

A Detailed Study of Asynchronous Transfer Mode Techniques

¹Sandip Banerjee & ²Dr. Suresh Dara

¹Research Scholar, Department of Computer Science & Engineering, Sri Satya Sai University of Technology and Medical Science, Sehore, MP (India)

²Department of Computer Science & Engineering, B.V RAJU Institute of technology, Narsapur, Medak, Telangana (India)

ARTICLE DETAILS

Article History

Published Online: 07 September 2018

Keywords

Telecommunication, Public Network, cost

ABSTRACT

TELECOMMUNICATIONS STANDARDS have usually followed sometime after the development of products for a particular need. This has led to subtle, or sometimes gross, incompatibilities between two functionally similar products, making interworking both inefficient and expensive. When telecommunication networks were largely considered as a local, or national, resource then the incompatibilities were not particularly significant; but, in many respects, the public network now needs more and more to be regarded as an international resource and such matters as interworking cannot be 'put on one side' until the local network is functioning. The cost of designing and installing the equipment is now so great that to justify replacing the current network with a new one requires that significant benefit be obtained by doing so and, further, that a similar need must not occur again for a long time (if ever). This paper describes in detail the background and functionality of a B.I.S.D.N., showing how both new and old, error-sensitive and delay-sensitive, services can co-exist on a single network. The prospects of implementing such a network are still a few years in the future, and it will not appear 'overnight': in the interim it will be necessary for the new and old networks to interoperate. At present, the plain old telephone service (POTS) is still, by far, the largest revenue earner for the network operators and such a service must continue to be provided with an equivalent, or higher, grade of service (GOS, or quality of service, QOS) in any new system that is installed.

1. Introduction

Traditional digital voice networks and data networks use in-compatible transmission techniques and have different performance criteria: voice telephony requires only moderate levels of data integrity, but has stringent constraints on the delay that may be incurred while passing through the network; for data communication, on the other hand, the constraints are the other way round. To resolve these differences voice communication uses synchronous techniques such as time division multiplexing (T.D.M.) [1] which has low latency at the exchanges since no queueing is involved: data communication is normally packetized; the traffic characteristics of the data being transmitted are hard to quantify in advance, so unlike voice, where a fixed amount of the communications resource can be allocated to each call, the data are split up into variable length packets and given access to the entire available bandwidth; transient overloads at the exchanges (switches) are handled by queueing.

It was clearly desirable for the services provided by the two systems to be merged into a single network using a single access point: the capital value of the local network is immense, and the cost of duplicating it, including the disruption this would cause, to provide data services is almost unimaginable. Fortunately, recent advances in technology have made it possible to use each copper pair in the access network (which was originally designed to support a single 3 kHz band-width channel) at a transmission rate of 144 kbit/s (sufficient for two 64 kbit/s voice channels plus some spare). In 1984 the C.C.I.T.T. published a series of recommendations

(the I-series) that defined an integrated services digital network (I.S.D.N.) which were subsequently refined in 1988. The recommendations standardized a total of six different types of channel (labelled A to E and H), the most important of these being: the B channel, a 64 kbit/s channel for voice or data; the D channel, a 16 or 64 kbit/s channel for signalling or packetized data; and the H channel, which can run at 384, 1536 or 1920 kbit/s for higher-rate access. To standardize access further, two main channel combinations were defined: the basic rate, which contains 2B + ID channels and is suitable for use with the standard copper pair; and a primary rate, which contains 30B + ID channels in Europe and 23B + ID channels in the U.S.A. and Japan. [2] The I.S.D.N. defines the B channel purely in terms of its size and position in the frame for transmission to the exchange (O.S.I. layer 1); no meaning is attached to any of the bits within the channel: this makes it suitable for use with any service requiring not more than 64 kbit/s bandwidth; services requiring more than this must use either multiple circuits or one of the H channels (the two are not interchangeable), but in either case the network does not 'see' any of this higher-layer protocol adaptation.

