

UC Irvine

ICS Technical Reports

Title

LAN with collision avoidance : switch implementation and simulation study

Permalink

<https://escholarship.org/uc/item/7xg648hn>

Authors

Huang, Hung Khei
Suda, Tatsuya
Noguchi, Yasuhiro

Publication Date

1991

Peer reviewed

Notice: This Material
may be protected
by Copyright Law
(Title 17 U.S.C.)

Z
689
C3
no. 91-13



LAN with Collision Avoidance:
Switch Implementation and Simulation Study¹
Technical Report 91-13

Hung Khei Huang, Tatsuya Suda, Yasuhiro Noguchi
Department of Information and Computer Science
University of California, Irvine
Irvine, CA 92717

¹This material is based upon work supported by the National Science Foundation under Grant No. NCR-8907909.

This research is also in part supported by University of California MICRO program.

Abstract

Packet collisions and their resolution create a performance bottleneck in random access LANs. A hardware solution to this problem is to use collision avoidance switches. These switches allow the implementation of random access protocols without the penalty of collisions among packets.

The simplest network based on collision avoidance is the Broadcast Star network, where all the stations are connected to a central switch. A more sophisticated architecture based on collision avoidance is the CAMB (Collision Avoidance Multiple Broadcast) Tree network, where concurrent broadcast is possible.

This paper presents a design of a collision avoidance switch for the CAMB Tree using TTL devices. Simulation study exploring the performance of the Broadcast Star network in both synchronous and asynchronous operations is also presented in this paper.

1. Introduction

Advantages of random access protocols include simplicity and ease of implementation at the stations in a network. Random access protocols also exhibit small transmission delays under light traffic conditions. However, they have a performance bottleneck under heavy traffic conditions. In that situation, a large number of collisions occurs, resulting in a low channel utilization. Channel capacity is wasted in the transmission of collided packets.

In order to solve the performance bottleneck in random access protocols, a new network architecture based on collision avoidance, called collision avoidance tree network, has been proposed and studied [1 - 4]. The tree network uses collision avoidance switches. These switches allow implementation of random access protocols without the penalty of collisions among packets.

The simplest form of the tree network is called Broadcast Star network. (In [2, 5], it is referred to as the Hubnet.) It has one central collision avoidance switch. Each station is connected to the central switch through a channel. This channel consists of two logical channels: uplink channel to send data from a station to the switch, and downlink channel to send data from the switch to a

station. Contention among stations is solved by the collision avoidance switch. When two or more packets contend for the output of a switch, it is guaranteed that one of the packets acquires the switch and is successfully transmitted through the switch. Thus, no channel capacity is wasted in the transmission of collided packets.

The Broadcast Star network is first proposed and studied in [1 - 3, 6], and various studies followed. Network protocols and a circuit design for a collision avoidance switch in a Broadcast Star is given in [1, 6]. In [2, 5], the Hubnet is studied. The Hubnet is a Broadcast Star network implemented on a tree topology network. Papers [7 - 9] develop queueing analyses for the performance of the Hubnet. Approaches in the analysis of polling systems are used in these papers. New access protocols for the Hubnet are proposed and analyzed in [8, 10]. In [11], synchronous operation of the Broadcast Star is considered, and its performance is analyzed. An exact analysis is developed for the network with an infinite number of stations, and an approximate analysis is developed for the network with a finite number of stations.

A more general form of a collision avoidance network, called the CAMB (Collision Avoidance Multiple Broadcast) Tree network, has been proposed in [3]. In the CAMB Tree, the switches are organized into a tree topology, and the stations are the leaves of the tree. A packet reaches its destination by climbing the tree and being broadcast by the switch that is the root of the minimal subtree containing both the source and the destination stations. The CAMB tree allows multiple broadcasts in the network. Station and switch protocols for the CAMB tree is further studied in [12, 13]. A switch design based on photonic devices and a simulated performance for this network are given in [4, 12]. A performance study on the CAMB tree network through theoretical analysis and simulations are given in [14, 15]. Comparisons with the Hubnet is also done, showing a better performance of the CAMB tree. Variation of the collision avoidance network using point-to-point network architecture can be found in [16, 17].

In this paper, we present a switch implementation for the CAMB tree based on TTL standard modules. The performance of the Broadcast Star network is also studied through simulations. Two modes of network operation, synchronous and asynchronous operations, are considered and

compared. It is shown that a synchronous network achieves better performance (smaller delay and higher throughput).

In section 2 a description of the Broadcast Star network is presented, and the CAMB Tree network is explained. A switch hardware design for the CAMB tree network is given in section 3. Section 4 presents a simulation study on the performance of the Broadcast Star network. Finally, in section 5 concluding remarks are given.

