in each of the slots that are to be aggregated. These marker signals are sent in the same frame (or frames if the marker is longer than 1 octet) and each contains the slot number in which it is transmitted. The partner ramp recognizes these marker signals and for each slot notes the frame in which the marker is received. From the relative position of the frame in which the marker in a given slot is received the ramp can determine the skew associated with that slot. Rearrangement of slots within frames can be determined from the slot numbers carried in the markers. Since the path across the network of a B-channel remains fixed after it has been established, any skew or rearrangement associated with a given slot remains constant until the network call associated with that slot is terminated. To remove skew and rearrangement, the receiving ramp keeps at any instant a number of frames arranged in a matrix form, with the ith column holding octets from the ith slot and each row holding octets from the same frame, see figure 1. The process of reception is to scan across this matrix taking an octet from each column, that is slot, associated with an active channel. To take out any rearrangement, the columns are scanned in an order that, for each U-channel being received, corresponds to the order of transmission,

that is in ascending order of transmitting slot number as determined from the markers. To take out any skew, octets are taken from different rows, in such a way that if the markers for any U-channel had been read out in this fashion they would have appeared in the same scan. After each scan, the rows of the matrix are moved up one place, the top octet in each column disappearing and the next frame being added as the new bottom row, and the scanning process repeated.

ing received a marker with the slot synchronized flag set, a ramp is in a position to use the slot for data. This it does by raising the slot ready flag in its transmitted marker; immediately thereafter the U-channel is widened to include the new slot and the B-channel will carry data. A similar procedure will happen in the opposite direction. U-channels are built up in this fashion slot by slot and the width of a U-channel may be increased at any time by adding slots. The delete flag is

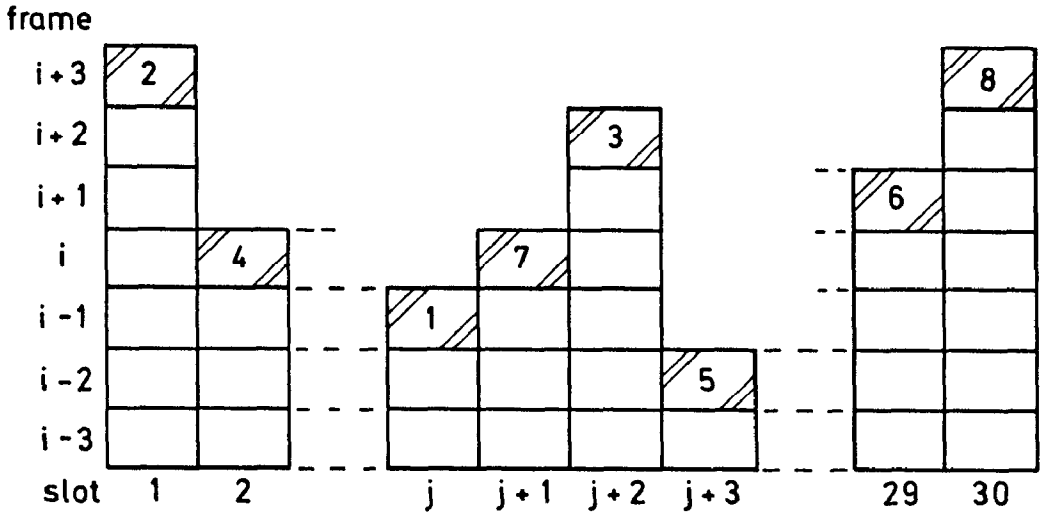

Shows octets taken in scan for i^{th} frame with rearrangement.

Fig. 1: Frame scanning procedure.

The scheme described above allows U-channels of fixed bandwidth to be created. The ramps, however, are able to deal with U-channels that have widths that vary with time. To achieve this, ramps transmit marker signals at regular intervals within the data, that is every nth frame, where n is fixed, a ramp transmits markers rather than data. Markers contain three single-bit flags defined as follows:

1. slot synchronized — receiving side of ramp has located markers in slot;
2. slot ready — slot is incorporated in U-channel;
3. slot delete — slot is deleted.

When a B-channel has been established across the network, the ramps at each end begin transmitting 'idle' octets, interspersed with markers at fixed intervals, in the slot associated with the B-channel (this slot is not, of course, in general the same at each end). At this stage the marker flags are unset. When a ramp finds a received marker in a slot, it raises the slot synchronized flag in its transmitted marker for that slot. Having received a marker with the slot synchronized flag set, a ramp is in a position to use the slot for data. This it does by raising the slot ready flag in its transmitted marker; immediately thereafter the U-channel is widened to include the new slot and the B-channel will carry data. A similar procedure will happen in the opposite direction. U-channels are built up in this fashion slot by slot and the width of a U-channel may be increased at any time by adding slots. The delete flag is used to remove slots. When a ramp is instructed to drop a slot from a U-channel, it raises the slot delete flag in the transmitted marker for that slot and immediately thereafter the slot ceases to be part of the U-channel and will no longer be used to carry data. On receiving a slot delete flag a ramp responds by setting the slot delete flag in its own transmitted marker for the particular slot. Thereafter the associated B-channel call across the network can be closed.

Using the scheme described above, channels (U-channels) having bandwidths of multiples of 64 Kbps can be established across an ISDN. The bandwidth of these channels can be varied dynamically, in units of 64 Kbps, without disturbing the flow of data.

3 Implementation

The aggregation scheme outlined above has been implemented in a ramp design that makes use of Inmos transputer technology [2]. The Inmos transputer is a complete computer on a chip running at 20 Mhz. The

chip contains a CPU, 4 Kbytes of memory and four duplex serial channels that run at speeds up to 20 Mbps. The four channels are independent and can run together duplexed, that is as 8 unidirectional channels, in DMA mode. It is a straightforward matter to add external memory to the chip. The main feature of interest in the technology is the ease with which multi-processor configurations can be built. At the hardware level, transputers can be joined by wiring together the serial channels; no external components are required for this linkage.

Software for multi-processor configurations can be written in the Occam programming language [3]. In Occam, software is structured as independent processes linked together by message-passing channels. These channels have exactly the same form whether they connect processes running in the same transputer or in a set of linked transputers.

The ramp implementation undertaken at the Rutherford Appleton Laboratory interfaces a 50 Mbps Local Area Network (LAN) of the slotted ring type, called a Cambridge Fast Ring (CFR) [4], to a primary rate ISDN channel. This implementation makes use of eight transputers configured as shown in figure 2. The links shown as thick lines carry data traffic between the LAN interface and the ISDN interface; the links shown with thin lines carry command and monitoring messages between the processors. Transputers 2 and 3 (transmit) and transputers 4 and 5 (receive) implement the aggregating scheme described in Section 2 above, transputer 6 deals with the LAN interface, and transputer 1 deals with queuing and address mapping for data being transmitted on ISDN. Transputers 0 and 7 deal with supervision of the ramp and with the signalling channel which is used for ISDN call setup.

Transputers are not fast enough to deal with the aggregating algorithm on an octet-by-octet basis at 2 Mbps. They are, however, fast enough to deal with frames of words, where words consist of 4 octets. Thus transputer 2 passes frames of 30 words to transputer 3. Transputer 3 converts this word frame into 4 ISDN octet frames, by taking the first octet from each of the 30 words to form the first ISDN frame and then the second octet for the next frame and so on. This simple process, together with input, output, and control processes runs in the internal (fast) memory of transputer 3. The task of transputer 4 is the reverse, that is taking in 4 octet frames to produce each word frame for transputer 5. The procedure is more complex since there are 4 possible starting frames for forming the words and these may vary from slot to slot. Transputer 5 is responsible for indicating to transputer 4 how the alignment of octets into words is to be performed for each slot. Transputers 2 and 5 implement the aggregating algorithm for fixed length packets of N words. Transputer 2 forms word frames by scanning across the 30 slots, determining for each slot to which U-channel it belongs and then taking the next word from the packet currently being transmitted on that U-channel. If no packet is being transmitted, an idle word is placed in the slot. At fixed intervals, frames of marker words are transmitted. In addition to the transmitting slot number and the three flag bits, these contain a fixed bit pattern to facilitate recognition by the receiver. Transputer 5 receives word frames and scans across these in the manner described in the previous section to produce packets of correctly ordered words. After a new slot has been setup by a call across the network, only idle words and markers are transmitted. Transputer 5 will receive one of the four possible forms of idle word, corresponding to the four possible alignments of octets into words. From inspection of the actual idle word received it can instruct transputer 4 on the appropriate alignment of octets into words for the slot.

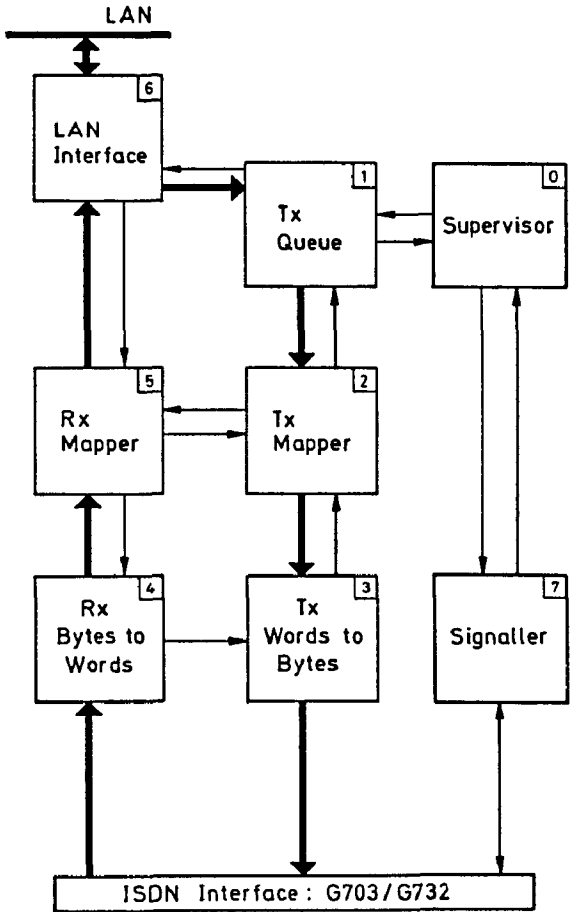

Fig·2 : Ramp configuration.

Transputer 1 organizes queues of packets for each U-channel. It must be able to determine the appropriate U-channel for each packet from addressing information carried in the packet. Any required address mapping may be carried out in this processor. Queuing of packets

for U-channels is organized on a priority basis, currently with two priorities.

Transputer 7 deals with the signalling channel, the D-channel. Quite inappropriately, HDLC bit-stuffing is required on this channel and this necessitates the separation of this channel so that a suitable chip can be used to implement this bit-stuffing. Transputer 7 controls this chip and handles the transmission and reception of signalling packets on the D-channel. Supervision of the operation of the ramp is performed by transputer 0. This takes commands from the network or from a terminal and instructs processes in the various transputers of the ramp to carry out these commands. Typical commands would be to create a U-channel to a particular ISDN address, to change the width of this channel by adding or removing slots and to associate packet addresses with a particular U-channel. The supervisor also collects monitoring information on the state of the ramp, relays this to other computers on the LAN and displays summary performance information on a monitor. To allow processes in the ramp program, which resides in eight separate processors, to exchange messages, each transputer contains a message *exchange* process. These exchange processes act together to provide a packet network within the transputer configuration to carry addressed messages between processes. The configuration of processes running in transputer 2 is shown in figure 3.

The ramp is built on standard, double-height Eurocards. Five of these cards are used to carry the eight transputers, associated memories, and LAN and G.732 interfaces. A separate small box houses the line interface card that implements the G.703 interface.