There are several problems, however, with the I.S.D.N. architecture that make it potentially unsuitable for many communication needs in the not too distant future. The internal workings of an I.S.D.N. do not naturally form a single network, once information reaches the exchange it is normally split into the appropriate data class and routed on the appropriate type of network (for example, voice traffic over a T.D.M. network and data over a packet-switched network); the B and H channels are highly restrictive in nature as they each have a

fixed capacity that cannot be exceeded, whilst if a service requires less than the capacity of a channel the balance is wasted; in a similar manner, the suitability of the channels for packetized data is very poor, the D channel has very low capacity and the B and H channels do not support their stochastic traffic nature very efficiently. The addition of new services is also quite complex; this follows from the individual carrier-networks that tend to make up an I.S.D.N.; adding a new service may well require a new type of network in addition to more software at each exchange. Finally, and perhaps most importantly, the I.S.D.N. does not define any access rate higher than 2 Mbit/s, this means that services such as high definition video cannot be carried. Most of these problems can be solved with the newly emerging standards for a Broadband I.S.D.N. (B-1.S.D.N.); this defines both the interface and the network itself in a far more flexible manner which provides not only for much higher data transfer rates but also makes the installation of new services much simpler for network managers.

2. Background

When a broadband version of the I.S.D.N. was first proposed. It was unclear what mechanism would be required to support the new services which it would provide. At that time most of the high-bandwidth services that would need to be carried by such a network (for example video based services) could only be encoded using constant bit-rate (C.B.R. or Fixed B.R.) techniques, so it was suggested that the best approach should be to define yet more higher access rate channels, similar to the H channels: however, another of the main services that the network would have to support is high-speed connection between LANs, which tends to have a bursty traffic characteristic. When it subsequently became possible to encode video signals using a variable bit-rate (V.B.R.) scheme it became obvious that fixed channel allocations would be far too wasteful when segmenting the communications resource: further, for efficiency of channel allocation, it would probably be necessary to partition the S.T.M. frames in a fixed manner with each class of channel only being allowed to start on the appropriate boundary within the frame; this can lead to fragmentation, making it impossible to carry the higher-rate channels even when there is theoretically sufficient spare capacity.

It became clear that a new transfer mechanism would be required to make more efficient use of the capacity of the network. Two differing schemes were under investigation in the United States and Europe (with different applications in mind); it was proposed that the two should be analysed in more detail and a common solution adopted: the solution was initially referred to as the new transfer mode but was subsequently changed to asynchronous transfer mode (A.T.M.) because of the analogy with S.T.M.. The American proposal, fast packet switching (F.P.S.), was originally aimed at producing a high speed network suitable for connecting LANs together: the European proposal was more concerned with handling packetized voice samples in a delay-critical environment and is known as asynchronous time division (A.T.D.).

2.1 Fast Packet Switching

The C.C.I.T.T.'s X.25 and associated protocols were originally designed with relatively poor-quality copper-based transmission techniques in mind: the error rates are high, so the probability of a packet being transmitted from one end of the network to the other without some errors creeping in is not negligible; complex retransmission policies are employed across each link in the transmission path so that each packet is received correctly on one link before it is forwarded to the next; and windowing techniques are employed to enable sufficient throughput when the link delays are high. All these techniques require substantial amounts of software for verification and management of the protocol, and memory for buffering: as the reliability of the links increases to the extent that errors are no longer likely over the end-to-end path, the protocol becomes a hindrance to efficient use of the network. It has been shown that when the bit error ratio drops below 10^{-5} the use of end-to-end retransmission becomes no less efficient than link-by-link [3].

When fibre-optic links are used the bit error ratio can be reduced to levels significantly below 10^{-5} , so the proposal for fast packet switching involved the use of end-to-end retransmission. Moving error control and, also, flow control to the edges of the network significantly simplified the link protocols; the remaining functions were then easily implemented using hardware, resulting in extremely fast switches, and made it possible to carry delay-sensitive services such as voice. Since a primary service for the network was to be high-speed LAN interconnection, variable sized packets were proposed.

2.2 Asynchronous Time Division

Whilst F.P.S. was being derived from traditional packet switching techniques, asynchronous time division was evolving from the synchronous time division (S.T.D.) techniques used for traditional circuit switching. By adding more intelligence to the switch and by adding short headers to each block of data, the need for time-position dependence was removed and the flexibility of the system enhanced. Research was performed to enable A.T.D. to be defined as a layer 1 protocol in the O.S.I. reference model; the functionality of the header was reduced to the absolute minimum required for switching the frame of data: fixed sized frames of between 8 and 32 octets were proposed, so the header needs to contain only connection (routing) and priority information. Because of its origins, A.T.D. was always intended primarily as a low-latency transport mechanism for delay-sensitive services such as voice and video, with the emphasis on video [4].