2. Review of Broadcast Star and CAMB Tree Networks

2.1. Broadcast Star Network

In a Collision Avoidance Broadcast Star, the stations are connect by full duplex channels to a central switch, where each of these channels comprise an uplink and a downlink. See Fig.1. The switch may be viewed functionally as containing two components: the selector and the broadcaster. The selector selects one packet from the uplinks and transmits this packet on a single output line to the broadcaster. The broadcaster broadcasts the packet on all its downlinks. While a selector is busy (i.e., while a packet is currently being transmitted through the switch), all the arriving packets are ignored (or blocked). The important feature in this switch is that there is no collision. It is guaranteed that one packet will always be successfully transmitted, even when two or more packets contend for the right to use the switch. Thus, no channel time is wasted in the transmission of collided packets, eliminating the main disadvantage of a traditional random access network.

The station protocol for the Broadcast Star is very simple and similar to the ALOHA protocol. A station transmits a packet as soon as one becomes available. The station monitors its downlink for the start of its packet for a time equal to a propagation delay time to and from the switch (round trip propagation delay). If the station does not see the start of its packet, then it retransmits the packet immediately. If the station does see its packet, then the packet will be successfully transmitted in its entirety.

There are two possible operations of the network: synchronous and asynchronous operations. In synchronous operation (slotted Broadcast Star), the channel time is divided into slots, and stations

- Address Recognizer (AR), which checks the header of the selected packet to see if the switch is the proper ancestor. If it is, the AR sends the packet to the DS and to the parent switch, otherwise the AR sends the packet only to the parent switch.
- Downlink Selector (DS), which receives a packet from the AR and the downlink from the parent switch, and broadcasts the packet to all the children switches. In case of simultaneous arrivals of a packet from the AR and a packet from the parent downlink, priority is given to the packet from the parent downlink. In other words, if the DS is broadcasting a packet received from the AR upon the arrival of a packet from the parent downlink, the transmission of the packet from the AR will be aborted, and the packet from the parent is broadcast.

3.1. Packet Format

In our switch design, we assume the packet format and addressing scheme described in this section. See Fig.4. The first bit in a packet is a control bit. This bit is used to indicate if the packet has reached its proper ancestor. This bit is initially set to 0 by the sending station. It remains 0 until the packet reaches the proper ancestor. When a packet reaches its proper ancestor, the proper ancestor switch sets this bit to 1.

The next some number of bits (8 bits in Fig.4) are used as the destination address field. It is followed by the packet size field. The size field indicates the number of bits contained in the data field. This information is used by the stations (not by the switches) to determine whether the packet received has been aborted. If the packet size field is equal to the size of the data field, the packet is successfully transmitted, otherwise, abortion happened. Note that, depending on higher layer protocols, other overhead information such as the source address and the CRC may be necessary in the packet.

3.2. Addressing Scheme

The addressing scheme assumed in our switch design is shown in Fig.5. Stations are ordered in ascending order from the left to the right of the tree. Each station has a $2 \times (M - 1)$ bit long

address, where M is the height of the tree. Addresses are assigned to the switches in the following way. The level i switches have a $2 \times (i - 1)$ bit long address. The address of a switch on level i matches with the first $2 \times (i - 1)$ bits of the addresses of all of the stations in its subtree.

The above numbering of stations and switches makes it easy for a switch to know whether it is the proper ancestor of a packet. A switch on level i checks the $2 \times (i - 1)$ leftmost bit of the address field of a packet. If it is equal to its own address, then it is the proper ancestor of the packet. In this case, the packet is sent to both the parent switch and the downlink switches. Otherwise, the packet is only sent to the parent switch.

Fig.5 shows a transmission of a packet from station A (address 00000000) to station B (address 00000100) as an example. Station A 's immediate parent switch, Z (address 000000), checks the destination address of the packet. As the six leftmost bits of the destination address in the packet header is different from the switch address, it sends the packet only to its parent switch (Y). Switch Y (address 0000) then checks the four leftmost bits of the destination address of the packet. This time, as they match with the switch's address, switch Y knows that it is the proper ancestor of the packet. Thus, the packet is broadcast to its children. The packet is also sent to the parent switch (X).

Now, the switch X (address = 00) checks the two leftmost bits of the destination address field. They agree with the switch address. Therefore, the switch X knows that it is the proper ancestor of the packet and becomes ready to broadcast the packet to its children. However, if the switch Y actually broadcasts the packet, the packet is broadcast twice (once by switch Y and once by switch X). To avoid this undesirable situation where the same packet is broadcast more than once by different switches, a control bit has been added in the packet header (Fig.4). First, a switch checks the control bit in the packet header. If it is 1, the switch knows that the packet has already been broadcast by its proper ancestor, so it will transmit the packet only to the parent switch. Otherwise, it will check the address field to see if it is the proper ancestor of the packet. In the example shown above, the switch X first checks the control bit. As it is 1, the switch does not broadcast the packet.