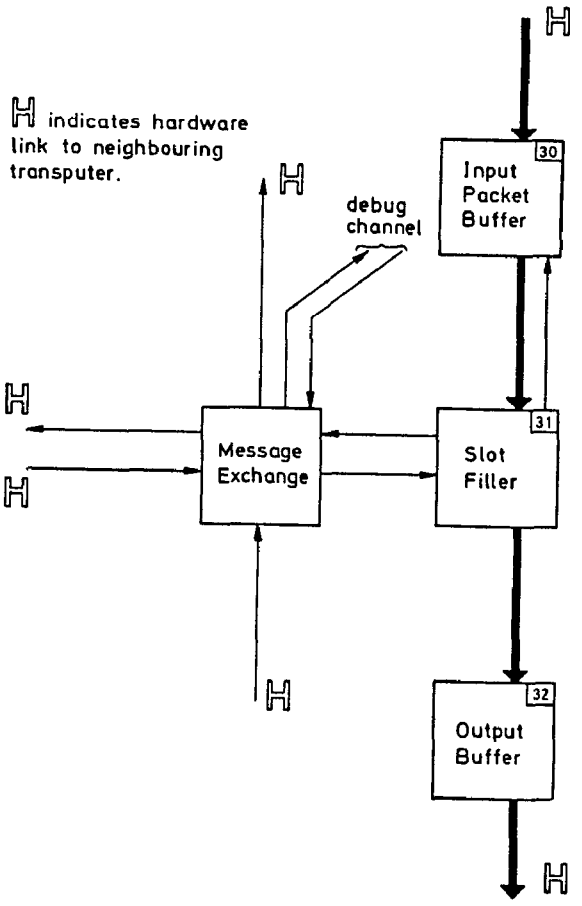

Fig. 3 : Process structure within transputer 2, the Tx Mapper.

4 Performance and possibilities

As part of the Alvey information technology initiative in the UK, a small pilot ISDN network was built by British Telecom and Logica, and made available to research groups. The network has a single exchange that provides primary rate access to these groups. This network has been used by the Unison collaborators to link LANs over a wide area [5]. Ramps have been used at four sites interfacing CFRs to ISDN. Some performance measurements have been made on these ramps. These measurements are reported here. It is planned to use the ramps in conjunction with British Telecom's operational primary rate ISDN service (Multiline IDA) later this year. It is intended to obtain further performance measurements using this network, including measurements of call setup and removal times.

Ramps operate on word frames which have a cycle time of 500 μsec. Markers are sent every 100 frames, that is every 50 msec, giving an overhead of 1%. The minipackets used on the CFR contain 32 bytes of data and two 16-bit addresses, giving a total length of 9 words. The ramp adds an extra word containing a 16-bit CRC and a 16-bit delimiter pattern, so the packet length transmitted on ISDN is 10 words. Thus the maximum rate of transfer of packets across the ISDN on a channel of 30 slots is 6000 packet/sec, less the 1% overhead for the markers. The ramp is able to sustain this rate simultaneously in both directions, that is support a flow of almost 12 000 packet/sec. The packet delay through a pair of ramps is of the order of 5 msec. The synchronization process on a slot normally takes less than 100 msec with a worst-case time of 200 msec. The ramp is, however, not able to synchronize all 30 slots at the same time. The times for the total procedure of creating a wideband channel and changing its bandwidth are of great interest. These await measurement on an operational ISDN network.

Used in the fashion described above, with the minipacket structure of the CFR, the ramp is interesting in the context of simulating an ATM-style network. This matter is discussed in a companion paper [6]. For use with current operational networks, such as ethernet

and token ring, fragmentation and reassembly of packets would be necessary and in this case a rather longer fixed length packet would probably be optimal for use on the ISDN. The current configuration of transputers is probably powerful enough to carry out the fragmentation process but, if this were not the case, there would be little problem in adding an additional transputer for this purpose.

References

[1] International Telecommunications Union, *I-series Recommendations.* CCITT VIIIth Plenary Assembly, Fascicle III.5, 1985.

[2] Inmos Limited, *Transputer Reference Manual.* Prentice Hall, 1988.

[3] Inmos Limited, *Occam 2 Reference Manual.* Prentice Hall, 1988.

[4] A. Hopper, S. Temple, and R. C. Williamson, *Local Area Network Design.* Addison-Wesley, 1986.

[5] D. L. Tennenhouse, I. M. Leslie, R. M. Needham, J. W. Burren, C. J. Adams and C. S. Cooper, *Exploiting Wideband ISDN: The Unison Exchange.* Proc. IEEE INFOCOM '87.

[6] D. L. Tennenhouse and I. M. Leslie, *A test-bed for wide-area ATM research.* ACM SIGCOMM '89 symposium on Communications Architectures and Protocols, Austin, Texas.