



⑪ Publication number : **0 666 666 A2**

⑫ **EUROPEAN PATENT APPLICATION**

⑳ Application number : **95480005.8**

⑤① Int. Cl.<sup>6</sup> : **H04L 12/24, H04L 12/42, H04L 29/00**

㉔ Date of filing : **24.01.95**

③① Priority : **07.02.94 US 192884**

④③ Date of publication of application :  
**09.08.95 Bulletin 95/32**

⑥④ Designated Contracting States :  
**DE FR GB**

⑦① Applicant : **International Business Machines Corporation**  
**Old Orchard Road**  
**Armonk, N.Y. 10504 (US)**


⑦② Inventor : **Cidon, Israel**  
**2600 Netherland Avenue, Apt. 2507**  
**Bronx, New York 10463 (US)**

Inventor : **Georgiadis, Leonidas**  
**36 Old Mill Road**  
**Chappaqua, New York 10514 (US)**  
Inventor : **Guerin, Roch Andre**  
**Scenic View No.4H,**  
**Rochambeau Drive**  
**Yorktown Heights, New York 10598 (US)**  
Inventor : **Shavitt, Yuval Yitzchak**  
**Meonot Segel Zuter 5/13**  
**Technion City, Haifa 32000 (IL)**  
Inventor : **Slater, Andrew Emlyn**  
**2262 Parker Lane**  
**Yorktown Heights, N.Y. 10598 (US)**

⑦④ Representative : **Therias, Philippe**  
**Compagnie IBM FRANCE,**  
**Département de Propriété Intellectuelle**  
**F-06610 La Gaude (FR)**

⑤④ **Method and apparatus for improved throughput in a multi-node communication system with a shared resource.**

⑤⑦ A non-quota access indicator is circulated among nodes in a multi-node quota based communication system with a shared resource, indicating maximum possible non-quota access to the shared resource to a given node receiving same. Upon arrival at a node, the indicator is saved and then updated to reflect the current status of that node as either starved or satisfied, the former being a condition of currently having quota remaining and a shared resource access requirement, and the latter being a condition of either currently having no remaining quota or having no current shared resource access requirement. After updating, the node immediately propagates the indicator to the next node in the system. When a node without quota requires access to the shared resource, it compares its requirement to the last stored indicator and accesses the shared resource if the stored indicator is equal to or greater than the access requirement.



*fig. 4*

## **BACKGROUND OF THE INVENTION**

### **Technical Field**

The present invention generally relates to multi-node communication systems with shared resources. More particularly, the present invention relates to improving throughput in multi-node communication systems with shared resources and employing a shared resource access quota scheme.

### **Background Art**

In the past, multi-node communication systems, such as local area networks, including both ring networks and bus networks, have managed node access to shared resources by limiting access thereto via mechanisms such as polling, random access, time slot and quota allocation mechanisms. The mechanisms proposed have generally attempted to provide fair access to the shared resources by the nodes. Some quota schemes have equalized the node throughput by allocating the same quota to each node.

However, allocating the same quota to each node does not achieve optimum efficiency with regard to system throughput. The task of determining and tracking the allocation of different quotas among nodes is difficult and requires a large amount of information exchange, especially in systems where loading conditions fluctuate or loading is asymmetrical. Thus, most quota schemes have centered on the allocation of the same quota to all nodes. One such quota scheme is described in U.S. Patent No. 4,926,418 issued to Cidon et al. and assigned to IBM (hereinafter, "the Cidon patent"), the disclosure of which is herein incorporated by reference in its entirety.

Thus, a need exists for improved efficiency in multi-node communication systems with shared resources employing a fixed quota scheme in order to allow for improved throughput in the case of an asymmetrically loaded system or a system whose load fluctuates.

## **SUMMARY OF THE INVENTION**

Briefly, the present invention satisfies the need for improved efficiency in multi-node communication systems by providing a method and apparatus for allowing nodes currently without access quota to access a shared resource so long as access by nodes with quota remaining is not affected.

The basic method requires a communication system having a plurality of nodes and a shared resource, and which employs a shared resource access quota scheme for node access to the shared resource. A non-quota access indicator is circulated

among the nodes. The non-quota access indicator indicates a maximum possible non-quota access by a given node, currently without quota, without affecting quota access by other nodes in the system. In response to receipt by the given node, the non-quota access indicator is updated to reflect a current status of the given node as either starved or satisfied. A status of starved is defined as having quota and an access requirement for the shared resource. A status of satisfied is defined as either having no quota remaining or no access requirement for the shared resource.

Where the shared resource is a bidirectional communication path, the non-quota access indicator may be circulated in a direction opposite that of system information traffic. Where the shared resource is a unidirectional communication path, an access requirement status indicator is circulated among the nodes. The status indicator indicates a current status for each node in the system as either currently starved or satisfied.

The present invention also includes a method for a given node in the system to indicate to other nodes a status change of the given node which affects possible non-quota access by the other nodes to the shared resource. The status change is between a status of starved and a status of satisfied or vice-versa. When a status change of the given node is detected, a non-quota access update indicator is issued therefrom and propagated.

The present invention also includes a communication system implementing the above principles.

These, and other objects, features and advantages of this invention will become apparent from the following detailed description of the various aspects of the invention taken in conjunction with the accompanying drawings.

## **BRIEF DESCRIPTION OF THE DRAWINGS**

FIG. 1 depicts a five-node communication system according to the present invention with a bidirectional ring architecture.

FIG. 2 depicts a five-node communication system according to the present invention with a unidirectional ring architecture.

FIG. 3 depicts a portion of a node in FIG. 1 where a buffer insertion ring architecture is used.

FIG. 4 depicts, in block diagram form, circuitry within the node of FIG. 3 for implementing the present invention.

## **BEST MODE FOR CARRYING OUT THE INVENTION**

FIG. 1 depicts a five-node bidirectional communication ring 10 including nodes 12, 14, 16, 18 and 20. Communication ring 10 comprises two distinct rings—clockwise ring 11 and counterclockwise ring 13. Be-

tween any two consecutive nodes are two communication links, for example, links 22 and 24 between nodes 14 and 16. Link 22 provides a path for information to travel from node 16 to node 14, and link 24 provides a path for information to travel from node 14 to node 16. For each ring there is a set of control signals, a quota allocation signal and an INFO signal (to be described subsequently), rotating on the other ring. In addition, a given node with information to transmit to another node will choose the shortest path provided on either ring for the transmission. A given quota allocation signal therefore travels in the opposite direction of the information traffic it regulates. A given quota allocation signal, when received at a given node, allocates quota to that node, i.e., it allows that node to transmit a certain amount of information on the relevant ring. The following example will focus only on a quota allocation signal and INFO signal circulating on counterclockwise ring 13 regulating information traffic on ring 11, the description for the other set of signals being similar.

Assume, for exemplary purposes, that nodes 20 and 12 each currently have information to transmit to node 14. Assume also that node 16 currently has no quota, but has information with a destination of node 20. Under the quota allocation scheme of the Cidon patent, node 16 must wait until its quota is refreshed via the quota allocation signal to transmit to node 20.

The present invention improves throughput on the ring by allowing node 16, although currently without quota, to transmit the information to node 20 prior to the next visit of the quota allocation signal if quota transmissions by the other nodes are not affected. This non-quota access can only be utilized by a node after its quota has been exhausted. A first aspect of the present invention provides a second signal, the INFO signal, traveling on counterclockwise ring 13. The INFO signal contains a HOPCOUNTER, indicating to a given node the current number of downstream nodes plus one that either do not currently have quota remaining or information to transmit (i.e., are currently satisfied). The term "downstream" is used in the sense of information flow. In other words, the HOPCOUNTER informs a given node how far ahead on ring 10 it may currently transmit information (or how many nodes away it can "hop" to) without quota. In general, a given node with quota remaining and information to transmit is referred to as being "starved", and a node currently either without information to transmit or having no quota remaining is referred to as "satisfied".

As described in the Cidon patent, benefits are realized by propagating the quota allocation signal in a direction opposite the information traffic it regulates. The INFO signal also preferably travels opposite the information traffic and separately from the quota allocation signal, in order to propagate it quickly through ring 10. Under heavy loading conditions, the INFO

signal traveling separately from the quota allocation signal may circulate several times around ring 10 in the time it takes the quota allocation signal to make one round trip. Thus, system throughput may be maximized if both the INFO signal and the quota allocation signal travel separately and in a direction opposite the information traffic regulated thereby.

When a node receives the INFO signal, it immediately stores the HOPCOUNTER in local memory (see FIG. 4). The node stores the HOPCOUNTER so that whenever it has information to transmit and is out of quota, it checks to see if the destination for the information is within the stored HOPCOUNTER distance. If a node is satisfied, the node increments the HOPCOUNTER by one when received and immediately propagates it to the next upstream node. When a starved node receives the INFO signal, the node resets the HOPCOUNTER to one and propagates the INFO signal to the next upstream node. In the present aspect, a node always assumes that an upstream neighbor may transmit to it, hence, resetting the HOPCOUNTER to one when starved. A maximum value for the HOPCOUNTER above which it is not incremented may be chosen, as a function of the number of nodes in the ring and the number of bits in length the HOPCOUNTER is.

Assume now that node 20 is satisfied and node 18 is satisfied with information destined for node 12. Assume also that node 20 has just received the INFO signal from node 12 over link 28. Node 20 stores the HOPCOUNTER (assume it has a value of one), increments it to a value of two since it is satisfied and propagates the INFO signal to node 18 over link 30. When node 18 receives the INFO signal, it immediately stores the HOPCOUNTER value, increments it by one since it is satisfied and propagates the INFO signal to node 16 over link 32. Node 18 may now transmit without quota to node 12 over links 26 and 34, since the destination of its information, node 12, is two nodes away and its most recently stored HOPCOUNTER has a value of two, i.e., the destination node is equal to or less than the most recently stored HOPCOUNTER value.

In a second aspect of the present invention, a given node generates an INFO signal whenever its status changes from satisfied to starved or vice-versa. The benefits of such an INFO signal "on demand" scheme, instead of only circulating a single INFO signal, are illustrated by the following two examples. First, consider a satisfied node having quota left, but no information to transmit. When the node gets an INFO signal, it immediately stores the HOPCOUNTER, increments it (since it is satisfied) and sends the INFO signal on. Just after propagating the INFO signal, the node receives information to transmit. The node's status has changed from satisfied to starved. Thus, the HOPCOUNTER just sent by that node is inaccurate and will remain so, with respect to that node,

until the INFO signal next arrives there. An on demand scheme would allow the node to issue another correct INFO signal immediately, rather than waiting to correct the first one.

Second, consider a starved node. The node receives the INFO signal, stores the received HOPCOUNTER, resets the HOPCOUNTER to one (since it has information to transmit and quota remaining) and propagates the INFO signal. Shortly after propagating the INFO signal, the node finishes transmitting its information or runs out of quota. The status of the node has changed from starved to satisfied. Thus, non-quota transmissions which might otherwise be allowed upstream will not take place until the INFO signal is corrected when it next arrives at the node.

In this second aspect of the present invention, the nodes in the two examples above would each issue a new INFO signal in response to its status change. A given node receiving an INFO signal would check to see if the received HOPCOUNTER is different from or the same as the most recently stored HOPCOUNTER value. If the received and stored HOPCOUNTERs are the same, the node does not update the HOPCOUNTER nor propagate the INFO signal. If the compared HOPCOUNTERs are different, the node stores the received HOPCOUNTER. Then, if the node is satisfied, it updates the HOPCOUNTER (i.e., increments by one) and propagates the INFO signal. If the node is currently starved, it stores the received HOPCOUNTER but does not propagate the INFO signal. In addition, in the case of a bidirectional system, if a node receives an INFO signal with a HOPCOUNTER having a value equal to the number of nodes in the system, it does not update the HOPCOUNTER or propagate the INFO signal, since such a value implies that all nodes are currently satisfied and all nodes are aware that all nodes are satisfied. The same is true in a unidirectional system where the received HOPCOUNTER has a value equal to twice the number of nodes in the system.

The above two aspects of the present invention assume a bidirectional communication ring. However, some communication systems are unidirectional. FIG. 2 depicts a five-node unidirectional communication ring 36 with nodes 38, 40, 42, 44 and 46. In a third aspect of the invention, an INFO signal is circulated in the direction of information traffic and contains a bit map with a bit for the status of each node. A satisfied node has a status of one and a starved node has a status of zero. When a given node, e.g., node 40, receives the INFO signal, it determines from the bit map how many consecutive nodes ahead of it have a bit value of one (a HOP indicator) and stores this information in local memory. The node then updates its own bit on the bit map and propagates the INFO signal. A given node may transmit information if its last stored HOP indicator value is equal to or greater than the number of hops it needs to make to reach its des-

tinuation node.

In a fourth aspect of the present invention, apparatus implementing the above principles is presented. The fourth aspect of the present invention will be described with reference to ring 10 of FIG. 1 further described as a buffer insertion ring architecture, as it is known in the art. FIG. 3 depicts an input/output portion 48 of node 14. Insertion buffer 50 can store at most one maximal size packet of information. Node 14 may transmit information packets from transmit buffer 52 at any time, provided insertion buffer 50 is empty and no other nodes are transmitting over link 24 (ring traffic is given priority). If ring traffic is arriving on link 23 while node 14 is transmitting an information packet, the incoming packet is stored in insertion buffer 50 until the transmission is completed. If information destined for node 14 is received on link 23, it enters receive buffer 54. When node 14 has no information packets to transmit, ring traffic is not delayed in insertion buffer 50.

FIG. 4 depicts INFO-subcircuit 56, in block diagram form, within node 14. INFO-subcircuit 56 comprises hopcounter register 58, hop register 60, comparator 62 and requestor 64. Hopcounter register 58 can be loaded, incremented, reset and read from. Hop register 60 can be loaded and read from. Comparator 62 compares two values and asserts a signal when certain conditions are met. Requestor 64 controls INFO-subcircuit 56. The operation of INFO-subcircuit 56 will now be described in detail.

When node 14 receives an INFO signal on link 22, hopcounter register 58 and hop register 60 are immediately loaded with the HOPCOUNTER. Requestor 64 responds to the INFO signal, active upon the loading of hopcounter register 58, by incrementing hopcounter register 58 if the SATISFIED signal is asserted. An assertion of the SATISFIED signal indicates that node 14 has exhausted its quota or transmit buffer 52 is empty. If the SATISFIED signal is not asserted, requestor 64 resets hopcounter register 58 to one. Immediately after the increment or reset operation, the INFO signal (with the HOPCOUNTER from hopcounter register 58) is sent on to node 12 over link 25.

In addition to updating the HOPCOUNTER, INFO-subcircuit 56 determines if a non-quota transmission by node 14 is possible. Comparator 62 compares the value in hop register 60 with a DestHops signal to determine if DestHops is less than or equal to the stored HOPCOUNTER in hop register 60. DestHops indicates the number of downstream "hops" the next information packet from transmit buffer 52 must make in order to arrive at its destination. If DestHops is less than or equal to the stored HOPCOUNTER, comparator 62 asserts a HopsOK signal, indicating to requestor 64 that the currently desired non-quota transmission is possible.

## Claims

1. In a communication system having a plurality of nodes and a shared resource and employing a shared resource access quota scheme for node access to said shared resource, a method for non-quota access to said shared resource by a given node without quota, said method comprising the steps of:  
 circulating a non-quota access indicator among said plurality of nodes, said non-quota access indicator indicating a maximum possible non-quota access by said given node to said shared resource without affecting quota access thereto by one or more of the other of said plurality of nodes; and  
 updating said non-quota access indicator in response to receipt thereof by said given node to reflect a current status of said given node as either starved or satisfied, said status of starved being defined as having quota and an access requirement for said shared resource, and said status of satisfied being defined as either having no quota remaining or no access requirement for said shared resource.
2. The method of claim 1, wherein said shared resource comprises a communication path, and wherein said non-quota access indicator indicates a maximum distance on said communication path available for said given node to transmit information.
3. The method of claim 2 wherein said non-quota access indicator indicates a number of consecutive nodes in said communication path directly ahead of said given node currently having a status of satisfied.
4. The method of claim 3, wherein said non-quota access indicator comprises a node counter, and wherein said step of updating comprises incrementing said node counter to reflect that said given node is currently satisfied.
5. The method of claim 3, wherein said non-quota access indicator comprises a node counter, and wherein said step of updating comprises setting said node counter to a predetermined number to reflect that said given node is currently starved.
6. The method of claim 3, wherein said communication path is bidirectional, and wherein said step of updating comprises updating said non-quota access indicator to reflect said current status of said given node unless said number of consecutive nodes equals a total number of said plurality of nodes on said bidirectional communication path.
7. The method of claim 3, wherein said communication path is unidirectional, and wherein said step of updating comprises updating said non-quota access indicator to reflect said current status of said given node unless said number of consecutive nodes equals twice a total number of said plurality of nodes on said unidirectional communication path.
8. The method of claim 1 further comprising the step of allowing said given node to access said shared resource in accordance with said non-quota access indicator.
9. The method of claim 8, wherein said step of allowing comprises the steps of:  
 comparing said non-quota access indicator to a current access requirement indicator of said given node, said current access requirement indicator indicating an extent of access to said shared resource currently required by said given node; and  
 determining if said given node may access said shared resource without quota based on said comparison.
10. The method of claim 9 wherein said step of determining comprises determining if said non-quota access indicator is equal to or greater than said current access requirement indicator, wherein said given node is allowed to access said shared resource only if said non-quota access indicator is equal to or greater than said current access requirement indicator.
11. In a communication system having a plurality of nodes and a shared communication path and employing a shared resource access quota scheme for node access to said shared communication path by said plurality of nodes, a method for a given node to indicate to the other of said plurality of nodes a status change thereof affecting possible non-quota access thereby to said shared communication path, said method comprising the steps of:  
 monitoring for a status change of said given node between a status of starved and a status of satisfied;  
 issuing a non-quota access update indicator from said given node in response to detecting said status change; and  
 propagating said issued non-quota access update indicator.
12. The method of claim 11, wherein said step of issuing comprises:  
 incrementing and issuing a previously stored non-quota access indicator if said status change

is from said status of starved to said status of satisfied; and  
 resetting and issuing said previously stored non-quota access indicator to a predetermined value if said status change is from said status of satisfied to said status of starved.

5

comprising a means for comparing said non-quota access indicator to a current access requirement indicator of said given node to determine if a current shared resource access requirement thereof is possible.

- 13.** The method of claim 11 or 12, wherein said step of propagating comprises the steps of:

comparing at another node said propagated non-quota access update indicator and a stored non-quota access indicator to determine if a difference therebetween exists;

10

storing at said another node said issued non-quota access indicator in place of said stored non-quota access indicator if said difference is determined; and

15

updating at said another node said propagated non-quota access indicator to reflect a current status of said another node and propagating same if said current status of said another node is satisfied.

20

- 14.** The method of claim 13 wherein said step of updating comprises:

25

incrementing said issued non-quota access indicator by one if said current status of said second node is said status of satisfied; and

resetting said issued non-quota access indicator to one if said current status of said second node is said status of starved.

30

- 15.** A communication system, comprising:

a plurality of nodes;

a shared resource, wherein said communication system employs a shared resource access quota scheme for node access to said shared resource; means for circulating a non-quota access indicator among said plurality of nodes, said non-quota access indicator indicating a maximum possible non-quota access by a given node to said shared resource without affecting quota access to said shared resource by one or more of the other of said plurality of nodes; and

35

means for updating said non-quota access indicator in response to receipt thereof by said given node to reflect a current status of said given node as either starved or satisfied.

40

45

- 16.** The communication system of claim 15, wherein said shared resource comprises a shared communication path, and wherein said non-quota access indicator indicates a consecutive number of nodes directly ahead of said given node on said shared communication path currently having a status of satisfied.

50

55

- 17.** The communication system of claim 15 further




fig. 1




fig. 2




fig. 3

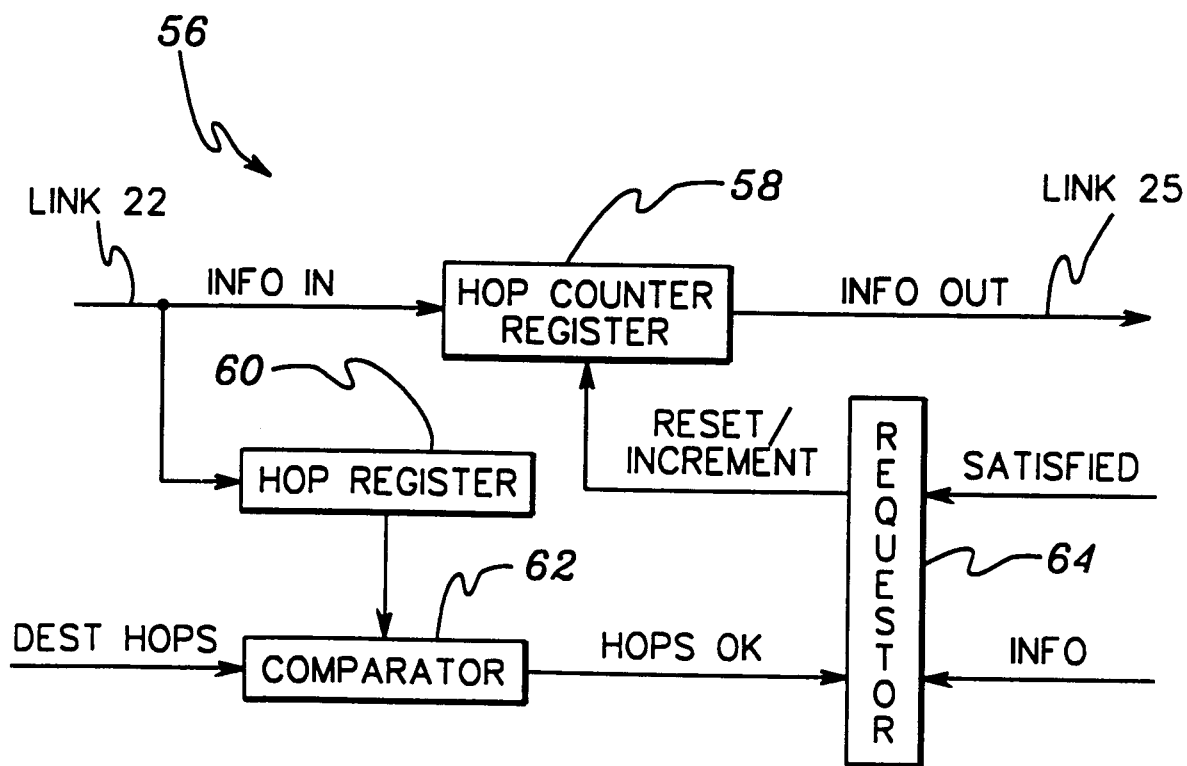


fig. 4