In the following, we present a possible implementation of the CAMB switch. In the design, we assume the packet format and the addressing scheme discussed above. Design of each of the three components (US, AR, DS) of the switch is explained below.

3.3. Switch Design

Uplink Selector (US)

Fig.6 schematizes the possible design of the US section of the CAMB switch. Each of the 4 uplinks consists of a control line and a data line. The control signal and the data stream are received by the selection block of the US portion. In our switch implementation, if more than one packet arrive at the same time, one is selected randomly. The function of the selection block is to provide a random selection of an active control line (no priority on one uplink over the others). This selection block consists of a priority register, multiplexers, a counter, a control logic (CL) and a selector.

The selected packet is then sent to a block of a serial-in parallel-out device, a counter, two flip-flops (FF1 and FF2) and two shift registers. The serial-in parallel-out device, the counter and one of the two flip-flops (FF1) are used to convert the destination address (in the packet header) from the serial format to the parallel format. The destination address is then sent (in parallel) to the Address Recognizer (AR) section of the switch. The destination address will be examined by the AR. The function of the shift registers is to store temporarily some number of packet bits while the destination address is being analyzed by the AR. The remaining flip-flop (FF2) stores the control bit (first bit in the packet header) during the transmission of the packet. This bit is passed to the AR.

The remaining part of the US section of the switch consists of a decoding logic and a multiplexer (MUX). They are used to detect the end of a packet transmission from the selected uplink. A change on control line from high to low indicates the end of packet transmission.

Address Recognizer (AR)

Fig.7 shows the schematic of the AR section of the switch. When the AR receives a packet (from

the US), it first checks the control bit of the packet stored in the FF2 of the US. If the control bit is 1, the packet has already been broadcast by its proper ancestor. In this case, the packet is sent to the parent switch only. Otherwise, namely, if the control bit is 0, the AR checks the destination address of the packet.

The packet destination address received from the US in a parallel format is compared against the switch address (determined by the Address Set) by the comparator. If the addresses do not match, the output of the comparator (denoted as REQ in Fig.7) is equal to 0. Otherwise, REQ is equal to 1.

Downlink Selector (DS)

The DS section is shown in Fig.8. The DS is responsible for broadcasting a packet to the switch children. The DS can be idle (no packet being broadcast), busy broadcasting a packet from the parent downlink, or busy broadcasting a packet from the AR.

The control logic receives a signal from the comparator (REQ) in the AR, from the flip-flop (FF2) in the US, and from the (downlink) control line of the parent switch. If the signal from the control line of the parent switch is 1, a packet is being transmitted from the parent switch. In this case, the multiplexer selects the packet from the parent switch and broadcasts it to all the children switches. Otherwise, there is no transmission from the parent switch. In this case, the multiplexer selects the packet from the AR and broadcasts it.

As described before, priority is given to the packet coming from the parent. Therefore, if the DS is broadcasting a packet from the AR when a packet from the parent downlink arrives, the transmission of the packet from the AR is aborted. When this abortion happens, the flip-flop (FF3) forces the control line go to low, indicating the end of the transmission of the packet from the AR. A station detects an abortion by checking the packet size indicated in the packet header and the actual size of the packet. After one clock cycle (allowing stations to distinguish the end of the aborted packet and the beginning of a new packet), the control line goes high again, indicating the beginning of the broadcast of the packet received from the parent.

4. Simulation Study of the Performance of the Broadcast Star Network

4.1. Simulation Model

In this section, we present performance study of the Broadcast Star network through simulations. In simulations, we consider both slotted and unslotted operations of the Broadcast Star. In both cases, a packet transmission time (the packet length divided by the channel speed) is used as a unit time. In the slotted case, the slot length is equal to a packet transmission time.

The performance measures of our interests include the average transmission delay, the variance of the transmission delay, and the loss probability of packets. The transmission delay D is the time from the beginning of the transmission of a packet to the time that the packet is successfully received by its destination. If there is a limit on the buffer capacity at a station, packets finding the buffer full upon arrival are lost. The loss probability is the probability of packets being discarded due to buffer overflow at a station.

The stations are assumed to be homogeneous. They are located at the same distance from the central switch, and the packet arrival statistics are same at all the stations. Let R denote the round trip propagation delay time on the channel to and from the switch. In case of the unslotted Broadcast Star, new packet arrivals form a Poisson distribution. In case of a slotted Broadcast Star, new packet arrivals form a geometric distribution. Namely, a packet arrives at a station with the probability q (per slot).

N is the total number of stations in the network. We consider both finite and infinite station population cases. In the infinite population case N is ∞ . Note that Nq is the total input rate to the Broadcast Star.

4.2. Simulation Results

We present simulation results. Figs.9 through 12 are for the Broadcast Star with infinite number of stations. Slotted operation of the network is assumed in Figs.9 and 10, and unslotted operation

is assumed in Figs.11 and 12. Figs.9 and 11 show the average transmission delay as a function of the throughput, and Fig.10 and 12 show the variance of the transmission delay for various values of R .

From Figs.9 and 11 it is observed that the average transmission delay increases as the total input rate increases. The average delay also increases as the value of R does. In the slotted system (Fig.9), the maximum throughput is 1.0. This is because, under conditions of heavy traffic, it is likely that all the channels are always busy, and thus there is a successful transmission in every slot. Therefore, the throughput of 1.0 is achieved. It is observed that the average delays increase rapidly when the throughput exceeds 0.8. In the unslotted system (Fig.11), it is shown that the maximum throughput is less than 1.0. The higher the value of R is, the smaller the achieved maximum throughput is. This is because when the throughput is higher, the number of conflicts is also higher, leading to a larger number of retransmissions.

The variance of the transmission delay exhibits a similar behavior in the slotted and unslotted Broadcast Star (Figs.10 and 12, respectively). The variance increases slowly until the throughput exceeds 0.8: beyond this point, the variance grows rapidly to infinity. The variance of the transmission delay is higher in the unslotted system, especially when R becomes larger. The same explanation given to the transmission delay behavior applies.

Figs.13 and 14 show the average transmission delay as a function of the total input rate (Nq) for the network with finite station population. Fig.13 is for the slotted Broadcast Star, and Fig.14 is for the unslotted Broadcast Star. The value of R in these figures is 5. In the slotted Broadcast Star, the average transmission delay is not very sensitive to the number of stations on the network. The maximum throughput achieved in the slotted Broadcast Star (approximately 0.9) is very close to the throughput of the network with infinite station population (Fig.9). However, in the unslotted Broadcast Star, the average transmission delay is very sensitive to the number of stations. When the number of stations is equal to 10, the maximum throughput is only 0.3, while the maximum throughput is 0.7 when the number of stations is equal to 50.

It is shown in Figs.9 through 14 show that the slotted Broadcast Star gives the smaller average

transmission delay than the unslotted one for all values of R . The difference between the slotted and unslotted becomes more significant as the value of R becomes larger. Also the maximum throughput achieved in the unslotted Broadcast Star is smaller than the slotted one. It is clear that the slotted Broadcast Star achieves a better performance than its unslotted counterpart. Therefore, we will focus on the slotted Broadcast Star in the following.

Figs.15 through 18 are for the (slotted) Broadcast Star with the finite number of stations. The size of the buffer at stations is assumed to be 1. In Figs.15 and 16, the round trip propagation delay R is assumed to be 0.05, and in Figs.17 and 18, R is assumed to be 5. Figs.15 and 17 show the loss probability due to buffer overflow, and Figs.16 and 18 shows the average transmission delay.

Figs.15 and 17 show that the loss probability increases slowly when the total input rate is not very large ($Nq < 1$). After that point ($Nq \geq 1$) the loss probability increases quickly, reaching 1. It is also observed that the loss probability increases as the propagation delay (R) increases.

Fig.16 and 18 show that the average transmission delay remains almost constant until the total input rate Nq becomes nearly equal to 1 ($Nq \approx 1$). Further, when Nq is small, the average delay is not very sensitive to the number of stations in the network. When Nq is large ($Nq > 1$), the average transmission delay starts to increase quickly and becomes much more sensitive to the number of stations in the network. This is explained by the fact that if more stations are contending for the switch, the chances of acquiring the switch is smaller, resulting in more blocked messages, and thus, more retransmissions. By comparing the case of $R = 0.05$ (Fig.16) and the case of $R = 5$ (Fig.18), it can be seen that when R (propagation delay) is longer, the average transmission delay is larger.

5. Concluding Remarks

In this paper we present a possible implementation of a switch for the CAMB tree network. In our design, we used TTL devices. We also investigated the performance of the Broadcast Star network through simulations. Simulation results show that the slotted Broadcast Star achieves better performance than the unslotted Broadcast Star, especially when the traffic load is heavy.

References

- [1] A. Albanese, "Star Network with Collision-Avoidance Circuits," *The Bell System Technical Journal*, Vol. 62, No.3, March 1983.
- [2] E. S. Lee and P. I. P. Boulton, "The Principles and Performance of Hubnet: A 50 Mbit/s Glass Fiber Local Area Network," *IEEE Journal on Selected Areas in Communications*, Vol. SAC-1, No.5, November 1983, pp.711-720.
- [3] T. Suda, Y. Yemini and M. Schwartz, "Tree Network with Collision Avoidance Switches," *Proc. of the IEEE Infocom*, 1984, pp.105-113.
- [4] S. Morris, T. Suda and T. Nguyen, "Tree LANs with Collision Avoidance: Photonic Switch Design and Simulated Performance," *Computer Networks and ISDN Systems*, 17, 1989.
- [5] E. S. Lee and P. I. P. Boulton, "Hubnet Performance Measurement," *IEEE Journal on Selected Areas in Communications*, Vol. SAC-1, No.5, November 1983, pp.711-720.
- [6] F. Closs and R. P. Lee, "A Multi-Star Broadcast Network for Local-Area Communications," in *Local Networks for Computer Communications*, A. West and P. Janson eds., North-Holland Publishing Company, cp IFIP, 1981.
- [7] A. E. Kamal, "A Performance Model for a Star Network," *Proc. of the IEEE Globecom*, 1986.
- [8] A. E. Kamal, "Star Local Area Networks: A Performance Study," *IEEE Trans. on Computers*, Vol. C-36, No. 4, April 1987.
- [9] A. E. Kamal and V. C. Hamacher, "Analysis of a Star Local Area Network with Collision Avoidance," in the *IEEE INFOCOM*, 1988.
- [10] H. S. Hassanein and A. E. Kamal, "Performance Evaluation of Prioritized Collision-Avoidance Star Local Area Networks," *IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*, June 1989.

- [11] T. Suda and K. Goto, "Performance Study of a Tree LAN with Collision Avoidance," Proc. of the IEEE Infocom, 1989.
- [12] T. Suda, S. Morris and T. Nguyen, "Tree LANs with Collision Avoidance: Protocol, Switch Architecture and Simulated Performance," Proc. of the ACM SIGCOMM Symposium, 1988.
- [13] T. Suda and S. Morris, "Tree LANs with Collision Avoidance: Station and Switch Protocol," Computer Networks and ISDN Systems, 17, 1989.
- [14] S. Marano and A. Volpentesta, "Performance Evaluation of Alberonet by Simulation and Theoretical Analyses," Proc. of the IEEE Infocom, 1987.
- [15] V. Ielapi, S. Marano and A. Volpentesta, "A Simulation Study for a Tree Local Area Network with Concurrent Transmissions," in Local Communications Systems: LAN and PBX, Elsevier Science Publishers BV, cp IFIP, 1987.
- [16] Y. Yemini, "Tinkernet: or, Is There Life Between LANs and PBXs?," Proc. of the IEE ICC, 1983.
- [17] F. Borgonovo and L. Fratta, "The S-ALOHA Throughput of a Tree-Topology Communication Network," Local Communication Systems: LAN and PBX, J.P. Cabanel, G. Pujolle and A. Danthine (editors), Elsevier Science Publishers B.V. (North-Holland), IFIP, 1987.




Fig.1 Broadcast Star Switch




Fig.2 Partitioning of CAMB Tree by Climbing Packet

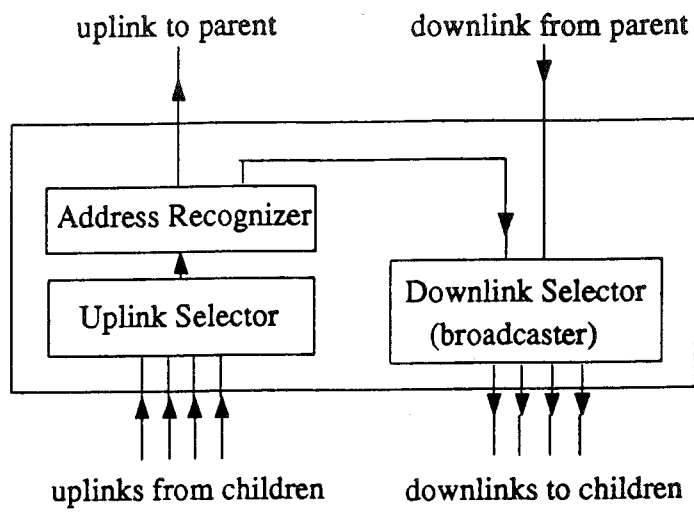


Fig.3 CAMB Switch

