

Description

This invention relates to communications multiplexers for use in communications systems such as communications networks.

Systems providing for the transport of encoded signals on a communications link are well known. Many such systems are used for data signals which are not time critical; examples of such systems are packet switching networks and many local area networks (LANs). Many others can carry speech and video signals, where transmission delays above certain limits are unacceptable, but not data signals.

IBM Technical Disclosure Bulletin, Vol 26 No 11 April 1984 p5991-2 describes a token ring communication system for transmitting several frames before a free token is released. Token ring systems allow only a single token at a time onto the ring; messages can be long and different messages usually have different lengths. Such systems are used for carrying data. The rules for the technique described in the IBM Bulletin ensure that a station can have access to the ring for a certain time, before releasing the token. There is no mechanism for ensuring that access delays are kept within set limits, and the ring could be unsuitable for carrying services such as voice or low bit rate video which have maximum delay requirements. There are obvious advantages in being able to use a single system for both data and speech and other synchronous services. For example, there would be a major cost saving over separate data and speech networks. However, there are numerous problems in achieving a satisfactory dual system particularly at the bit rates required to transport voice and low bit-rate video. Existing protocols for low bit rate systems would suffer from some or more of the following drawbacks: excessive delay and/or jitter; inability to offer guaranteed bandwidth for the duration of delay sensitive service calls; no effective overload control; and inefficient use of bandwidth.

The Bell System Technical Journal, Vol. 51 No 6 July/August 1972 pages 1147-1165, "An Experimental Data Block Switching System" - W J Kropf proposes a slotted ring type system with an anti-hogging mechanism implemented by a "hog-control" field in each block and a memory device in each station. An "A" station is provided which serves to close the ring and provide timing information for all messages on the ring. It also performs a centralised resetting function and overload control. Each "B" station has a "write-request" line which it uses to indicate the desire to transmit. "B" stations can write blocks into vacant slots on the ring and may continue to do so provided that subsequent blocks have an un-set control field (if they are vacant and the hog-control field is set then

"B" stations can only seize the block if they have not yet transmitted). If a full block is encountered and the "B" station requires to transmit then it will set the hog-control field. In this way, all "B" stations have an equal opportunity of access to the ring. This system does not allow for preferential service for any particular station and would not be suitable for supporting a mixture of voice and data services. The resetting of the hog-control field which determines station access to the ring is initiated by a single "A" station and complete failure of the ring would therefore result from an "A" station failure. Also, since this reset function is centralised, "B" stations do not have a direct method of monitoring ring loading and therefore do not have an indication of the service that they could support.

AT&T Bell Laboratories Technical Journal, Vol 63 No 2 February 1984 pages 307 to 334, "Integrated Voice/Data Services on Fastnet" - J W Mark & J O Limb - describes an "implicit" token passing system - no actual token is used but a rigid service order is observed, i.e. station 1, station 2,station n. The system requires two buses (one for forward and the other for reverse transmission). The stations are not specifically for voice or data (synchronous or asynchronous respectively) since voice and data "cycles" are described which occur at different times. Voice is given higher priority than data. Stations can write into a maximum number of free slots of the correct type. A head station exercises centralised access control which permits access to either voice or data services by issuing a free slot of the correct type. This type of system is relatively inflexible in that allocation of bandwidth to stations will tend to be unfair and will be biased towards stations located physically close to the head station. Also, since two unidirectional buses are used (and not a ring), stations are required to know which bus to use in order to transmit to any other station.

United States Patent No. 3,988,545 (Kuemmerle et al.) discloses a multiplexer in which data from synchronous sources is received in respective registers which are addressed for read-out at appropriate times of a frame. Data from asynchronous sources is formatted into packets and fed to a buffer store which is controlled for read-out in accordance with the gaps between the synchronous data channels in the frame.

It is often desired to transmit information between rings and loops, or other networks, on communications links. An aim of the present invention is to provide a link multiplexer which enables queues of blocks to be transmitted to a link without allowing any queue to hog the link.

According to the present invention there is provided a communications multiplexer comprising

a plurality of queues for message blocks to be transmitted onto a communication link characterised by control means arranged to provide an allocation of d blocks to each queue, to monitor transmission from the queues onto the link, and to inhibit transmission of further blocks from each queue which has transmitted d block onto the link, until a further allocation is made. Note that d can be any suitable integral value, including zero.

Preferably, the control means comprise a counter for each queue and the counters are reset when every queue is inhibited or empty.

Preferably, a maximum allowable interval between resets is set according to delay requirements for message blocks for time sensitive services.

Preferably, respective queues are provided for time sensitive services having different maximum delay intervals, and the maximum allowable interval between resets is no greater than the maximum delay interval for the queue having the smallest such interval.

Prior to accepting a new time-sensitive service call, the rate of resets may be compared with the maximum allowable interval between resets to determine whether the call can be accepted without the reset interval rising above the maximum, and on acceptance of the call, increasing the allocation d for its queue.

Preferably, each queue has an assigned priority and the control means polls the queues in priority order.

Preferably, a queue is provided for non-time sensitive services, said queue having a small, fixed allocation d .

Thus it will be seen that a queue which is transmitting is allowed to seize only a given number of empty blocks passing along the link before it has to pause to allow other queues which require to transmit to gain access to the link. Any queue which has paused is only allowed access to the link after the other queues requiring to transmit have had the opportunity of using their allocation of message blocks.

