An Overview of IEEE 802.12 Demand Priority

Alan Albrecht*, Joe Curico*, Dan Dove*
Steve Goody*, Michael Spratt
Networks and Communications Laboratory
HP Laboratories Bristol
HPL-94-17
February, 1994

100VG-Any LAN, 100Base-VG, demand priority, Local Area Network, (LAN), high speed network, IEEE 802.12

100VG-Any LAN is a Local Area Network developed and standardized by the IEEE 802.12 committee. Demand Priority, the Media Access Control of IEEE 802.12, is a 100 Mb/s hub to end node protocol, but extends to allow multiple hubs to form a single network. Demand Priority supports 2 priority levels, enabling it to support the emerging multimedia applications. This paper describes Demand Priority, and how it combines with the physical layers to support a wide range of media types, particularly all IEEE 802.3 10Base-T compliant cabling. The paper also positions Demand Priority relative to FDDI, ATM and higher speed CSMA/CD, and outlines potential advantages Demand Priority may have relative to these technologies.

Internal Accession Date Only
* HP Roseville Networks Division, Roseville, CA

1 Introduction

100VG-AnyLAN is a Local Area Network technology currently being developed and standardised by those participating in the IEEE 802.12 committee. Demand Priority, the Media Access Control (MAC) of IEEE 802.12, is a protocol between a hub and connected end nodes [1], [2]. Demand Priority also extends to cascading[1] across hubs. Key characteristics of the network are:

- 100 Mb/s network data rate.
- Support of Category 3 Unshielded Twisted Pair (UTP) cables, including bundled cables, Category 5 UTP cabling, