





INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ³ : H04Q 9/00		A1	(11) International Publication Number: WO 83/ 03180 (43) International Publication Date: 15 September 1983 (15.09.83)
(21) International Application Number: PCT/US83/00282			Published <i>With international search report.</i>
(22) International Filing Date: 3 March 1983 (03.03.83)			
(31) Priority Application Number: 355,021			
(32) Priority Date: 5 March 1982 (05.03.82)			
(33) Priority Country: US			
(71) Applicant: BURROUGHS CORPORATION [US/US]; Burroughs Place, Detroit, MI 48232 (US).			
(72) Inventor: GROW, Robert, M. ; 69 Cedar Hill Road, Newtown, CT 06470 (US).			
(74) Agents: CHUNG, Edmund, M.; Burroughs Corporation, Burroughs Place, Detroit, MI 48232 (US) et al.			
(81) Designated State: JP.			

(54) Title: TIMED TOKEN RING WITH MULTIPLE PRIORITIES

(57) Abstract

A timed token protocol (Fig. 3) for use in a local area loop communications network. The protocol provides three classes of service with a priority (32) relationship between classes. The three classes allow guaranteed bandwidth, interactive and batch services. The classes are implemented by timing (24) the rotation time of a write token to measure instantaneous load. Transmission of information is limited by class of service and measured write token rotation time. Also disclosed is a hardware embodiment of a loop communications station (Fig. 4) which implements the timed token protocol.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	LI	Liechtenstein
AU	Australia	LK	Sri Lanka
BE	Belgium	LU	Luxembourg
BR	Brazil	MC	Monaco
CF	Central African Republic	MG	Madagascar
CG	Congo	MR	Mauritania
CH	Switzerland	MW	Malawi
CM	Cameroon	NL	Netherlands
DE	Germany, Federal Republic of	NO	Norway
DK	Denmark	RO	Romania
FI	Finland	SE	Sweden
FR	France	SN	Senegal
GA	Gabon	SU	Soviet Union
GB	United Kingdom	TD	Chad
HU	Hungary	TG	Togo
JP	Japan	US	United States of America
KP	Democratic People's Republic of Korea		

-1-

TIMED TOKEN RING WITH MULTIPLE PRIORITIES

Background of the Invention

The present invention relates in general to a system and method for communicating information between distributed stations. More particularly, the present invention relates to a timed token loop which provides three or more classes of service with a priority relationship between classes. The three classes of service provided allow guaranteed bandwidth, interactive and batch services. These classes are implemented by timing the rotation time of a write token to measure instantaneous load and limiting transmission of information by class of service and the observed write token rotation time.

The synthesis of the architecture for future office automation systems will be effected to a large extent by the available communication mechanisms supporting those systems. Circuit switched communication as provided by PBXs is efficient for handling voice traffic, but is limited by the maximum bandwidth available for data traffic. The inefficiencies inherent in using circuit switched service



-2-

for communication of bursty data traffic are being solved for low data rate devices through submultiplexing in the PBX. In contrast, local area networks have been optimized for efficient transmission of data at high burst rates with 5 little consideration for handling of digital voice.

Broadband systems provide the capability for carrying both voice and data over the same media through separate logical networks for voice and data on the same physical media. This hybrid approach solves some problems 10 by using a single media, but does not allow the flexibility given by integration of data on the same logical network.

Distributed communication systems based on token loop structures are well known in the prior art. Thus, in the article by David J. Farber titled "A Ring Network", 15 Datamation, February 1975, pp. 44-46, a collection of minicomputers connected by a ring-like digital communication system is disclosed. Similarly, the work by Newhall and Farmer on token controlled rings is exemplified in a paper titled "An Experimental Distributed Switching System To 20 Handle Bursty Computer Traffic, in Proc. ACM Symp. Problems in the Optimization of Data Communications Systems (Pine Mountain, GA, Oct. 1969), pp. 31-34.

The work of both Newhall and Farber was primarily applicable to data type traffic where all stations were 25 given an essentially equal opportunity to transmit. Thus the drawback of both systems is that they cannot support a guarantee of bandwidth.

The inability to guarantee bandwidth is also a drawback of the Ethernet system and all CSMACD-type 30 protocols for local area networks.

The Cambridge ring, which is based on the work of Pierce, is more adapted to a circuit type of service. In such a system, you essentially have a boxcar (which

-3-

corresponds to available bandwidth) into which data may be inputted or extracted. A fixed allocation procedure may be implemented which reserves the boxcar for a station to use. The drawback of the Cambridge ring is that there's 5 no simple or obvious way to use the boxcar when its not being used by the station that has a fixed reservation on it.

It is the general object of the present invention to overcome these and other drawbacks of the prior art by 10 providing a timed token protocol method and apparatus which integrates the favorable characteristics of both circuit switched and packet switched communications and provides these characteristics to a station through a single physical and logical interface.

15 It is another object of the present invention to provide a protocol adaptable to either a physical or logical loop which provides three or more classes of service with a priority relationship between classes, the three classes allowing guaranteed bandwidth, interactive 20 and batch services.

It is still another object of the present invention to provide a timed token protocol for a loop 25 communications network which allows the bandwidth guaranteed to a first station to be used by that first station when needed and further allows the same guaranteed bandwidth to be used by another station when it is not needed by the first station.

It is yet another object of the present invention to provide a token controlled loop which provides efficient 30 utilization of available bandwidth by the stations configured on the loop.

It is an additional object of the present invention to provide a timed token protocol adaptable to either a

-4-

physical or logical loop which provides for stations configured on the loop, each being of one or more of multiple classes of priority.

It is a further object of the present invention
5 to provide a timed token protocol for stations interconnected in a continuous loop structure, the protocol providing for stations with different priority levels, wherein stations of the highest priority level are capable of allocating a guaranteed minimum bandwidth to themselves to the exclusion
10 of other stations of the same priority.

It is still a further object of the present invention to provide a timed token protocol for stations interconnected in a continuous loop structure, the protocol providing for stations with different priority levels,
15 wherein stations not of the highest but of equal priority level are guaranteed a pool of bandwidth which may be fairly shared between them.

These and other objects, features and advantages of the present invention will become more apparent from
20 the detailed description of the preferred embodiment when read in conjunction with the drawings.

Summary of the Invention

According to the invention, a timed token protocol is provided to allow for the integration of three or more
25 classes of service on the same loop communications network. The highest class of service (Class 1) is for information which requires a guarantee of bandwidth and/or deterministic delay and jitter characteristics. This class of service is used for information that is represented in the time domain
30 (e.g., pulse code modulated voice), or where minimal queuing delay is important (e.g., process control). The

-5-

second class of information (Class 2) is that which is non-real time, but of an interactive nature. Traditional data communications between terminals and computers falls into this classification. That is, Class 2 information 5 requires some minimum throughput, but the absolute guarantee of bandwidth is not a requirement. The third class of service (Class 3) is for batch information where no minimum throughput is required and transmission of information may be delayed until network load is light.

10 The timing of a write token is the mechanism used for measuring instantaneous load. The measurement of load allows for the establishment of the three classes of service. Each class of service uses a different set of rules for determining when information may be transmitted.

15 An initial value called the target token rotation time (TTRT) must be selected for the network. The network protocols are designed so that under 100% offered load, the write token will revolve at the TTRT. The TTRT must be chosen to be less than or equal to the rate at 20 which devices operating at a Class 1 priority require service.

25 As in all token controlled communications systems, the right to service new information into the network is controlled by passing a privilege to transmit (the write token) from one station to another. The passing of the write token is controlled by the communication protocol.

30 Class 1 (C1) information may be transmitted upon every write token reception. The amount of C1 information transmitted with each write token reception is limited by a bandwidth allocation procedure to be explained below.



-6-

Other classes of information control their transmission by using a timer for each class. The timer for Class 2 is the target token rotation time minus the time for transmission of the maximum length frame (TTRT - MAXFRAME) (or the Class 5 2 target time). The target time for Class 3 would be some percentage of the target token rotation time related to load minus the maximum sized frame transmission time; (e.g., .6 x TTRT - Maxframe).

Stations of Class 2 or lower priority control 10 their transmission of information by timing between arrivals of the write token. Each time the write token arrives at a station, priority class timers are reset to their initial value if no information is to be transmitted, and the write token is passed on unimpeded. If, at arrival of the 15 write token, timers have not decremented to zero and information is queued for transmission, queued information may be transmitted for any priority class (other than Class 1) with a non-zero timer value.

A station which simultaneously services multiple 20 classes of information transmits lowest priority first and highest priority last. For example, in a station which has Class 1, Class 2 and Class 3 traffic, if when the write token arrives the Class 3 (C3) timer has not decremented to zero, the residual value in the C3 timer is loaded into a 25 token holding timer, the C3 timer is reset, and all timers continue to decrement. New frames may be sourced at C3 priority level onto the ring as long as the token holding timer has not decremented to zero. When the holding timer has decremented to zero, the current frame in transmission 30 is completed. At that time the residual in the Class 2 (C2) timer is loaded into the holding timer and the C2 timer is reinitialized. Transmission of C2 information then proceeds as described for C3. Class 1 (C1) information is then

-7-

transmitted. The write token is transmitted after all the allowed C1 frames. If at write token arrival the C3 timer had decremented to zero, the protocol begins with service to Class 2. If both the C3 and C2 timers had decremented 5 to zero, transmission begins with C1.

At design or configuration time, parameters must be chosen to allow for the quality of service desired for the classes of information. A value representing allocatable bandwidth is used for the control of Class 1 10 traffic. This value (ALLOC) is less than the target token rotation time. The amount of difference depends upon two factors. One is the amount of bandwidth wished for Class 2 devices to guarantee minimum throughput (C2POOL), and the other is the latency (LATEN) of the physical or 15 logical ring. Latency is the amount of time it takes for the token to go around the ring with zero load. The sum of allocatable bandwidth, the Class 2 pool, and latency is equal to the target token rotation time (TTRT = ALLOC + C2POOL + LATEN). Normally the timer value for class 3 20 would be less than the available bandwidth for Class 1 (C3TIMER < ALLOC). In such case, no minimum bandwidth is guaranteed for Class 3. If the C3 timer is greater than allocatable bandwidth, some minimum bandwidth has been 25 guaranteed for Class 2 or Class 3.

By limiting traffic through the described mechanisms, Class 1 service is that of a minimum guaranteed bandwidth, since the guarantee is for each Class 1 station that has received a portion of the allocatable bandwidth to be able to transmit some fixed amount of information every 30 rotation of the write token. If the write token is rotating faster than the target token rotation time, the bandwidth usable by Class 1 would therefore be greater than the guaranteed minimum.



-8-

Class 2 devices have a pool of bandwidth guaranteed to them which is shared fairly by all devices designed to use that quality of service. In addition, the timing of the write token allows any unused Class 1 bandwidth to be used by lower priority stations. This includes both allocated and unused bandwidth as well as unallocated Class 1 bandwidth. Class 3 priority has no guaranteed throughput, but would only transmit when the load falls below the arbitrary percentage established for that class.

In the preferred embodiment, the mechanism for allocation of bandwidth for Class 1 service is to use a special token which is transmitted around the loop. When a Class 1 station desires to allocate additional bandwidth to itself, the write token is captured and a bandwidth allocation token is transmitted. Each station on the loop takes this bandwidth allocation token, adds to it the amount of bandwidth which is currently allocated by that station for Class 1, and forwards the allocation token to the next station. When the bandwidth allocation token returns to the station attempting to allocate bandwidth, the write token is regenerated and passed on to the next station. If the current allocation of bandwidth returned in the bandwidth allocation token plus the desired allocation is less than the allocatable bandwidth, then the allocation is granted. The bandwidth desired is then added to the total allocated for that station.

In addition to the present invention providing a protocol method for the integration of different classes of service on the same communication network, a preferred hardware embodiment of the station logic is also described. Details of the hardware implementation is contained in the detailed description of the preferred embodiment.



-9-

Brief Description of the Drawings

Figure 1 illustrates the interconnection of a plurality of stations in a communications loop as employed in the present invention.

5 Figure 2 illustrates a plurality of Class 1 stations interconnected in a loop structure, each station having a current Class 1 allocation of the number of units indicated.

10 Figure 3 shows the formats of the data frame, recovery token, bandwidth allocation token and write token as used in the preferred embodiment of the present invention.

15 Figure 4A illustrates a single loop with redundant transmission paths, having been restructured into two separate fragmentary loops after the occurrence of the two indicated faults and the invocation of loopback. Each station shown is Class 1 priority and has a current Class 1 allocation as indicated.

20 Figure 4B is similar to Figure 4A, but illustrates the configuration after the failure between stations C and D has been repaired, wherein the two separate loops of Figure 4A have begun to function as a single loop.

Figure 5 shows a block diagram of the loop logic included in each station attached to the loop.

25 Detailed Description of the Preferred Embodiment

The protocol of the present invention is designed to operate in a token controlled loop (or ring) communications network. Those skilled in the art will appreciate that the protocol of the present invention may 30 be implemented either on a physical loop or on a logical loop (or ring) superimposed on a physical bus. Such a loop communications network may utilize a single loop (or ring) architecture or one incorporating a second redundant loop

-10-

with network components that route the data around any network faults. The following description will consider the implementation of the present invention in a single loop architecture. However, those skilled in the art will appreciate that the present invention may be readily 5 adapted for use in loop architecture employing a second redundant loop.

The timed token protocol of the present invention allows for the integration of different classes of service on the same communications network. The preferred 10 embodiment of the protocol supports three basic priority classes of service. The first class (Class 1) is for that type of information which requires a guarantee of bandwidth and deterministic delay and jitter characteristics. This class of service is used for information which is 15 represented in the time domain. Thus, for example, Class 1 information may include real-time audio applications such as pulse code modulated (PCM) voice. Such Class 1 information is sometimes referred to as synchronous information.

20 The second priority class of information (Class 2) supports non-real time applications requiring interactive response. Typical of such an application is traditional data communications between terminals and a computer. Thus, in this class some minimum throughput is required, but an 25 absolute guarantee of bandwidth is not a requirement. This class of traffic is sometimes referred to as asynchronous.

The third class of service (Class 3) is that of batch traffic which requires no minimum throughput, but can be transported through the network when load is light. 30 Background communications such as electronic mail or file transfers are typical applications where Class 3 service would be appropriate. The protocol of the present invention



-11-

allows for the creation of additional priorities within this class of service for batch traffic. Thus, Class 4 Class N, all subclasses of Class 3, may be established.

5 As will be discussed below, the timing of a write token is the mechanism used for measuring instantaneous load on the loop. This measurement of load allows for the establishment of the classes of service discussed above. Each class of service uses a different set of rules when 10 information may be transmitted.

The speed at which the write token rotates around the loop determines the rate at which the network services each station, the network's queuing delay, and the bandwidth lost to passing the write token from station to station. An 15 initial value called the target token rotation time (TRTT) must be selected for the network. The loop protocols guarantee that the short-term average rotation time will vary from the minimum orbit time with no load, to the target token rotation time (TRTT) under full load conditions. The 20 TRTT is important since it defines the maximum average write token rotation time. Thus, under 100% offered load, the write token will revolve at the TRTT. In the preferred embodiment of the present invention, the TRTT is loaded into each station during network initialization.

25 The TRTT is related to the rate at which devices operating at a Class 1 priority require service. For Class 1 priority, information may be transmitted once every write token arrival. Thus, the protocol guarantees that Class 1 (C1) information may be sent at least once every TRTT for each 30 Class 1 station on the loop that has received a guarantee of bandwidth. This means that the TRTT determines the guaranteed service rate for C1 transmission.



-12-

As previously mentioned, the TRTT is related to the rate at which devices operating at a Class 1 priority require service. For example, if service is only required every eight milliseconds, a 64K bits per second information rate would be transmitted using a packet size of 64 bytes (or 512 bits). Thus, the time to accumulate 512 bits at a rate of 64K bits/second is 8 milliseconds. If the information generation rate was 32K bits per second instead of 64K bits per second (bps), a packet size of 32 bytes (or 256 bits) would be used with this eight millisecond TRTT. Similarly, a chosen service time of four milliseconds would require 32 bytes (or 256 bits) per packet to support a 64K bps information rate, or 16 bytes (or 128 bits) per packet for a 32K bps information rate.

The value used for the target token rotation time must be less than or equal to the selected service time in order to guarantee that each station is serviced within the selected service time. If error free bandwidth is desired for Class 1, the error rate of the media must be taken into consideration. In such case, the target token rotation time would be correspondingly less than the rate at which packets are produced.

As in all prior art token controlled communication systems, the right to source new information into the network of the present invention is controlled by passing a privilege to transmit (or write token) from one station to another. In the present invention, the passing of the write token is controlled by the communication protocol.

FIG. 1 shows a single loop communications system as employed in the present invention. In such a system, station A 10 passes a write token to station B 12,



-13-

station B 12 passes the write token to station C 14 and station Z 16 passes the write token to station A 10, and the cycle is thereafter repeated.

5 Each station on the loop may interconnect devices for handling one or more priority classes of data. Thus, a single station may handle classes 1 or 2 or 3, or any combination of the three classes.

10 Class 1 information may be transmitted upon every write token reception by a Class 1 station. The amount of C1 information that a Class 1 station may transmit upon each write token reception is limited by a bandwidth allocation procedure to be explained below. Other classes of information transmission are controlled by using a timer for each class, each station capable of transmitting 15 non-Class 1 information must have a timer for each non-class 1 priority of information it is capable of transmitting.

20 The initial value (or target time) for each Class 2 timer is set to the target token rotation time (TTRT) minus the time for transmission of the maximum length frame (TRTT-MAXFRAME). As will be obvious to those skilled in the art, MAXFRAME is dependent on the quality of crystals used in the network, the amount of padding inserted between frames, the coding and synchronization system utilized, and other characteristics of the network. If the 25 network uses a central clock, the maximum length frame could be arbitrarily long. In the preferred embodiment of the invention, the maximum length frame has been chosen to be 1024 bytes in the information field.

30 In the preferred embodiment, the timer value for Class 2 is chosen to assure efficient network utilization and



-14-

prevent starvation of stations. Thus, if the timer value for Class 2 is not TTRT-MAXFRAME, the network will give small frames a greater probability of transmission. This would tend to encourage inefficient network utilization 5 since the overhead to information ratio is higher for smaller frames. It also might prevent some Class 2 stations from transmitting at all.

The target time for every Class 3 timer is set to some percentage of the target token rotation time 10 minus the maximum sized frame transmission time. For example, the target time for Class 3 might be $0.6 \times \text{TRTT} - \text{MAXFRAME}$. The percentage of the TRTT used in the calculation of the Class 3 target time is related to the load on the network. Thus, each timer within the same 15 Class is set to the same initial value, the value being set when the network is configured.

The timer value for any other priorities of batch information would follow the same pattern as used for Class 3, with each priority using a smaller percentage 20 of TTRT for its initial timer value. For example, if a Class 4 priority was included, the initial Class 4 timer value (or target time) might be $0.4 \times \text{TRTT} - \text{MAXFRAME}$.

Note that in the present invention all timers in all stations are continually decrementing until they have 25 decremented to zero (at which time a timer ceases to decrement).

Stations of Class 2 or lower priority control 30 their transmission of information by timing the write token from arrival to arrival. Each time the write token arrives at a station, priority class timers are reset to their initial values. If no information is to be transmitted, the write token is passed back onto the loop unimpeded. If, at arrival of the write token, the timers have not decremented



-15-

to zero and information is queued for transmission, queued information may be transmitted for any priority class (other than Class 1) whose corresponding timer has a non-zero value. Thus, a station can transmit Class 5 2 (or Class 3) information as long as the time since the previous write token arrival is less than the Class 2 (or Class 3) target time. If the write token arrives later than a Class 2 (or Class 3) target time since its previous arrival, the station must refrain from transmitting any 10 Class 2 (or Class 3) information onto the loop.

A station which simultaneously services multiple classes of information transmits lowest priority first and highest priority last. For example, in a station which has Class 1, Class 2 and Class 3 traffic, when the write 15 token arrives, if the timer for Class 3 has not decremented to zero, the residual value in the Class 3 timer is loaded into a token holding timer, the Class 3 timer is simultaneously reset to its initial value, and the Class 3 timer and token holding timer immediately begin to decrement. Note that during this time the Class 2 timer and 20 all timers in other stations are continuously decrementing. Multiple new frames may be sourced at Class 3 priority level onto the loop as long as the token holding timer has not decremented to zero.

25 When the token holding timer has decremented to zero, the current Class 3 frame in transmission is completed. At that time, the residual value in the Class 2 timer is loaded into the token holding timer, the Class 2 timer is simultaneously reset to its initial value, and the 30 Class 2 timer and token holding timer immediately begin to decrement. Note that during this time the Class 3 timer and all timers in other stations are continually decrementing. Transmissions of Class 2 information then proceeds in the

-16-

same manner as described for Class 3, the transmission of Class 2 information ceasing after completing transmission of the frame being transmitted when the token holding timer goes to zero.

5 Note that if at the time the write token arrives there is no Class 2 information to transmit, the Class 2 timer still would not be reset until after Class 3 is serviced.

10 After Class 2 transmission is ended, the station next transmits Class 1 information. Details of when Class 1 information may be transmitted will be discussed below. After all of the allowed Class 1 frames have been transmitted, the write token is retransmitted back onto the loop.

15 Note that in the above description, if at the arrival of the write token the C3 timer had already decremented to zero or there was no Class 3 information to transmit, the station would immediately reset the C3 timer to its initial value and begin with service to Class 2. If 20 when the write token arrived both the Class 3 and Class 2 timers had decremented to zero or there was no Class 2 or Class 3 information to transmit, the station would immediately reset the C2 and C3 timers and transmission would begin with Class 1 information.

25 To improve the fairness within a class of service, a method of accumulating lateness is preferred. Such a method may be optionally implemented in the present invention as follows. Thus, after transmission of Class 2 (or Class 3) information, if on the next rotation of the 30 write token the timer for a class that transmitted decrements

SUBSTITUTE SHEET



-17-

to zero, it will cause a class lateness register (not shown) to increment until write token arrival. This lateness must be removed before transmission can occur for that class again. The measure of token earliness (the residual of the target timer normally transferred 5 to the token holding timer) is subtracted from the lateness register. This continues until the lateness has been compensated for by earliness. Then transmission can resume either on the same rotation that lateness is 10 completely removed or on the next rotation after lateness is completely removed.

When the network is configured, parameters must be chosen to allow for the quality of service desired for the classes of information. As previously mentioned, the 15 amount of Class 1 information that may be transmitted upon each write token reception is limited by a value representing allocatable bandwidth. This value, ALLOC, is less than the target token rotation time. The amount that ALLOC differs from the target token rotation time is dependent on three 20 factors. These three factors are the amount of bandwidth to be allocated for Class 2 devices to guarantee minimum throughput (C2POOL), the latency (LATEN) of the physical or logical ring, and any bandwidth for system administration (SYSADM). System administration would include functions 25 required to construct a logical loop or for bandwidth allocation procedures. Latency (LATEN) is the amount of time it takes for the token to go around the loop with zero load. The sum of allocatable bandwidth (ALLOC), the Class 2 pool (C2POOL), latency (LATEN) and system administration 30 (SYSADM) is equal to the target token rotation time. Thus,

$$TTRT = ALLOC + C2POOL + LATEN + SYSADM$$

In practice, latency and system administration will be a small factor for most systems. Thus, allocatable bandwidth



-18-

(ALLOC) is selected primarily on how much Class 2 traffic is anticipated on the loop. By selecting the size of C2POOL, you determine the minimum throughput per class 2 station. The remainder (ALLOC) is the bandwidth which 5 is available for all Class 1 (C1) traffic.

Normally, the timer value for Class 3 and any lower priority classes will be less than the available bandwidth for Class 1 (C3TIMER < ALLOC). In such case, no 10 minimum bandwidth is guaranteed for Class 3. If this is not the case, then the actual Class 2 pool (C2POOL) is

$$\begin{aligned} \text{C2POOL} &= \text{TTRT} - \text{LATEN} - (\text{C3TIMER} - \text{ALLOC}) \\ &= \text{TTRT} - \text{LATEN} - \text{C3TIMER} + \text{ALLOC} \end{aligned}$$

Thus, if the Class 3 timer value is greater than the allocatable bandwidth, some minimum bandwidth has been 15 guaranteed for Class 2 or Class 3.

Class 1 service is that of a minimum guaranteed bandwidth. Each Class 1 station that has received a portion of the allocatable bandwidth (ALLOC) guarantees it 20 will be able to transmit some fixed amount of information corresponding to the size of its allocation for every rotation of the write token. Thus, if the write token is rotating faster than the target token rotation time, the bandwidth usable by Class 1 stations would be greater than the guaranteed minimum bandwidth available for Class 1 25 information.

The pool of bandwidth guaranteed to Class 2 stations (C2POOL) is shared by all stations designed to use Class 2 service. In addition, the timing of the write token allows any unused Class 1 bandwidth to be used by lower 30 priority stations. This is accomplished by lower priority stations timing the write token and measuring load as previously described. Thus, when the write token arrives



-19-

at a lower priority station, the token holding timer basically says how much bandwidth was not used on the last rotation. Note that this has no bearing on what is allocated or what was even allocatable; it is just 5 what was used, and so by measuring the current load (via the C2 and C3 timers) lower priority stations are able to use that unused Class 1 bandwidth which is either not allocated to Class 1 or allocated and not used. In other words, the transmission algorithm used by Class 2 and 3 10 stations does not even look at ALLOC; only the allocation algorithm for Class 1 looks at ALLOC.

Class 3 and lower priority stations have no guaranteed pool of bandwidth available to them and hence have no guaranteed throughput. Thus, they may only transmit 15 when the load falls below the arbitrary percentage established for that class. As previously discussed, Class 3 and lower priority stations make this determination based on the target time (or timer value) established for that class. For example, if the Class 2 target time is 3 milliseconds 20 and the Class 3 target time equated to 6 milliseconds, then whenever Class 2 load fell below 75 percent Class 3 would be able to transmit.

There are situations where a Class 1 station may want to increase the amount of Class 1 bandwidth allocated 25 to it. For example, if a device such as a PBX is attached to a station, each time you establish a call or terminate a call the station would want to change the amount of allocated Class 1 bandwidth for that station to reflect the current load.

30 Those skilled in the art will appreciate that the allocation of bandwidth for Class 1 service can be



-20-

accomplished in one of at least five ways:

- through configuration limitation.
- Through a central device which controls the allocation of bandwidth.

5 - By including in the write token a field representing the current allocated bandwidth.

 - By a special token which circulates around the loop and collects the currently allocated bandwidth from all stations.

10 - By a distributed message passing algorithm for determining the current allocation of bandwidth.

In the preferred embodiment of the present invention, the mechanism used for allocation of bandwidth for Class 1 service is to use a special token which is transmitted 15 around the loop.

When a station desires to allocate additional Class 1 bandwidth, it captures the write token and transmits a bandwidth allocation token (FIG. 3). This bandwidth allocation token includes a field for 20 accumulating the allocated bandwidth of all stations (ALLOCATED). This field is initialized to the current allocation of the station transmitting the bandwidth allocation token. Each Class 1 station on the loop takes the bandwidth allocation token, adds to the allocated field 25 (ALLOCATED) the amount of bandwidth which is currently allocated by that station to Class 1, and forwards the bandwidth allocation token to the next station. When the bandwidth allocation token returns to the station attempting to allocate bandwidth, the write token is 30 regenerated and passed on to the next station. If the current allocation of bandwidth returned in the bandwidth allocation token (ALLOCATED) field plus the desired additional allocation is less than or equal to the allocatable bandwidth (ALLOC), then the allocation is



-21-

granted. In such case the additional bandwidth desired is then added to the total bandwidth currently allocated for that station.

As an example of the above operation, assume we have a loop as shown in FIG. 2 with all stations 5 A, B, ..., E being of Class 1, each station having the number of Class 1 units of bandwidth allocated to it as shown in the figure. Further assume that the loop has a capacity of 12 allocatable units of Class 1 bandwidth (= ALLOC). If station C wants to allocate additional 10 bandwidth to itself, it captures the write token, and sends out a bandwidth allocation token with 0 in the allocation field ALLOCATED (since C currently has 0 units allocated to it). When the allocation token is received by station D, 5 units are added to the ALLOCATED field 15 and the allocation token is transmitted to station E. This process continues until the bandwidth allocation token arrives back at station C with a value of 9 in the ALLOCATED field. Station C compares the value in the allocation field ALLOCATED against the loop capacity of 20 12 (as specified by ALLOC). Since there is a difference of 3 units between the value of the ALLOCATED field and ALLOC, station C may allocate up to three additional units of Class 1 bandwidth. If C was trying to allocate only one additional unit, it would change its current allocation 25 register from 0 to 1. Subsequently, if station C received a bandwidth allocation token, it would add 1 unit to the ALLOCATED field; viz., its current allocation. Note that when station C is finished using its 1 unit of allocated 30 Class 1 bandwidth, it will change its current allocation register back to 0 units. Similarly, when other Class 1



-22-

stations no longer require part or all of their currently allocated bandwidth, they will also reduce the value in their current allocation register accordingly.

Under one particular failure condition it is
5 possible for a Class 1 station to receive a bandwidth allocation token wherein the value in the ALLOCATED field exceeds the total allocatable Class 1 bandwidth (ALLOC). Thus, for example, a loop having redundant transmission
10 paths and using loopback for reliability will fragment into separate loops when multiple failures occur (FIG. 4A). A description of the loopback mechanism is contained in U. S. Patent 4,190,821, issued to Thomas R. Woodward on Feb. 26, 1980, which is incorporated by reference. Upon repair of one of the failures, the two separate loops will begin
15 functioning as a single loop (FIG. 4B). Before repair, the two loop fragments operate using the same bandwidth allocation procedure. Hence both loop fragments in FIG. 4A have less than the ALLOC amount (= 12 units) of bandwidth allocated. But when repair of the fault between stations C
20 and D rejoins the loops (FIG. 4B), the current allocation is larger than ALLOC.

As an example, assume in FIG. 4B that station C wanted to allocate additional bandwidth. When the allocation token arrives at station A, the ALLOCATED field would specify
25 11 units. In such case, station A would detect that the currently allocated Class 1 bandwidth exceeded ALLOC, and would respond by terminating all sessions since it could only have a current allocation of 1 unit. Since station A terminated all sessions, it would change its current
30 allocation to 0 and when the allocation token arrived at station B the ALLOCATED field would still specify 11 units. Station B would not be required to terminate its sessions since it could still maintain its current allocation of 1

-23-

unit without the total current allocation exceeding ALLOC (= 12 Units). However, the request for additional allocation made by station C could not be granted since if granted the total current allocation would exceed

5 ALLOC.

Alternatively, station A could deallocate 1 unit of bandwidth and maintain a portion of its sessions. In such case, station A would add 1 unit to the ALLOCATED field of the allocation token. Thus, when the allocation token reached station B, the ALLOCATED field would specify 12 units. Station B would detect that that currently allocated Class 1 bandwidth exceeded ALLOC (if station B maintained its current allocation), and respond by terminating all of its sessions since it could no longer have a current allocation of 1 unit. Thus when the allocation token arrived at station C, the ALLOCATED field would specify 12 units, and the request made by station C could not be granted since ALLOCATED was already equal to ALLOC.

20

In the preferred embodiment of the present invention, under normal conditions only one write token circulates around the loop. Station failures or transmission errors may cause the loss of the write token. The determination of when the write token has been lost is quite simple in the present timed token loop because of the target token rotation time. Thus, the absolute worst case time for rotation of the write token is less than or equal to two times the target token rotation time (TRTT). This worst case time occurs under the extremely unlikely conditions that all of the bandwidth is currently allocated to Class 1 (no C2POOL) and all of the bandwidth is being used by Class 1 stations on one rotation and on the next rotation all stations cease to transmit.

SUBSTITUTE SHEET



-24-

Therefore, if the time since the previous arrival of a write token at a station (of any class or classes) is greater than two times the target token rotation time, the station knows that the write token has been lost. In 5 response to detecting the loss of the write token, the station will initiate a bidding cycle to recover the token. The bidding cycle is initiated by the station transmitting a recovery token (FIG. 3) with its individual address in the destination address (DA) field. The next station 10 receiving the recovery token passes it on if the destination address in the recovery token is greater than its own individual address, replaces the destination address with its own individual address if the destination address is smaller than its own individual address, or wins 15 the bid and regenerates the write token when the destination address in the recovery token is equal to its own individual address. Thus, the station with the highest individual address wins the bid and regenerates the write token.

Those skilled in the art will appreciate that the 20 bidding algorithm may easily be modified so that the station having the lowest individual address will win the bid and regenerate the write token.

It should be noted that the same bidding process used to regenerate a lost write token may also be utilized 25 to generate a write token when the network is initialized.

The discussion will next briefly consider the formats of the tokens and data frame (FIG. 3) and the station's response upon the receipt of incoming information.

All stations on the loop look at all incoming 30 information. If the destination addressed (DA) specified in a data frame equals the receiving station's individual



-25-

address, the data frame will be buffered if there is no room in the station's buffer. The network allows for creation of status indications (in the EFD field) for error control and flow control on data frames. If 5 a received data frame has a bad cyclic redundancy check (in the frame check sequence field), a status bit in the EFD field will be set to NAK. If the station has no room to buffer the frame, a status bit in the EFD field will be set to indicate the buffer full condition. On the 10 other hand, if the frame is buffered a status bit in the EFD field will be set to ACK. If the destination address of a frame is not that of the station, the frame is repeated to the next station without change, except in the case described below.

15 Each station also looks at the source address of every received data frame. If the source address is equal to the station's individual address, the frame has made one complete revolution of the loop. The status indicated in the EFD field is captured and used to control 20 retransmission. The station invalidates the frame by either placing an abort sequence in the frame or by removing the frame from the loop. This prohibits multiple receptions of the same frame by its continuing to rotate around the loop under very light load.

25 The following description describes the characteristics of the hardware used to implement a station. Although a detailed description of the hardware is not included, the description is sufficient to enable those skilled in the art to make and use the invention.

30 Referring to FIG. 3, information is carried on the loop (or ring) in frames which include unique identifiers for start of frame and end of frame. The start frame delimiter (SFD) indicates whether the frame is one of

-26-

the tokens used for management of the loop or is a data frame. The end frame delimiter (EFD) includes a status field which is used to indicate positive or negative acknowledgment of the frame, buffer status and other status indications. The data frame has two addresses, the destination address (DA) and the source address (SA). User data and any user control information is carried in the information field (INFO) of the data frame. The frame check sequence field (FCS) is used to detect errors in transmission.

10 The various fields in the token formats have functions similar to those described for correspondingly named fields in the data frame format. The exact positioning of the fields in the token and data formats and the meaning assigned to the various status bits is 15 not critical to the operation of the present invention. Thus, those skilled in the art will appreciate that modifications to the formats shown in FIG. 3 may be made without departing from the spirit of the present invention.

20 In the preferred embodiment of the present invention, information transmitted on the loop is encoded in a dual frequency format which produces three values, "1", "0", and "violation". The violation value is used to define SFD and EFD. However, those skilled in the art will 25 appreciate that other data transmission formats may be more appropriate to a particular application.

FIG. 5 shows a block diagram of the ring logic provided for each station on the loop. In FIG. 5, solid lines indicate actual flow and broken lines indicate control 30 information. The broken line entering the bottom of the

-27-

transmit priority logic 32 is the control interface from the data device(s) (not shown) supported by and included in each station. The data lines leaving the bottom of the receive buffer logic 42 and entering the 5 bottom of the transmit buffer logic 34 are data interfaces which allow the data device(s) (i.e., facsimile units, storage devices, etc.) to receive/transmit information from/to the ring logic. As will be obvious to those skilled in the art, these date interface lines 10 are normally coupled to some logic on the data bus of the data device(s). Signals received from the loop are fed into a decoder 20 where a clock is extracted and the data is clocked into the first in/first out queuing unit (FIFO) 22. The choice of whether a centralized or decentralized 15 clock is utilized is not important. In the preferred embodiment of the invention a decentralized clock is utilized; viz., each station runs off of an independent clock. In such a situation, the difference in rates between the received clock and a station's internal clock 20 can cause the need to insert or delete bits between frames. This is accomplished by the FIFO 22. Thus, the FIFO 22 will add or delete bits between frames to compensate for rate 25 disparities between crystals in different stations. The design of the FIFO 22 is dependent on the clocking system being used, the quality of crystals and other parameters. The design of a FIFO 22 to accommodate the particular network's characteristics will be obvious to those skilled in the art.

It should be noted that the FIFO's 22 length is 30 an important factor in determining the maximum frame length MAXFRAME. Thus, MAXFRAME is dependent on how many bits of

-28-

padding the frame creation logic 36 inserts between frames, the length of the FIFO 22 and other factors such as how many stations are in the loop and the method of clocking used. The tradeoff between these variables will be apparent to those skilled in the art.

5 The data output from the FIFO 22 sources four parallel sections of logic: the token management logic 24; the frame synchronization logic 26; the address recognition logic 28; and the error detection logic 30. The frame 10 synchronization logic 26 scans the data coming in from FIFO 22 looking for start frame delimiters (SFDs). Upon detection of a start frame delimiter, the frame synchronization logic determines whether the frame is a data frame or a token based on the coding of the start 15 frame delimiter. The design of the frame synchronization logic is dependent on the coding scheme used to indicate start frame delimiters (SFDs) and end frame delimiters (EFDs). Thus for example, if a HDLC/BDLC type of zero 20 insertion was used to create the unique patterns identifying SFD and EFD, the frame synchronization logic 26 would have to find those flag patterns and also remove the inserted zeros in non-flag fields. In the preferred embodiment, 25 illegal coding sequences are used to indicate delimiters and identify whether a data frame or one of the three types of tokens has been received.

If the frame synchronization logic 26 detects a token, the token management logic 24 is notified by the frame synchronization logic 26. The token can be one of three types: a write token, a recovery token, or a bandwidth 30 allocation token. In response to the frame synchronization logic 26 notifying the token management logic 24 that a write token has been received, the token management logic 24 will reset an internal timer (not shown) which is doing

-29-

timeouts for recovery of the write token. That is, the internal timer will be reset to its initial value of two times the TT RT. Note that this internal timer is also set to its initial value when a station joins the loop.

5 If a write token is received, the transmit priority logic 32 is notified by the token management logic 24. The transmit priority logic 32 includes the C2 timer, C3 timer, token holding timer, and lateness register, none of which are shown. Upon the transmit priority logic 32 being notified of the receipt of a write token, the transmit priority logic 32 will load the token holding timer and reset the C2 and C3 timers in accordance with the procedure specified earlier. If information has been queued for transmission in the transmit buffer logic 34 and the timers for the classes queued have not decremented to zero, information will be taken from the buffers (not shown) in the transmit buffer logic 34 and run through the frame creation logic 36 where a start frame delimiter, the source address, the frame check sequence and end frame delimiter are appended.

20 The design of the frame creation logic 36 is dependent on the data structure used in the transmit buffer logic 34, the ordering of fields of the data frame, and the method used for creating frame delimiters. In the preferred embodiment, the only information stored in the transmit buffer logic 34 is the information field and the destination address. Thus, the frame creation logic 36 puts on a start frame delimiter (SFD), adds the source address (SA), generates the frame check sequence (FCS) and adds the end frame 25 delimiter (EFD). With the functions performed by the frame creation logic 36 now understood, suffice it to say that the 30

SUBSTITUTE SHEET



-30-

design of the frame creation logic 36 will be obvious to those skilled in the art.

The frame constructed in the frame creation logic 36 is transferred to the output control logic 38 which in turn supplies the data frame to the encoder 40. The output control logic 38 is a multiplexor, responsive to inputted control signals, for gating to the encoder 40 one of the output control logic's 38 data inputs in accordance with the received control signals. Its design is well known to those skilled in the art. The encoder 40 encodes the data frame into whatever data transmission format is being utilized (i.e., NAZ, dual frequency) and transmits the encoded data frame onto the loop. The transmission of information continues until the allowed holding time for the write token (as indicated by the token holding timer) has expired or until all queues information in the transmit buffer logic 34 has been transmitted. The output control logic 38 then causes the write token to be regenerated by the token management logic 24, and the regenerated write token is gated from the token management logic 24 through the output control logic 38 to the encoder 40.

If the token received is a recovery token the token management logic 24 uses the control signals from the address recognition logic 28 to determine whether the destination address (DA) in the recovery token is less than, equal to, or greater than the station's address. The address recognition logic 28 includes a simple comparator to perform the comparison and supplies the token management logic 24 with control signals indicating the result of the comparison. If the destination address is greater than the station's address, the recovery token is gated through

-31-

the output control logic 34 into the encoder 40 and back onto the loop unimpeded. If the destination address in the recovery token is less than the station's address, the recovery token's DA is replaced with the station's 5 address and the modified recovery token is gated out onto the loop. If the destination address in the recovery token is equal to the station's address, the station has won the bid to recover the write token. In such case, a new write token is generated (by the same mechanism used 10 to generate a write token after completing the transmission of data frames) and gated out onto the loop. As previously mentioned, this bid mechanism allows only the station with the highest address to regenerate a write token after it is lost.