3. Asynchronous Transfer Mode

Rather than select one of the two alternatives, the C.C.I.T.T. elected to choose a compromise solution: some of the features of each of the proposals were chosen to form the A.T.M. implementation to be used for a Broadband I.S.D.N.. It was decided, for example, that fixed-size packets, known as cells (to avoid confusion with traditional packet switching, where the packets are of variable length), would be used, but their size would be longer than proposed in the A.T.D. definition; a small amount of extra functionality was

added to the header over that proposed for A.T.D., but explicit priority information was omitted.

In O.S.I. parlance, A.T.M. provides layer 1 and some layer 2 functionality; the precise relationship between the A.T.M. reference model and the O.S.I. model is still to be fully defined (C.C.I.T.T. recommendation I.321) and, in any case, the mapping is fairly vague: many functions which have traditionally been part of the data-link layer are now omitted entirely or moved to functions in the transport layer. In traditional packet-switched networks the protocol stack at intermediate nodes in a packet's route covered layers one, two and three; for an A.T.M. network, however, the A.T.M. layer is the highest protocol layer within intermediate nodes. Figure 1 shows the layers within the A.T.M. protocol stack. The Physical Layer has two sub-layers: the physical media sub-layer, which is the lower layer and is responsible for bit timing and transmission; and the upper layer is the transmission convergence sub-layer, which handles aspects such as cell delineation and header error control. The A.T.M. layer has further sub-layers of its own, which include: Generic Flow Control, which only operates at the edges of the network and is used to throttle back sources when the network is in danger of overload; cell header generation and extraction; virtual-circuit and virtual-path number translation when switching the cells; and cell multiplexing or demultiplexing between supported services.




Figure 1: The layers of the A.T.M. protocol stack

3.1 Cell Structure

The A.T.M. cell structure is defined in Recommendation I.361; the cell is of fixed size with a 48 octet payload field and a 5 octet header. The format of the payload remains fixed and unmodified during the whole time that the cell remains in the A.T.M. network; since error control is performed at the edges, and then only when a service requires such a feature, the field does not contain an explicit error-correcting, or -detecting, capability. The header has two formats, depending on where the cell is within the network: between user equipment and the network (U.N.I., user-network interface) provision is made for generic flow control; for communication between nodes within the network (N.N.I., network- network interface) no such provision is made and the relevant bits are made available to increase the number of virtual circuits.

The header at the U.N.I. is shown in Figure 2; it contains:

G.F.C. Generic Flow control. This field is to enable the network to regulate the arrival of cells from a source external to the network. Details of the values to be used have yet to be defined 4 bits.




Figure 2: The cell-header format for the U.N.I.; dimensions in bits.

V.P.I./V.C.I. Virtual Path Indicator, Virtual Circuit Indicator. These fields are normally used in conjunction with each other; together they are used to route cells through the network from one edge connection to another. The values assigned have only local significance and may change at various intermediate nodes. 8 bits V.P.I., 12 bits V.C.I..

P.T. Payload type. This field is used to describe the type of data that the cell contains; assignments of the values have yet to be made, but it might, for example, be used to distinguish signalling traffic from user data, or to indicate some form of priority. 2 bits.

C.L.P. Cell Loss Priority. A cell with this field set is more likely to be discarded at times of network overload than one with the field cleared. 1 bit.

H.E.C. Header Error Control. This field contains a CRC-8 code which is used to protect the header (only) against corruption during transmission. Whilst it is part of the A.T.M. header, it is computed and tested in the Physical Layer functions before cells are passed up to the A.T.M. layer. 8 bits.

res Reserved. A single bit field that is reserved for future use. 1 bit.

The header for use at the N.N.I. is shown in Figure 3; it is identical to the header for the U.N.I. with the following exceptions:

- There is no G.F.C. field.
- The V.P.I. field is expanded to 12 bits to provide increased routing capacity within the network.