The invention will now be described, by way of example, with particular reference to the accompanying drawings. In the drawings:

Figure 1 is a schematic diagram illustrating a communications network in which a communications multiplexer of the present invention may be used;

Figure 2 shows the structure of a slot of the network of Figure 1;

Figure 3 is a schematic diagram illustrating one example of a station on the network of Figure 1;

Figures 4 to 7 are flow diagrams of various functions carried out at a station of the network of Figure 1;

Figure 8 is a schematic diagram illustrating an embodiment of a communications multiplexer in accordance with the invention; and

Figure 9 is a diagram showing a network of rings and links.

Referring to Figure 1 of the drawings, a communications network system comprises a plurality of stations 10 which can communicate via a transmission ring or loop 11. The stations 10 can incorporate many forms of digital based devices such as data processing equipment, video devices, facsimile or telephone equipment and each station may concentrate traffic from several devices. Also one station could for example provide access to the public switched telephone network. The loop is designed to operate according to what is known as a slotted ring protocol. In this type of arrangement one or more slots or blocks each of which comprises a predetermined number of bits circulates around the ring and can be seized by a station which wishes to transmit data to another station. The ring carries a fixed integer number of equal length slots which are established at switch-on and maintained continuously by a station which acts as monitor. In case of failure of the monitor station, another station can take over the monitor function.

In the present arrangement when a station wishes to transmit data to another station it is permitted to seize up to a total of d empty slots (which need not be successive slots) and when it has used that number d it is automatically placed in a state, known as a pause state, in which it is not permitted to seize further slots. That station is not allowed to seize further slots until it is reset to an active state. If other stations are waiting to transmit data the original station cannot be reset to an active state until those further stations have had an opportunity of using up their allocation of empty slots. The way in which this is achieved will become apparent from the following description.

Referring to Figure 2 of the drawings each slot comprises two sections each made up of a number of bits. Section 20 is a control field 20 and the second section 21 comprises the data bits which carry the information to be transmitted to another station. The control field 20 has a number of bits in which bit 1 is a full/empty bit, bit 2 is a trial bit, bit 3 is a monitor bit, bit 4 is a priority indication bit, bit 5 is a broadcast bit and bits 6 and above define a code representing the destination address (DA). The destination address defines only the station to which the information carried in the data field is direction. Additional addressing to route that information to the appropriate location within the station is contained within the data field 21.

In operation when a station wishes to transmit data to another station it identifies an empty slot circulating on the ring 11, renders the bit 1 to a

condition indicating a full slot and inserts the appropriate destination address DA into bits 6 onwards. The data to be transmitted is inserted into the data field 21. The slot then travels around the ring and the destination station recognises the destination address and accepts the data in the data field. The slot is emptied at the destination station, the bit 1 is returned to its condition indicating an empty slot and that empty slot is then passed on to the next adjacent station which can make use of the slot if it has data to transmit.

Any station which has been rendered active may seize up to d empty slots. Once the threshold value d has been reached, which is sensed by a counter at the station, the station enters a pause or inhibited state in which it cannot seize any further slots even though it may have further data to transmit. The station remains in that pause state until it is reset to an active state as will be described below. It will be appreciated that subject to the d threshold restriction an active station can have several slots on the ring at any time. During a pause state the station can continue to receive data slots from other stations but cannot transmit data.

When a station is in a idle state or in a pause state it permanently carries out a function even though it is not transmitting data. This action is as follows. Provided that station is not marked as having been set to an outstanding reset condition, which will become apparent later, then whenever it identifies a passing empty slot whose trial bit (bit 2 of the control field) is off it turns that trial bit on in that slot, loads its own address into the DA bits and sets itself as having an outstanding trial. The full/empty bit of the slot is not set so the slot is available to other stations for transmission. If any station subsequently senses an empty slot in which the DA bits match its own address and in which the trial bit is on it responds as follows depending upon the outstanding trial condition. If the station has an outstanding trial indicated, the full/empty bit (bit 1 of the control field) of the slot is converted to a full condition. That station then reverts to an active or idle state according to whether or not a packet is waiting, sets its d counter to zero and loads into a d threshold store the next d threshold value. Finally it cancels the outstanding trial indication and marks itself as having an outstanding reset. If an outstanding trial is not indicated then the trial bit of the slot is turned off and the DA bits are all made zero.

A station which senses a slot whose full/empty bit is on and whose trial bit is on also responds as follows depending upon its outstanding reset condition. If it does not have an outstanding reset the station reverts to an active or idle state, zeros its d counter and loads into its threshold store the next d threshold value. Finally it cancels any outstanding trial marking. If an outstanding reset is indicated

then the trial bit of the slot is turned off together with the full/empty bit and the DA bits are made zero. Finally the outstanding reset indication is cancelled.

Thus, it will be seen that a station which wishes to transmit data can, when it is made active, use up to d slots before having to pause from the transmission of data. It is then not allowed to seize further empty slots. However, the station can set a trial bit in an empty slot in order to test whether any other station on the ring is waiting to transmit. If there is such a station waiting then that station will seize the empty slot and delete the trial bit which was set by the original station and then that further station will have an opportunity of transmitting its allocated number of d slots. When it has transmitted its allocated number it will also test the ring. This procedure carries on until a stage is reached where there are no more stations waiting to transmit and an empty slot with its trial bit set will return to the originating station and this originating station responds by setting the full/empty bit (bit 1) to indicate a full condition, and transmitting the slot. All stations then respond to this combination of a slot marked as full and with its trial bit on and their d allocations are reset. This is known as resetting or refreshing the ring. It should be noted that the time between the first node reaching a pause state and the ring being refreshed is usually short in comparison with the time between each refresh.

It will be appreciated that the d values can be chosen to provide each station with the bandwidth which it requires. If it is assumed that each station has the same d value then the instantaneous demand for bandwidth will vary from station to station being higher for example at stations with video transmissions. Some stations may therefore use their d allocation much quicker than others. The protocol described above ensures that the first station to use its allocation of d slots pauses momentarily thereby giving other stations the opportunity to transmit any data that has accumulated before the ring is refreshed and the high demand station can continue. In this way stations are guaranteed an equal share of the available bandwidth if required. A ring with fully loaded stations will share its bandwidth equally with a minimum guaranteed bandwidth being available at any station. This is an important feature of the present arrangement.

In practice it can be shown that delays are minimised if the minimum guaranteed bandwidth for each station is closely matched to its requirements. This can be achieved by allowing the d value of a station to be selected to correspond with its expected demand. Thus, for example, a station can be given as a minimum bandwidth the equivalent of one data channel in addition to which it can

claim extra bandwidth for telephony and video. Consider a station which is using a single video channel of 32 slots in a repetition period which is not more than 2 msec. The internal counter of the station is set so that it can tell when 32 slots have been transmitted following which it enters the pause condition and thus allows other stations to transmit until it is reset. If the ring is lightly loaded refreshing will occur after a few ring rotations and stations will then be able to seize extra bandwidth. It can be determined that the refresh period shall not be longer than 2 msec so that each station examines the refresh period and if it is close to 2 msec such that there would not be sufficient empty slots for a node to seize 32 slots to make an additional video call (and leave a minimum of empty slots so that there was not in fact 100% utilisation of ring) then a node would not be able to transmit this additional load and that would be blocked. All stations are able to determine for themselves whether the refresh period is too close to 2 msec for them to seize extra bandwidth.

The control of the *d* value is carried out using a *d* threshold store at each station which holds two stored values of *d*. One value, representing the current *d* allocation, is the reference value controlling entry into the pause state. The other is the "next *d* threshold", which replaces the reference value on every reset. Both read and write access is provided on "next *d* threshold" allowing a separate load control module to control the value of *d* using a higher level protocol remote from the ring protocol. This is described in more detail below with reference to Figure 3. If the maximum allowable reset interval is increased, the load control module can then increase the value of *d* accordingly. In effect the protocol provides a guaranteed bandwidth for established calls.

The details of the control of the next *d* threshold are as follows. First, next *d* threshold is initialised with some suitable small value eg. next *d* threshold equals 2. The purpose of this is to allow stations to support the low speed data service and signalling requests. The value of 2 would imply that as the time between successive resets approaches 2 msec the data service would still be 128 kbits/sec available at each station assuming a 128 bit slot information field. The fixed service is referred to as the background service.

When further low-speed data connections are set up on the ring the threshold *d* is not increased above the background value of 2. If there are no real-time services, the *N* stations will be able to supply 2*N* slots (assuming a background of 2 for all stations) before all stations reach their pause states and a reset occurs. Hence a ring loaded only by data connections has short reset intervals, although the ring may be fully loaded, and each

station obtains a high bandwidth. The load of real-time calls on the ring can be built up in a controlled way as follows. When a new call set up request is made (speech, video or fast file transfer) the current average reset time is examined. If this indicates that there are sufficient empty slots within the 2 msec period to support the new service then the new call request is permitted. The next *d* threshold is then updated by an amount suitable for the additional connection (eg. 2Mbit/sec video would require an additional 32 so the next *d* threshold would then be 32 plus background). As real-time calls are established, therefore, the effect is to reduce the bandwidth available to lowspeed data calls. Similarly, on termination of a call the next *d* threshold is decremented accordingly. In fact additional thresholds can be set for the average reset interval and for each service time so that a service request is blocked if the current average reset interval is above the appropriate threshold.

The maximum allowable period between resets is determined according to the requirements of real-time service on the network. For speech calls not on the public network, maximum allowable delays are around 2 msec. If no other real-time service (such as video) has more stringent delay requirements, then 2 msec is taken as the maximum delay interval.

If for some reason the average reset time becomes longer than 2 msec the background service may be temporarily inhibited (eg by decreasing the next *d* threshold by an amount equal to the background value).

Restoration of the background occurs as follows. Whenever the latest reset time drops below some suitable threshold (less than 2 msec) the background is added. This carries on until the full background is returned. In this way the reset times can be controlled to permit services which are sensitive to packet loss.

It will be observed that once a call has been set up and bandwidth allocated for it, that call will be able to continue until it is terminated by the called or calling party (unless there is node failure, for example). This is in contrast to conventional ring systems such as the Cambridge ring where overloads result in buffer overflows, with consequent packet loss, or the slowing down of attached devices; this is not acceptable for real-time services, where overloads should result in unsuccessful call attempts rather than degradation or premature release of calls in progress.

Referring to Figure 3, this shows a possible design of a station 10. The station comprises access controller 51 which operates the protocol described above. Controller 51 comprises a reference store holding the two stored values of *d* for the station and provides for transmission of slots onto

the ring 11 from transmit buffer 53. It also provides for reception of slots from the ring into receive buffer 55. Updates to "next d threshold" are received from load monitor 57. A pulse is sent to load monitor 57 on each reset. The load monitor counts how many reset signal pulses are received per maximum allowable reset interval, say 2 msec. Information on reset rate is also transmitted to controller 59 which receives call set-up requests in the normal way from attached devices. A new call requiring to send M slots per 2 msec interval can be accepted if the number of resets P per 2 msec is greater than or equal to M/L where L is the number of slots on the ring. If the reset interval rises above 2 msec, the load monitor 57 may reduce the threshold d by the background allocation.

On indication from the transmit buffer 53 that a further slot can be accepted, select poller 61 polls synchronous and asynchronous buffers 63 and 65 containing ready-assembled slots to see which waiting slot has higher priority, and this slot is transmitted to the transmit buffer 53 for onward transmission to the ring 11. The rate of acceptance of low priority slots is controlled by the ring reset rate as will now be explained. On each reset, a reset signal is sent to the poller 61. This resets a counter within the poller which enables the acceptance of low priority packets until the counter senses that the background allocation has been transmitted. Entry of low priority packets into the transmit buffer 53 then proceeds whenever the transmit buffer 53 indicates that a further packet can be accepted and provided that no higher priority packet is waiting. Thus the rate of entry of low priority slots into buffer 53 is controlled by the rate of reception of reset signals by poller 61, and the correct mix of low and high priority packets is supplied to buffer 53 by poller 61. Load monitor 57 may inhibit the supply of reset pulses to poller 61 if the reset interval increases above the 2 msec maximum.

Slots for buffers 63 and 65 are provided from synchronous packet assembler/disassembler (SPAD) 67 and asynchronous packet assembler/disassembler (APAD) 69, respectively. SPAD 67 creates slot-sized packets from time division multiplexer (TDM) equipment 71, comprising terminating equipment (eg telephones) 73 and line units 75 which send signalling information to controller 59. On each call set up an appropriate header is set up by controller 59 and supplied to SPAD 67 for use on all slots for that call. All SPAD slots have high priority and are stored in buffer 63.

The maximum allowable period between resets is determined according to the requirements of real-time services on the network. For speech calls not on the public network, maximum allowable de-

lays are around 2 msec. If no other real-time service (such as video) has a more stringent delay requirement, then 2 msec is taken as the maximum delay interval.

APAD 69 creates slot-sized packets from an arriving data packet stream from data terminating equipment 77 and generates and verifies a frame check sequence. Normally, data slots have low priority. If a particular data connection requires guaranteed bandwidth, signals are sent to controller 59 and load monitor 57 to increase the value of d in access controller 51. In this case, the corresponding slots are marked for high priority.

Figure 4, 5, 6 and 7, respectively, are flow diagrams summarising the control, receive, transmit and reset functions carried out by access controller 51.

The ring network described above may be used as a local area network (LAN) where the length of the ring is typically a few kilometres, and around 20 to 30 stations on the ring communicate with workstations, computers, VDUs, telephones etc around the office block or site served by the LAN. Alternatively the ring may be much smaller (although the number of stations may be similar) and, together with other similar rings, form part of a high speed switch for speech, data etc. For public switching, the maximum interval between resets would need to be smaller than the figure of 2 msec quoted above for a LAN. To meet CCITT requirements, resets would need to occur at least every 125 microseconds.

Each station, as shown in Figure 3, is provided with slots from different sources, for example speech slots from equipment 71 and data slots from equipment 77. In a switch, slots of different types from one ring need transferring to another.

Figure 8 shows a multiplexer suitable for multiplexing slots from several buffers containing queues of slots, onto a communications link. The multiplexer uses essentially the same protocol as the ring described above. This is advantageous because it enables slots to be transmitted over the whole network system without having to use gateways to rebuild slots to meet different requirements for different parts of the system. Overheads in moving from ring to link are therefore avoided.

Multiplexer 30 comprises a number of input first in -first out (FIFO) buffers comprising queues of fixed length blocks or slots awaiting transmission onto a link 40. Three queues 31, 32 and 33 are shown in Figure 8 out of a total of, for example, six or seven queues. The embodiment of multiplexer 30 described below is adapted to use a protocol very similar to that of the slotted ring network described above the reference to Figure 1. The differences between the protocols arise because the link multiplexer has a centralised controller,

whereas the stations on the ring are independently controlled. The controller of the link multiplexer knows the status of all its queues and can generate reset signals at appropriate intervals. Stations on the ring in pause states have to carry out trials to find when a reset can occur. A number of rings 11 used for switching or as LANS may be interconnected by multiplexers 30 as shown in Figure 9.

As shown in Figure 8, queues 31, 32 and 33 communicate via a common data and control bus 34 with a transmit FIFO buffer 35, which transmits slots from the queues onto communications link 40.

A queue is provided for each type of service requiring access to link 40. Queue 31 comprises data slots, for which access to the link is not time critical. Other services require access to the link within set periods, according to the nature of the service. Speech, for example, may need to send one slot of 160 bits, having a 128 bit information field, every 2 milliseconds to satisfy error and grade of service specifications. Voice slots and any other service needing to send a slot every 2 milliseconds are provided with a queue 33. Queue 32 is for low variable bit rate (VBR) video slots. Further queues (not shown) are provided for further synchronous fixed and variable bit rate services having different delay limits. Each queue has a designated priority. In this example, queue 33 (voice) has the highest priority and queue 31 (data) the lowest priority. Priority is determined by position, the highest priority queue being on the left and the lowest on the right. Alternatively queue priority could be stored in a register.

The multiplexer has a controller 42 which communicates with bus 34 and comprises access controller 41 and load monitor 36. Access controller 41 includes a polling mechanism 43 which polls the queues in priority order to see if they have a slot to transmit. The polling mechanism also includes a counter (not shown) for each queue and a reference store consisting of a read and write memory (not shown) which stores the current threshold, or allocation of slots, for each queue and also a "next threshold" value. Load monitor 36 communicates with the access controller 41 and receives from controller 41 data idle signals on line 37 and reset signals on line 38. Load monitor 36 also outputs threshold signals to controller 41 on line 39. Using these signals, the load monitor 36 in conjunction with access controller 41 serves to vary dynamically the bandwidth allocated to each of queues 31, 32, and 33, and thus guarantee a minimum bandwidth allocation to each queue.

In operation, queues 31, 32, and 33 are supplied with constant length message slots requiring transmission to the outgoing link 40.

As described above for a slotted ring, each queue is allowed to transmit a number, d , of slots,

and when a queue has transmitted its allocation of d slots, it is inhibited from transmitting further slots until the counter associated with that queue has been reset. A reset occurs only after all queues have had the opportunity of transmitting their full allocation of slots. Some queues may be empty, or become empty, before their full allocation of d slots has been transmitted onto link 40; these, however, together with queues which have transmitted their d allocation, will have their counters reset as soon as no queue can transmit any further slot. Of course, unlike the slotted ring, all slots queuing are full slots and no writing in the control field is necessary: slots are transmitted onto the link without any bits being altered. Each queue may have a different value for d , and the value of d for each queue can be varied.

The operation of the multiplexer will now be described in greater detail.

The transmit buffer 35 needs to be supplied with slots from queues 31, 32 and 33 at the rate at which it can unload slots onto link 40. Polling mechanism 43 carries out the following sequence in order to transmit slots to buffer 35 at the required rate. The sequence is initiated on detection by poller 43 of an input ready signal (logic 1) on the transmit buffer 35. The poller 43 checks the highest priority queue (queue 33) for an output ready signal and also checks the status register to see if the queue is inhibited from sending further slots (ie if it is in a "pause" state - which has logic 0). If the queue 33 has a slot to transmit and it is not inhibited (both logic 1 states), an output enable is put on its buffer and the slot is put onto bus 34 and delivered to buffer 35 for transmission onto link 40. The counter in access controller 41 associated with queue 33 is incremented by one. When the poller 43 next detects an input ready signal on the output buffer 35, this procedure is repeated. Once the full allocation of d slots for queue 33 has been transmitted to link 40, the counter for that queue contains the value d , the status register for queue 33 marks that queue as inhibited and the queue 33 enters the "pause" state. In this state, queue 33 is prevented from transmitting further slots until its counter is reset. Poller 43 now fails to detect two logic 1 states for queue 33 and moves to the queue to the right (not shown) which has the next highest priority. If that queue has no slots to transmit, the next queue to the right, say queue 32, is polled and a waiting slot is transmitted onto line 40. If, during transmission of that slot a new slot arrives at the higher priority queue to the left (previously empty), that slot will be the next to be transmitted as the poller always returns to the highest priority queue when an input ready signal is detected on buffer 35. The system therefore progresses towards a situation in which all queues are either in a

pause (ie inhibited) state or an idle state in which no slots are waiting for transmission.

When the polling mechanism finds that no slots can be selected it causes a reset of all the queue counters and at the same time it sends a reset pulse to load monitor 36 via reset line 38. Load monitor 36 keeps a count of the number of reset pulses received in a given time interval. Each occurrence of a reset corresponds to a slot on outgoing link 40 remaining unused, and the frequency of resets determines the spare capacity which is available on link 40. The load monitor 36 uses the rate of resets to determine call acceptance for synchronous time-critical services such as speech as described below.

The value of threshold d for each queue is the maximum number of slots that the queue can transmit onto link 40 between resets. Thus the maximum time between resets is the time for the full allocation of slots to be transmitted from all queues. To meet the stringent time constraints placed on synchronous services such as 64 kbit/sec speech, it is necessary for delays between consecutive slots associated with a particular call (each having a 128 bit information field) to be less than, for example, 2 msec. The maximum time between resets must be selected according to the delay limit of the queue having the most stringent time constraints. If one queue must send a slot every 62.5 microseconds and other queues have maximum delays of 2 milliseconds or more, 62.5 microseconds is taken as the maximum allowable interval between resets. Suppose, for example, that queue 33 is servicing 64 established 64 kbit/sec speech calls, each needing to offer 1 slot every 2 msec, and the maximum allowable interval between resets is 62.5 microseconds. In this case, the threshold for queue 33 needs to be 2 in order to handle the calls without unacceptable delays. There is thus a guaranteed minimum acceptance rate of 64 slots per 2 msec interval, and queue 33 is guaranteed the necessary bandwidth to service its calls. Load monitor 36 is used to monitor synchronous services and to make dynamic adjustments to the thresholds for queues of synchronous services to match the number of calls in progress whilst maintaining the reset interval below the maximum allowable. The former of these is implemented via line 39 between the load monitor 36 and the polling mechanism in access controller 41. There are two stored values of each queue threshold maintained by the poller 43 in its reference store. One value, representing the current bandwidth allocation (d) for the queue, is the reference value controlling entry into the pause state for the current reset interval. The other value, known as next threshold, is loaded in place of the existing value, as the new bandwidth allocation (d), on every reset.

Both read and write access is permitted for load monitor 36 on next threshold. When a new connection for queue 33 is set up, the number of message slots which the new call will require to send within a given service delay limit is sent to the load monitor 36. In the above example, the message slot has an information field of 128 bits, and a new 64 kbit/s voice call will require to send one slot within each 2 millisecond period which, for efficiency reasons, is chosen as the service delay limit. The load monitor 36 adjusts the threshold to maintain the correct guaranteed call acceptance rate. Thus, if the maximum allowable reset interval is 62.5 microseconds, the threshold would need to be incremented by one for each additional 32 64 kbit/s voice connections established on link 40.

Before adjusting the threshold, load monitor 36 must check that acceptance of the new voice connections will not increase the reset interval above the maximum allowable reset interval. A connection for a new call for queue 33 will only be established if the frequency of resets exceeds the guaranteed minimum. For example, given a maximum allowable reset interval of 62.5 microseconds, the target minimum number of resets in any 2 millisecond period is 32. Hence a new 64 kbit/sec voice connection offering one slot every 2 milliseconds can only be established if there are currently at least 33 resets per 2 milliseconds.

It will be appreciated that queues are set up according to the nature of the services using the multiplexer. For example, synchronous constant bit rate (SBR) services requiring a maximum link access delay of 125 microseconds will be assigned to one queue and SBR services with a maximum delay limit of 2 milliseconds to another. Additional queues may be provided for further SBR services and dynamically varying bit rate (VBR) services with different average maximum delay intervals. Normally only one queue need be provided for data.

The link multiplexer of Figure 8, unlike the ring of Figure 1, has a centralised controller 42, (comprising load monitor 36 and access controller 41). Control by load monitor 36 is restricted to delay critical synchronous services. For data queue 31, there is a fixed threshold. This is chosen to be a small value, eg 1, in order to ensure some acceptable minimum bandwidth on link 40 even when resets are spread apart at the target maximum interval. However, when link 40 is loaded mainly or entirely by data, resets will be very frequent. Hence a high bandwidth is obtained by data queue 31 when few time-sensitive services are connected to the multiplexer.

In the situation where synchronous connections are established on link 40 previously loaded only by data message slots, the reset count will initially

be high even though link 40 may have been fully loaded. This high reset count indicates that more time critical calls can be accepted and the threshold for queue 33 increased. Hence set up of speech calls proceeds and the threshold for their queue 33 is suitably increased. This reduces the reset count which means that data queue 31 receives a smaller bandwidth allowance. Thus a dynamically adjusted boundary is provided for allocation of bandwidth on link 40 for synchronous and asynchronous use.

In order to provide an estimate of spare data capacity, the polling mechanism in the access controller 41 supplies a pulse via line 37 to the load monitor 36 for each idle reset of data queue 31 (i.e. a reset which occurred when the data threshold had not been reached). Thus, if a two millisecond period contained, for example, 10 idle resets this would indicate that a new data connection could be established which required to send a maximum of 10 slots every two milliseconds. However this maximum is not guaranteed; the main purpose of the count is to prevent unnecessary loading of data queue 31 by new connections when existing capacity is fully utilised.

To establish a variable bit rate (VBR) connection, the threshold for queue 32 is updated. Load monitor 36 adjusts the threshold increment for an estimated mean rate of transmission of the connection. Video services are subject to time constraints in the same way as speech services; however in order to keep costs down for video customer, high grade video channels are only made available to customers when the system is lightly loaded. This is achieved by setting the threshold (d) for queue 32 lower than the maximum bandwidth. Most of the time, the frequency of resets will be much faster than the minimum frequency and the video queue will be able to put the necessary number of slots onto link 40. If, however, the other queues become busy and the time between resets approaches the maximum, only a minimum bandwidth is available to the video customer. Permission to establish a new VBR connection will use the current reset count just as for synchronous fixed bit rate connections such as voice calls.

In addition to methods for restricting the load to preserve a bounded reset interval, the polling mechanism registers an overload condition if it is unable to perform a reset within the maximum interval. At this point it immediately causes all queue counters to be reset thereby allowing high priority queues to resume service.

Output buffer 35 is responsible for sending a frame alignment slot at regular intervals to enable a distant receiver on link 40 to be synchronised to the outgoing message slot stream. Output buffer 40 also incorporates a register to provide link 40 with

a dummy full slot to be put on the link whenever no slot is received from any of the queues 31, 32 and 33, ie whenever no slot can be selected from the queues and a reset is taking place

Figure 9 shows three rings 11, used for switching purposes, interconnected by links 40. Each ring has a station 10', providing access to links 40. Stations 10' are the same as stations 10 shown in Figure 3 except they also include buffers for queues of slots (31, 32, 33 etc), a transmit buffer 35, and a controller 42 for each link 40 terminating at each station 10'.

Claims

1. A communications multiplexer comprising a plurality of queues (31, 32, 33) for message blocks to be transmitted onto a communication link (40) characterised by control means (42) arranged to provide an allocation of d blocks to each queue, to monitor transmission from the queues (31, 32, 33) onto the link (40), and to inhibit transmission of further blocks from each queue which has transmitted d block onto the link, until a further allocation is made.
2. A multiplexer as claimed in claim 1, characterised in that the control means (42) comprises a counter for each queue (31, 32, 33) for counting the number of blocks transmitted onto the link (40) and the counters are reset when every queue is inhibited or empty.
3. A multiplexer as claimed in claim 2, characterised in that a maximum allowable interval between resets is set according to delay requirements for message blocks for time sensitive services.
4. A multiplexer as claimed in claim 3, characterised in that respective queues are provided for time sensitive services having different maximum delay intervals, and the maximum allowable interval between resets is no greater than the maximum delay interval for the queue having the smallest such interval.
5. A multiplexer as claimed in claim 3 or claim 4, characterised in that prior to accepting a new time-sensitive service call, the number of resets per maximum allowable interval between resets is used to determine whether the call can be accepted without the reset interval rising above the maximum, and on acceptance of the call, increasing the allocation d for its queue.
6. A multiplexer as claimed in any one of claims



be high even though link 40 may have been fully loaded. This high reset count indicates that more time critical calls can be accepted and the threshold for queue 33 increased. Hence set up of speech calls proceeds and the threshold for their queue 33 is suitably increased. This reduces the reset count which means that data queue 31 receives a smaller bandwidth allowance. Thus a dynamically adjusted boundary is provided for allocation of bandwidth on link 40 for synchronous and asynchronous use.

In order to provide an estimate of spare data capacity, the polling mechanism in the access controller 41 supplies a pulse via line 37 to the load monitor 36 for each idle reset of data queue 31 (i.e. a reset which occurred when the data threshold had not been reached). Thus, if a two millisecond period contained, for example, 10 idle resets this would indicate that a new data connection could be established which required to send a maximum of 10 slots every two milliseconds. However this maximum is not guaranteed; the main purpose of the count is to prevent unnecessary loading of data queue 31 by new connections when existing capacity is fully utilised.

To establish a variable bit rate (VBR) connection, the threshold for queue 32 is updated. Load monitor 36 adjusts the threshold increment for an estimated mean rate of transmission of the connection. Video services are subject to time constraints in the same way as speech services; however in order to keep costs down for video customer, high grade video channels are only made available to customers when the system is lightly loaded. This is achieved by setting the threshold (d) for queue 32 lower than the maximum bandwidth. Most of the time, the frequency of resets will be much faster than the minimum frequency and the video queue will be able to put the necessary number of slots onto link 40. If, however, the other queues become busy and the time between resets approaches the maximum, only a minimum bandwidth is available to the video customer. Permission to establish a new VBR connection will use the current reset count just as for synchronous fixed bit rate connections such as voice calls.

In addition to methods for restricting the load to preserve a bounded reset interval, the polling mechanism registers an overload condition if it is unable to perform a reset within the maximum interval. At this point it immediately causes all queue counters to be reset thereby allowing high priority queues to resume service.

Output buffer 35 is responsible for sending a frame alignment slot at regular intervals to enable a distant receiver on link 40 to be synchronised to the outgoing message slot stream. Output buffer 40 also incorporates a register to provide link 40 with

a dummy full slot to be put on the link whenever no slot is received from any of the queues 31, 32 and 33, ie whenever no slot can be selected from the queues and a reset is taking place

Figure 9 shows three rings 11, used for switching purposes, interconnected by links 40. Each ring has a station 10', providing access to links 40. Stations 10' are the same as stations 10 shown in Figure 3 except they also include buffers for queues of slots (31, 32, 33 etc), a transmit buffer 35, and a controller 42 for each link 40 terminating at each station 10'.

Claims

1. A communications multiplexer comprising a plurality of queues (31, 32, 33) for message blocks to be transmitted onto a communication link (40) characterised by control means (42) arranged to provide an allocation of d blocks to each queue, to monitor transmission from the queues (31, 32, 33) onto the link (40), and to inhibit transmission of further blocks from each queue which has transmitted d block onto the link, until a further allocation is made.
2. A multiplexer as claimed in claim 1, characterised in that the control means (42) comprises a counter for each queue (31, 32, 33) for counting the number of blocks transmitted onto the link (40) and the counters are reset when every queue is inhibited or empty.
3. A multiplexer as claimed in claim 2, characterised in that a maximum allowable interval between resets is set according to delay requirements for message blocks for time sensitive services.
4. A multiplexer as claimed in claim 3, characterised in that respective queues are provided for time sensitive services having different maximum delay intervals, and the maximum allowable interval between resets is no greater than the maximum delay interval for the queue having the smallest such interval.
5. A multiplexer as claimed in claim 3 or claim 4, characterised in that prior to accepting a new time-sensitive service call, the number of resets per maximum allowable interval between resets is used to determine whether the call can be accepted without the reset interval rising above the maximum, and on acceptance of the call, increasing the allocation d for its queue.
6. A multiplexer as claimed in any one of claims



1 to 5 characterised in that each queue (31, 32, 33) has an assigned priority and the control means (42) polls the queues (31, 32, 33) in priority order.

7. A multiplexer as claimed in any one of claims 1 to 6 characterised in that a queue (31) is provided for non-time sensitive services, said queue having a small, fixed allocation d.