15 If the token received is a bandwidth allocation token, the allocation field (ALLOCATED) in the token is added to the current allocation for the station (which is stored in the token management logic 24) by the token management logic 24. If the result of this calculation 20 is greater than total amount of Class 1 bandwidth which is allowed to be allocated on the ring (ALLOC), the token management logic 24 notifies the transmit priority logic 32 and the unaltered bandwidth allocation token is immediately transmitted out on the ring through the output 25 control logic 38. As previously discussed, the latter situation can occur under the anomaly when two loops are joined together. If the result of this calculation is less than or equal to the total allocatable Class 1 bandwidth, the bandwidth allocation token is forward out onto the loop 30 with the new total in the ALLOCATED field.

If the frame coming out of the FIFO 22 is not a token (viz., it's a data frame or a frame fragment), the address recognition logic 28 will buffer the frame in the



-32-

receive buffer logic 42 if the frame's destination address is equal to (one of) the address(es) of the station, or cause the frame to be aborted if the station was the source of the frame (SA = station address). Note that a 5 station may recognize multiple addresses; viz., its station address, a broadcast address and (in some network implementations) a generic address. If the station's address is neither the source nor the destination specified in the frame, the frame is transmitted through the station 10 unimpeded.

The address recognition logic 28 uses timing signals from the frame synchronization logic 26 to supply signals from the address recognition logic 28: to the token management logic 24 indicating the result of recovery token 15 address comparison; to the frame creation logic 36 for aborting frames, and to the receive buffer logic 42 for frame and acknowledgment reception.

If the frame's destination address is equal to the station's address(es), the receive buffer logic 42 will 20 attempt to store the frame. If the station was the source of the frame, the frame creation logic 36 inserts an abort sequence into the frame. The abort sequence in the preferred embodiment is a frame delimiter (not shown in FIG. 3). After modifying the frame to indicate an abort 25 condition, the station retransmits the modified frame back onto the loop.

When comparing for destination addresses, the address recognition logic 28 will match both a unique 30 station identifier as well as any broadcast addresses which the station has been conditioned to recognize. If the destination address of the frame matches the station's unique identifier or broadcast address, the received buffer logic 42 causes an update of the status field of the frame



-33-

that is being received based on internal buffer conditions and control signals the receive buffer logic 42 receives from the error detection logic 30 (which computes the CRC). The receive buffer logic 5 42 accomplishes this status update by signaling the frame creation logic 36 to modify the end frame delimiter. In both the case of broadcast frame and non-broadcast frames, after updating the status the modified frame will be transmitted back onto the loop.

10 Those skilled in the art will appreciate that it is possible to implement all protocol functions (except token recovery and bandwidth allocation) in 1 bit of delay from FIFO 22 output to encoder 40 input by performing functions in parallel.

15 Having shown and described the preferred embodiment of the present invention, those skilled in the art will realize that various omissions, substitutions and changes in forms and details may be made without departing from the spirit of the invention. It is the intention, 20 therefore, for the invention to be limited only as indicated by the scope of the following claims.

-34-

What is claimed is:

1. A method of allocating bandwidth in a loop communications network, said network including a loop-connected set of stations providing data flow unidirectionally from one of said stations to the next, each of said stations being capable of passing data through itself along the loop, extracting data from the loop or injecting data into the loop, each of said stations having associated with it at least one identifier address, each of said stations being assigned a Class 1 priority, 5 a Class 2 priority, or both of said Class 1 and Class 2 priorities, the right of each of said stations to source new information into the network controlled by passing a write token from one of said stations to another, said method comprising the steps of:
 - 10 a. each of said Class 2 stations timing the rotation time of the write token to measure instantaneous load; and
 - 15 b. each of said Class 2 stations limiting the transmission of information onto the loop based on the observed write token rotation time.
- 20 2. The method in accordance with Claim 1 further including the step of allocating to the set of stations including all Class 1 stations a first portion of the bandwidth available on said loop.



-35-

3. The method in accordance with claim 2 further including the step of each Class 1 station requiring bandwidth reserving a portion of the first portion of bandwidth, wherein the sum of the reserved 5 Class 1 bandwidth does not exceed the first portion of the bandwidth.

4. The method in accordance with claim 3 further including a step wherein each Class 1 station that has reserved bandwidth and subsequently determines that it does not require its reserved bandwidth may surrender a 5 portion of its reserved bandwidth.

5. The method in accordance with claim 2 further including the step of allocating to the class of stations including all Class 2 stations a second portion of the bandwidth available on said loop.

6. The method in accordance with claim 5 wherein: each Class 2 station upon the receipt of a write token measures the time duration since the previous receipt of a write token; and 5 if the measured time duration is less than a predetermined time and that Class 2 station has queued information to transmit, the Class 2 station having received the write token will transmit said queued information onto the loop.



-36-

7. The method in accordance with claim 6 wherein that said Class 2 station having received the write token will transmit its said queued information onto the loop until the earlier of the transmission of all of its said 5 queued information or the expiration of a time equal to the difference between said predetermined time and said measured time duration.
8. The method in accordance with claim 6 or 7 further including the step of said Class 2 station transmitting a new write token onto the loop after transmitting said queued information.
9. The method in accordance with claim 5 wherein: each Class 2 station upon the receipt of a write token measures the time duration since the previous receipt of a write token; and 5 if the measured time duration is at least equal to a predetermined time, that Class 2 station will transmit a new write token onto the loop.



-37-

10. The method in accordance with claim 9 wherein:
if after transmission of information the
measured time duration at the next write token reception
is greater than the predetermined time, the Class 2
station will accumulate this lateness time;
5 on the next receipt of the write token by said
Class 2 station, if the measured time duration is at least
equal to said predetermined time, the difference in time
will be added to the lateness time;
10 if the measured time duration is less than the
predetermined time, the Class 2 station will subtract the
difference from the lateness time; and
15 if the accumulated lateness time is not greater
than zero, the Class 2 station will then transmit any queued
information.

11. The method in accordance with claim 1 further
including the steps of:
each station measuring the time since that said
station's last receipt of a write token; and
5 if the measured time since that said station
last received a write token exceeds a lost token time,
that said station will initiate a write token recovery
bidding process.



-38-

12. The method in accordance with claim 3 further including the steps of:

each Class 1 station checking whether the bandwidth reserved by Class 1 stations exceeds the first portion of the bandwidth; and

5 if the Class 1 bandwidth reserved exceeds the first portion, each said Class 1 station will surrender a portion of its reserved bandwidth.

13. A station for use in a loop communications network, said network including a loop connected set of said stations providing data flow unidirectionally from one of said stations to the next, each of said stations being 5 capable of passing data through itself along the loop, extracting data from the loop or injecting data into the loop, each of said stations having associated with it at least one identifier address, the right of each of said stations to source new information into the network 10 controlled by passing a write token from one of said stations to another, wherein the improvement in each of said stations comprises:

means for assigning each of said stations a Class 1 priority, a Class 2 priority, or both a Class 1 15 and 2 priority;

means in each Class 2 priority station for timing the rotation time of the write token for measuring instantaneous load on the loop; and

20 means in each Class 2 priority station for limiting the transmission of information onto the loop in response to the observed write token rotation time.

-39-

14. The station in accordance with claim 13 wherein each of said Class 1 stations includes means for specifying a first portion of the bandwidth available on said loop, said first portion corresponding to the 5 bandwidth allocated to the set of stations including all Class 1 stations.

15. The station in accordance with claim 14 wherein each Class 1 station further includes means for reserving a portion of the first portion of bandwidth, wherein the sum of reserved Class 1 bandwidth does not exceed the 5 first portion of the bandwidth.

16. The station in accordance with claim 15 wherein each Class 1 station that has reserved bandwidth includes means for determining that it no longer requires its reserved bandwidth and means for surrendering a portion 5 of its reserved bandwidth.

17. The station in accordance with claim 14 wherein each of said Class 2 stations includes means for specifying a second portion of the bandwidth available on said loop, said second portion corresponding to the pool of bandwidth 5 allocated to the set of stations including all Class 2 stations.

SUBSTITUTE SHEET



-40-

18. The station in accordance with claim 17 wherein each Class 2 station includes:

means for measuring the time between the receipt of one of said write tokens and the previous receipt of one of said write tokens; and

5 means, responsive to the measured time duration being less than a predetermined time and said Class 2 station having queued information to transmit, for the Class 2 station having received said write token 10 to transmit said queued information onto said loop.

19. The station in accordance with claim 18 wherein each of said Class 2 stations further includes means, responsive to having received one of said write tokens, said means for transmitting its said queued

5 information onto the loop until the earlier of the transmission of all of its said queued information or the expiration of a time equal to the difference between said predetermined time and said measured time duration.

20. The station in accordance with claim 18 or 19 wherein each of said Class 2 stations includes means for transmitting a new write token onto the loop after transmitting said queued information.

SUBSTITUTE SHEET



-41-

21. The station in accordance with claim 17
wherein:

each of said Class 2 stations includes means,
responsive to the receipt of a write token, for
5 measuring the time duration since the previous receipt
of a write token; and

means, responsive to the measured time
duration being at least equal to a predetermined time,
for transmitting a new write token onto the loop.

22. The station in accordance with claim 21
wherein each of said Class 2 stations further includes:

means, responsive to the measured time
duration being greater than the predetermined time at
5 the next write token reception following transmission of
information, said means for accumulating this lateness
time;

means, responsive to the next receipt of
one of said write tokens and a measured time duration at
10 least equal to said predetermined time, for adding the
difference in time to the lateness time;

means, responsive to the measured time
duration being less than the predetermined time, said
means for subtracting the difference from the lateness
15 time; and

means, responsive to the accumulated lateness
time not being greater than zero, for transmitting any
queued information.

-42-

23. The station in accordance with claim 13 further including:

means for measuring the time since that said station's last receipt of a write token; and

5 means, responsive to the measured time since that said station last received a write token exceeding a lost token time, for initiating a write token recovery bidding process.

24. The station in accordance with claim 15 wherein each Class 1 station includes:

means for checking whether the bandwidth reserved by Class 1 stations exceeds the first portion 5 of the bandwidth; and

means, responsive to the reserved Class 1 bandwidth exceeding the first portion, for surrendering a portion of its reserved bandwidth.

SUBSTITUTE SHEET






Fig. 1




Fig. 2



2/4

Fig. 3

START FRAME DELIMITER (SFD)	DESTINATION ADDRESS (DA)	SOURCE ADDRESS (SA)	INFORMATION (INFO)	FRAME CHECK SEQUENCE (FCS)	END FRAME DELIMITER (EFD)
-----------------------------	--------------------------	---------------------	--------------------	----------------------------	---------------------------

DATA FRAME FORMAT

START FRAME DELIMITER (SFD)	DESTINATION ADDRESS (DA)	FRAME CHECK SEQUENCE (FCS)	END FRAME DELIMITER (EFD)
-----------------------------	--------------------------	----------------------------	---------------------------

RECOVERY TOKEN FORMAT



START FRAME DELIMITER (SFD)	DESTINATION ADDRESS (DA)	SOURCE ADDRESS (SA)	ALLOCATED BANDWIDTH (ALLOCATED)	FRAME CHECK SEQUENCE (FCS)	END FRAME DELIMITER (EFD)
-----------------------------	--------------------------	---------------------	---------------------------------	----------------------------	---------------------------

BANDWIDTH ALLOCATION TOKEN FORMAT

START FRAME DELIMITER (SFD)	END FRAME DELIMITER (EFD)
-----------------------------	---------------------------

WRITE TOKEN



Fig. 4AFig. 4B




Fig. 5

INTERNATIONAL SEARCH REPORT

PCT/US83/00282

International Application No

I. CLASSIFICATION OF SUBJECT MATTER (If several classification symbols apply, indicate all) ³

According to International Patent Classification (IPC) or to both National Classification and IPC

INTL. CL. ³ H04Q 9/00;

U.S. CL. 340-825.51; 370-86

II. FIELDS SEARCHED

Minimum Documentation Searched ⁴

Classification System	Classification Symbols
U.S.	340-825.5, 825.51 370-94, 16, 84, 86, 87, 88, 89, 95, 17, 100, 102, 103, 105, 106 375-106, 117, 116

Documentation Searched other than Minimum Documentation
to the Extent that such Documents are Included in the Fields Searched ⁵III. DOCUMENTS CONSIDERED TO BE RELEVANT ¹⁴

Category ¹⁴	Citation of Document, ¹⁴ with indication, where appropriate, of the relevant passages ¹⁵	Relevant to Claim No. ¹⁵
Y	US, A 3,903,371 (Colton et al) 02 September 1975	1;11
Y	US, A 4,001,504 (Hendrickson) 04 January 1977	1;11
Y	US, A 4,161,719 (Parikh et al) 17 July 1979	1;11
Y	US, A 4,316,284 (Howsan) 16 February 1982	1;11
P, Y	US, A 4,344,180 (Cummiskey) 10 August 1982	1;11

* Special categories of cited documents: ¹⁵

"A" document defining the general state of the art

"E" earlier document but published on or after the international filing date

"L" document cited for special reason other than those referred to in the other categories

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but on or after the priority date claimed

"T" later document published on or after the international filing date or priority date and not in conflict with the application, but cited to understand the principle or theory underlying the invention

"X" document of particular relevance

IV. CERTIFICATION

Date of the Actual Completion of the International Search ¹⁶

15 June 1983

Date of Mailing of this International Search Report ¹⁷

23 JUN 1983

International Searching Authority ¹⁸

ISA/US

Signature of Authorized Officer ¹⁹

Donald J. Yusko