Figure 3: The cell-header format for the N.N.I.; dimensions in bits.

3.2 Routing

The primary role of the cell header is to provide sufficient information for the cell to be routed through the network. Since the size of an international network has to be considered it is impractical to use datagram headers and destination addresses within each cell (the 20 address bits available in the header at the U.N.I. would provide for approximately 1 million possible destinations, which is insufficient to cover the needs of U.K. alone), so a virtual circuit approach has been adopted.

To simplify routing of connections at the trunk exchange, the routing bits have been partitioned into two parts: a virtual circuit, with each connection having its own circuit identification, and a virtual path, which is then used to group connections using the same route, or part of a route, through the network. At the trunk exchanges, connections are routed solely on their virtual path identifier, but at local exchanges and elsewhere the full address part may be used: normally, at an exchange the circuit number will change to represent the value required for the local link; at trunk exchanges only the V.P.I. changes, while the V.C.I. remains constant; elsewhere, both the V.P.I. and the V.C.I. may change between input and output.

For some services, notably user-network and intra-network signalling, establishing a virtual circuit before communicating might be inefficient, or even impossible (the signal to create the circuit would need to be sent, but this requires that a circuit already exists.); in such cases permanent, or semi-permanent, circuits can be established. It is even possible that some user-user services will also work better in a connectionless environment; in this case a connectionless type service can be established by reserving a V.P.I. value and then assigning the V.C.I. values to represent

individual destinations within the network: individual communicating pairs are resolved using the A.T.M. adaptation layer. To implement this the network has to be partitioned into zones [5].

3.3 The A.T.M. Adaptation Layer

The A.T.M. adaptation layer is the functional layer above the A.T.M. layer; as such, it strictly lies outside the formal definition of the A.T.M. network: its functionality is service dependent and the fields that it 'defines' are only interpreted by equipment that lies beyond the edges of the network itself. However, despite the fact that implementation is not mandatory, all services will require some form of adaptation, and most are likely to fall into one of the four classes of adaptation layer framing. The layer handles such functions as: segmentation and reassembly of messages; error recovery and retransmission; handling of lost or mis-inserted cells; flow control; and timing control.

Four basic classes of adaptation are defined in Recommendations I.362 and I.363. The features supported are shown in table 1.

	Class A	Class B	Class C	Class D
Timing relation between source and destination	Required		Not Required	
Bit Rate	Constant	Variable		
Connexion Mode	Connexion Oriented			Connexionless

Table 1: Suggested A.T.M. Adaptation Layer classes.

4. Enabling Technology

The ability to build the high-speed networks required for broadband I.S.D.N. is a direct result of advances in hardware technology. Not only is greater capacity available because of the increase in the speed at which it works but, also, the probability of failure is decreasing: if the rate of failure were not to decrease then a system running ten times as fast would be likely to fail ten times as often; in practice the situation is better than this.

4.1 Fibre Optics

One of the main technologies that makes a B.-I.S.D.N. possible is the use of fibre optics with semiconductor-laser technology. Copper-based communication links have always been subject to electromagnetic interference (E.M.I.), which is commonly of a bursty nature and can distort the connection for several bit periods so that the signal is beyond recovery; single bit errors tend to be more of a secondary issue, but they can still cause problems. The non-negligible resistance and leakage of the copper connection causes attenuation; on runs of more than a few kilometers repeaters are required to boost the signal back to its original level and to 're-shape' the pulse edges.

With fibre optics the primary source of interference, E.M.I., no-longer applies, so the probability of an uncorrectable burst of data occurring is almost zero.

Attenuation within the fibre is significantly lower than for an equivalent length of copper and consequently repeaters are needed less frequently, thereby increasing the reliability of a link. In addition, mono-mode fibres have (as their name implies) only one mode by which the light signal can propagate down the fibre; this eliminates multi-path distortion and further increases the possible distance between repeaters. The current capacity of fibre-optic systems is not limited by the fibre itself, but by the speed with which the lasers and detectors can be made to operate reliably; it is to be expected that this will continue to increase in the future as the technology matures.