Patentansprüche

1. Kommunikationsmultiplexer mit einer Vielzahl von Warteschlangen (31, 32, 33) für Nachrichtenblöcke, die auf ein Kommunikationsglied (40) übertragen werden sollen, gekennzeichnet durch eine Steuereinrichtung (42), die angeordnet ist, um eine Zuordnung von d Blöcken zu jeder Warteschlange bereitzustellen, um eine Übertragung von den Warteschlangen (31, 32, 33) auf das Glied (40) zu überwachen und um eine Übertragung weiterer Blöcke von jeder Warteschlange zu hemmen, welche d Blöcke auf das Glied übertragen hat, bis eine weitere Zuordnung durchgeführt wird.
2. Multiplexer nach Anspruch 1, dadurch gekennzeichnet, daß die Steuereinrichtung (42) einen Zähler aufweist für jede Warteschlange (31, 32, 33) zum Zählen der Anzahl von auf das Glied (40) übertragenen Blöcken und daß die Zähler rückgesetzt werden, wenn jede Warteschlange gehemmt oder leer ist.
3. Multiplexer nach Anspruch 2, dadurch gekennzeichnet, daß ein maximales zulässiges Intervall zwischen den Rücksetzungen eingestellt ist gemäß den Verzögerungserfordernissen für Nachrichtenblöcke für zeitempfindliche Betriebsarten.
4. Multiplexer nach Anspruch 3, dadurch gekennzeichnet, daß jeweilige Warteschlangen bereitgestellt sind für zeitempfindliche Betriebsarten mit unterschiedlichen maximalen Verzögerungsintervallen und das maximale zulässige Intervall zwischen Rücksetzungen nicht größer ist als das maximale Verzögerungsintervall für die Warteschlange mit dem kleinsten solchen Intervall.
5. Multiplexer nach Anspruch 3 oder 4, dadurch gekennzeichnet, daß vor dem Annehmen eines neuen zeitempfindlichen Betriebsartenaufruf die Anzahl der Rücksetzungen je maximalem zulässigem Intervall zwischen Rücksetzungen verwendet wird, um zu bestimmen, ob der

Aufruf angenommen werden kann, ohne daß das Rücksetzungsintervall über das Maximum ansteigt und bei der Annahme des Aufrufs die Zuordnung d für seine Warteschlange erhöht wird.

6. Multiplexer nach einem der Ansprüche 1 bis 5, dadurch gekennzeichnet, daß jede Warteschlange (31, 32, 33) eine zugeordnete Priorität hat und die Steuereinrichtung (42) die Warteschlangen (31, 32, 33) der Priorität nach abfragt.
7. Multiplexer nach einem der Ansprüche 1 bis 6, dadurch gekennzeichnet, daß eine Warteschlange (31) bereitgestellt ist für nicht-zeitempfindliche Betriebsarten, wobei die Warteschlange eine kleine feste Zuordnung d hat.

Revendications

1. Un multiplexeur de communications comprenant plusieurs files d'attente (31, 32, 33) pour des blocs de message à transmettre sur une liaison de communication (40) caractérisé par un moyen de commande (42) agencé de façon à fournir une allocation de d blocs à chaque file d'attente, à surveiller une transmission depuis les files d'attente (31, 32, 33) vers la liaison (40) et à empêcher une transmission d'autres blocs à partir de chaque file d'attente qui a transmis d blocs sur la liaison jusqu'à ce qu'une nouvelle allocation soit effectuée.
2. Un multiplexeur selon la revendication 1, caractérisé en ce que le moyen de commande (42) comprend un compteur pour chaque file d'attente (31, 32, 33) pour compter le nombre de blocs transmis sur la file d'attente (40) et que les compteurs sont restaurés lorsque chaque file d'attente est empêchée ou vide.
3. Un multiplexeur selon la revendication 2, caractérisé en ce qu'un intervalle maximal admissible entre des restaurations est défini selon des exigences de retard pour des blocs de message pour des services sensibles au temps.
4. Un multiplexeur selon la revendication 3, caractérisé en ce qu'il est prévu des files d'attente respectives pour des services sensibles au temps dont les intervalles maximaux de temps sont différents, et que l'intervalle maximal admissible entre des restaurations n'est pas supérieur à l'intervalle maximal de retard de la file d'attente dans laquelle cet intervalle est le plus petit.



5. Un multiplexeur selon la revendication 3 ou la revendication 4, caractérisé en ce que le nombre de restaurations par intervalle maximal admissible entre restaurations est utilisé, avant d'accepter un nouvel appel de services sensibles au temps, afin de déterminer si l'appel peut être accepté sans que l'intervalle de restauration ne dépasse le maximum et pour augmenter, lors de l'acceptation de l'appel, l'allocation d pour sa file d'attente. 5 10
6. Un multiplexeur selon l'une quelconque des revendications 1 à 5, caractérisé en ce qu'une priorité est assignée à chaque file d'attente (31, 32, 33) et le moyen de commande (42) interroge les files d'attente (31, 32, 33) dans l'ordre de priorité. 15
7. Un multiplexeur selon l'une quelconque des revendications 1 à 6 caractérisé en ce qu'il est prévu une file d'attente (31) pour des services non sensibles au temps, ladite file d'attente possédant une petite allocation fixe d. 20

25

30

35

40

45

50

55