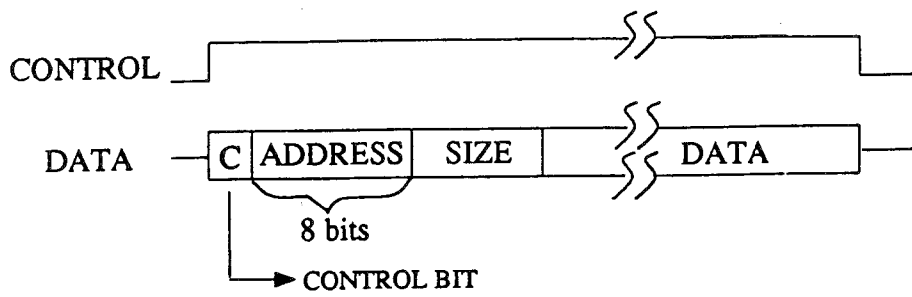


Fig.4 Packet Format

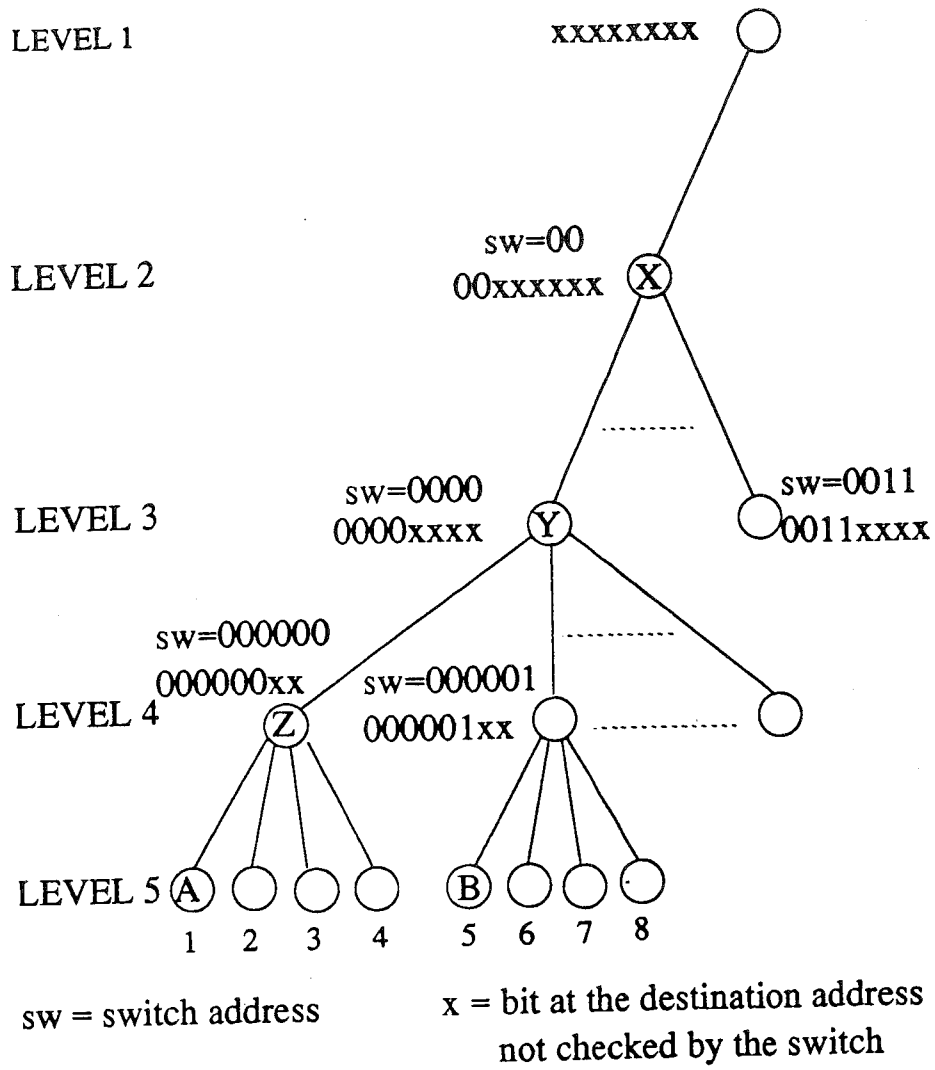
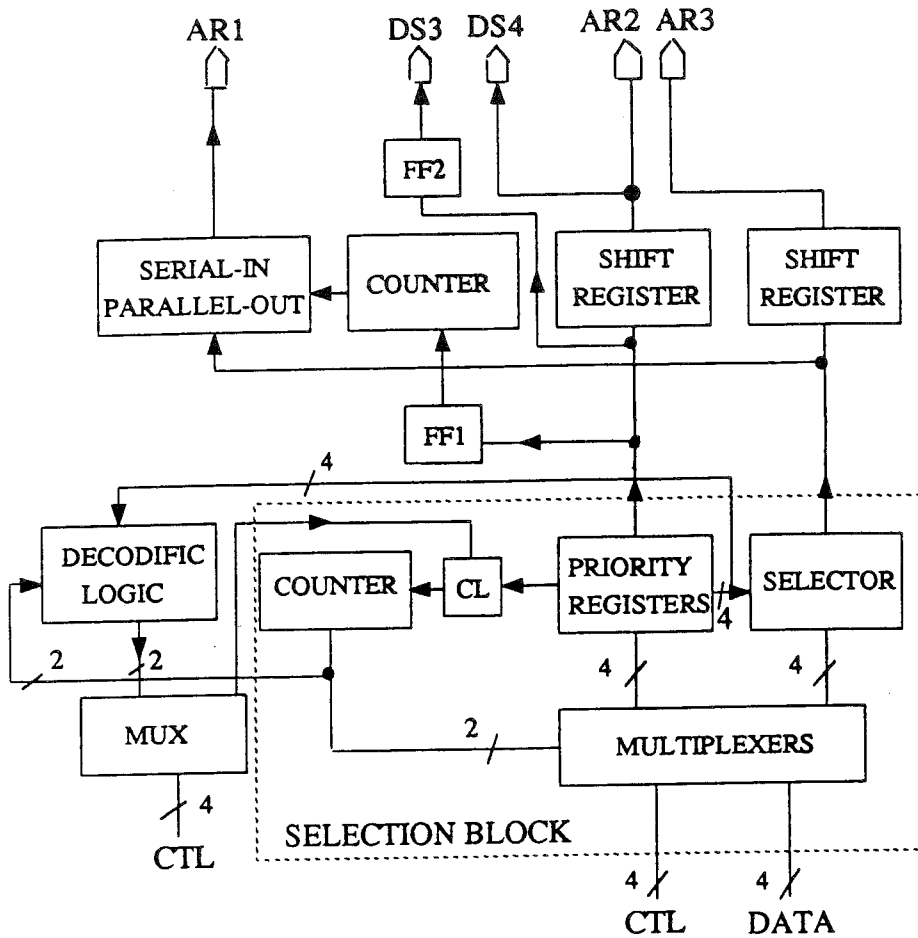