4.2 Switching Technology

To support the very large amount of traffic that a B.-I.S.D.N. will be required to carry, new types of switching technology are necessary. The concept of a circuit running through an exchange, carrying a single call from input to output, was weakened with the introduction of digitized voice services: it is almost totally destroyed by the packetization of the samples; circuits no-longer exist in any physical sense, only the logical sense remains. In traditional packet-switched networks the amount of processing that was performed on each packet at each switch in the network made it feasible for the switching process to be performed in software, without the need for complex hardware to parallelize the process. In order to achieve the higher throughput of a B.-I.S.D.N. the lower layers of software were simplified to the

extent that most of the remaining functions could be implemented in hardware; the switching routines would be

swamped if a software process were used.




Figure 4: 2-input, 2-output switching element that forms the basis of the Starlite switch. a) crossed over outputs, b) straight through.

A highly parallel switching architecture is the Starlite switch [6]. It uses fixed-size packets and routes these in parallel waves across the switching architecture; it is non-blocking provided that two, or more, packets are not routed to the same destination at the same time. The basic element is a 2-input, 2-output switching device that can either pass both inputs straight through or can cross them over to the opposite outputs (Figure 4): the main features of the switch are built up from replicas of this element, making the whole fabric well suited to implementation using V.L.S.I. technology). There are two main phases to the switching process (Figure 5): a sort phase, which uses the elemental switches arranged as a Batcher network to perform a 'perfect shuffle', and an expansion phase, which switches the sorted packets to the correct outputs. The expansion phase is sometimes known as a Banyan Network, leading to the whole

switch being known as a Batcher-Banyan Network. To cope with packets addressed to the same output a trap phase has to be used between the sorter and the expander (Figure 6); this re-routes packets, that would otherwise be lost, back to the input of the sorter for processing on the following pass; the loss rate of the exchange can be determined by the amount of capacity within the trap that can be re-routed in such a manner. The concentrators are used to keep the number of switching elements as low as possible: the sort phase requires $n(\log_2 n)(\log_2 n + 1)/2$ basic elements for n inputs; the expander requires $m(\log_2 m)/2$ basic elements for its m inputs: for large switches the probability that all inputs would be active at once is very small and large economies can be made by using concentrators to reduce n .

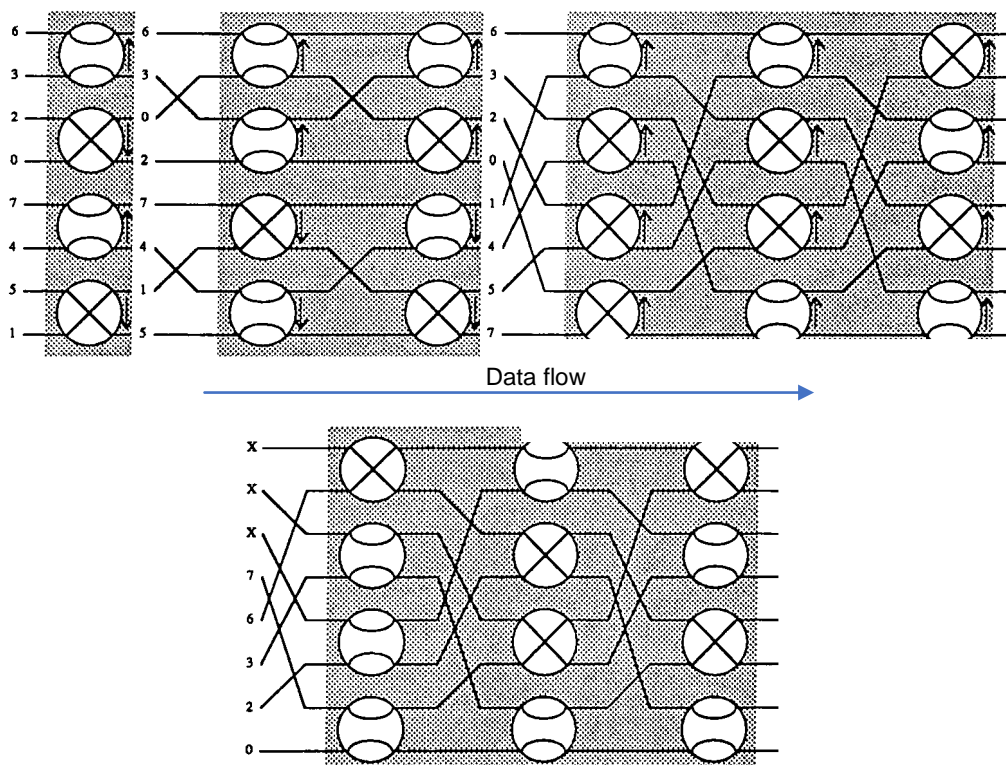


Figure 5: Connexion graphs for the two main phases of a Starlite switch. The top graph performs the perfect shuffle to sort the inputs into order (the arrows point to the output that should have the larger of the two inputs); the lower graph is the expansion phase that uses the bit fields of the output port to determine the route.

Another possible switch architecture that has been proposed is the Orwell Torus [7]; this would be used mainly in

low capacity exchanges where its greater flexibility would be more important than very high throughput. The protocol uses

multiple slotted rings, running in parallel, and novel access control mechanisms, to bound access delays for time-critical

services. The protocol is discussed in more detail in the next chapter.




Figure 6: Stages of a Starlite switch. The concentrator phase is used to reduce the number of inputs to the sorter; the trap phase to catch two or more cells that are routed to the same output which would get blocked in the expander; all but one are recycled back to the input for switching on a subsequent pass.

5. Traffic and Services

In order to dimension correctly the capacities of various elements within a B.-I.S.D.N. the types and expected loads of services offered need to be analysed in detail; provision for future growth of the services also needs to be included; this is likely to prove extremely difficult to calculate since the services requiring most bandwidth are, as yet, untried, and their likely penetration into the market are, at present, unknown. It is the necessity to reduce the risk of forecasting errors that makes the requirement for a single transmission network so important; errors in under-forecasting for growth in one service may well be compensated by over-forecasting in another.

5.1 Connexion Control

If all the services carried by a B.-I.S.D.N. used constant bit-rate encoding then calculating whether a service could be safely carried without affecting the delays and causing overloads would be a relatively simple extension of the techniques used for traditional voice networks (the main difference being that services could require differing amounts of bandwidth and that these would not be integral multiples of some base rate); in such circumstances connexion control is not particularly difficult. However, some services, such as interactive computer-data transfer, are naturally bursty in nature and others, such as compressed video signals, often require a variable bit-rate coding scheme; when several of these services are carried simultaneously on a single network then some network capacity can be saved by understanding the statistical processes involved and working out the likelihood of an overload occurring.

There are three statistical parameters to a variable bit-rate coding scheme that are of interest to network operators: mean, standard deviation (or variance) and peak bit rates. It is, at present, unclear which of the latter two is the most important in terms of dimensioning the network: one approach suggested for connexion control is for the call to be accepted

purely on the basis of its peak bandwidth requirement; while this would enable cell loss due to overload to be completely avoided, significant benefits available from statistical multiplexing would be lost, making the network inefficient. When connexion acceptance is based on statistical calculations then the 'acceptance surface' (a multi-dimensional surface indicating the load limits for various combinations of traffic) is no-longer a plane, but becomes distorted due to the differing standard deviations: attempts have been made to calculate this surface in advance [8], and others have suggested that the surface might be determined when the network is in operation (for example, by using neural networks [9]).

5.2 Charging and Policing

Two areas that are still a major subject of debate are charging customers for use of the network and policing connexions, particularly V.B.R. connexions, to ensure that they do not try to transmit more than had been originally stated. The two areas are similar in that the same techniques and problems apply to both, although different techniques may be applied. The main options available are to monitor the mean rate or the peak rate, or possibly some function of both: charging only for the peak bandwidth used is, however, difficult to justify unless the bandwidth is guaranteed; conversely, charging for the mean bandwidth used does not reflect the spare capacity that has to be set aside to cope with the statistical fluctuations. Policing is important, particularly when statistical connexion acceptance is performed, to ensure that the traffic load presented to the network matches that which was negotiated when the call was initiated. The problems in both cases are concerned with the sheer bulk of information that has to be processed; a simple network may well be carrying several million cells every second, and to monitor each one presents an insurmountable burden for the policing and charging functions.