Fig.5 Addressing Scheme



US UPLINK SELECTOR

Fig.6 Uplink Selector Block Diagram

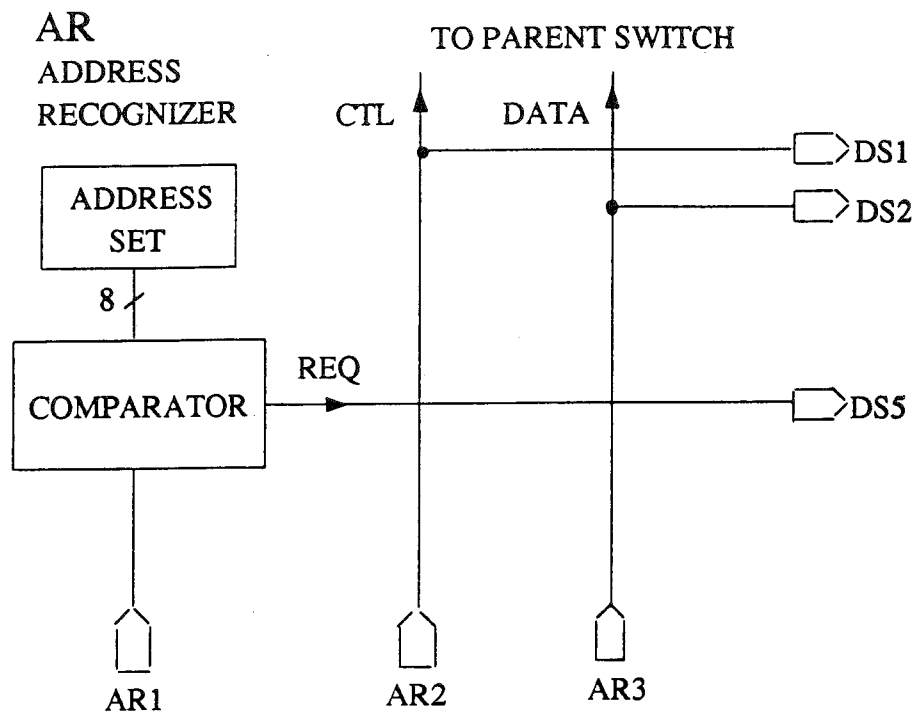


Fig.7 Address Recognizer Block Diagram

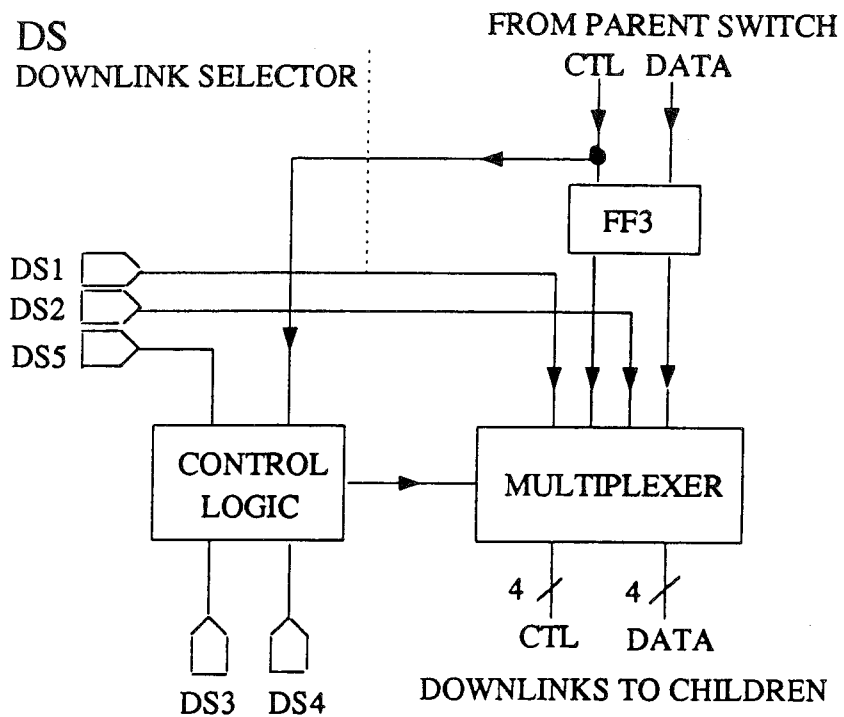


Fig.8 Downlink Selector Block Diagram




Fig.9 Transmission Delay in a Broadcast Star with Infinite Population (slotted)




Fig.10 Variance of Transmission Delay in a Broadcast Star with Infinite Population (slotted)




Fig.11 Transmission Delay in a Broadcast Star with Infinite Population (unslotted)




Fig.12 Variance of Transmission Delay in a Broadcast Star with Infinite Population (unslotted)




Fig.13 Transmission Delay in a Broadcast Star with Finite Population, $R=5$ (slotted, infinite buffer)




Fig.14 Transmission Delay in a Broadcast Star with Finite Population, $R=5$ (unslotted, infinite buffer)




Fig.15 Loss Probability in a Broadcast Star with Finite Population, $R=0.05$ (slotted, 1 buffer)




Fig.16 Transmission Delay in Broadcast Star with Finite Population, $R=0.05$ (slotted, 1 buffer)




Fig.17 Loss Probability in a Broadcast Star with Finite Population, $R=5$ (slotted, 1 buffer)

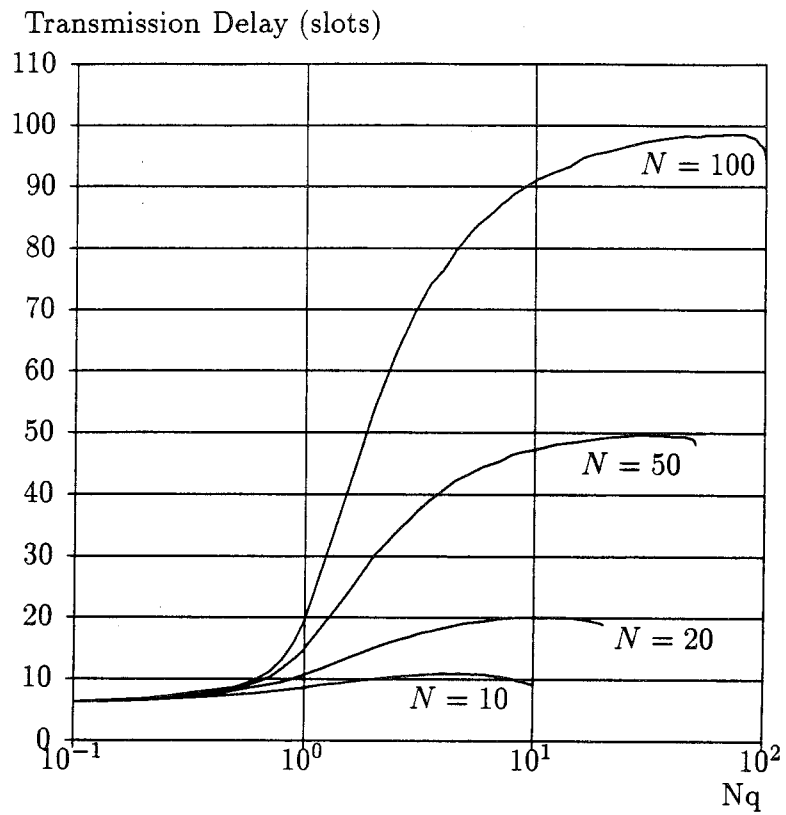


Fig.18 Transmission Delay in Broadcast Star with Finite Population, $R=5$ (slotted, 1 buffer)