References

1. B. Lubachevsky, A. Shwartz, and A. Weiss, "Rollback Sometimes Works ... If Filtered (Abstract)," in Proceedings of the 1989 Winter Simulation Conference, pp. 630-639, December 1989.
2. P. L. Reiher, F. Wieland, and D. Jefferson, "Limitation of Optimism in the Time Warp Operating System," in Proceedings of the 1989 Winter Simulation Conference, pp. 765-770, 1989.
3. L. Gould, I. Bowler, and A. Purvis, "Real-Time, Multi-Channel Digital Filtering on the Transputer," in Proceedings of the 1989 International Symposium on Computer Architecture and Digital Signal Processing, 1989.
4. D. W. Jones, Chien-Chun Chou, D. Renk, and S. C. Bruell, "Experience with Concurrent Simulation," in Proceedings of the 1989 Winter Simulation Conference, pp. 756-764, 1989.
5. D. M. Nicol, C. C. Micheal, and P. Inouye, "Efficient Aggregation of Multiple LPs in Distributed Memory Parallel Simulations," in Proceedings of the 1989 Winter Simulation Conference, pp. 680-685, 1989.
6. C. A. Buzzell, M. J. Robb, and R. M. Fujimoto, "Modular VME Rollback Hardware for Time Warp," in Distributed Simulation, pp. 153-156, January 1990.
7. R. J. Lipton and D. W. Mizell, "Time Warp vs. Chandy-Misra: A Worst-case Comparison," in Distributed Simulation, pp. 137-143, January 1990.
8. W. Cai and S. J. Turner, "An algorithm for Distributed Discrete-event Simulation - The "Carrier Null Message" Approach," in Distributed Simulation, pp. 3-8, January 1990.
9. R. C. De Vries, "Reducing Null Messages in Misra's Distributed Discrete Event Simulation Method," IEEE Transactions on Software Engineering, vol. 16, pp. 82-91, January 1990.
10. D. Baezner, G. Lumow, and B. W. Unger, "Sim++: The Transition to Distributed Simulation," in Distributed Simulation, pp. 211-218, January 1990.
11. N. J. Bailey, A. Purvis, I. W. Bowler, and P. D. Manning, "An Highly Parallel Architecture for Real-time Music Synthesis and Digital Signal Processing Application," in Proceedings of the International Music Conference, Glasgow, 1990.
12. J. Dupraz and M. De Prycker, "Principles and Benefits of the Asynchronous Transfer Mode," Electrical Communication, vol. 64, no. 2/3, pp. 116-123, 1990.
13. S. J. Nichols, Simulation and Analysis of Adaptive Routeing and Flow Control in Wide Area Communication Networks. PhD thesis, University of Durham, March 1990.
14. D. P. Helmbold and C. E. McDowell, "Modeling Speedup (n) Greater than n," IEEE Transactions on Parallel and Distributed Systems, vol. 1, pp. 250-256, April 1990.
15. J. Chauhan, T. King, and A. C. Micallef, Specification of the Orwell Protocol. British Telecom Research Laboratories, Martlesham Heath, Ipswich, Suffolk, UK. IP5 7RE, May 1990. Revision C.1(05/90).
16. P. W. Glynn and P. Heidelburger, "Analysis of Parallel Replicated Simulations under a Completion Time Constraint," A.C.M. Transactions on Modeling and Computer Simulation, vol. 1, pp. 3-23, January 1991.
17. J. Appleton, "Traffic Shaping in Asynchronous Transfer Mode Networks," in Proceedings of the Eighth U.K. Teletraffic Symposium, pp. 19/1-19/4, April 1991.
18. J. R. Chen and P. Mars, "Adaptive ATM Call Access Control Using Learning Algorithms," in Proceedings of the Eighth U.K. Teletraffic Symposium, pp. 18/1-18/5, April 1991.
19. N. J. Bailey, A. Purvis, I. W. Bowler, and P. D. Manning, "Some Observations on Hierarchical, Multiple-Instruction-Multiple-Data Computers," in Proceedings of Euromicro, Vienna, September 1991.
20. N. J. Bailey, On the Synthesis and Processing of High Quality Audio Signals by Parallel Computers. PhD thesis, University of Durham, October 1991.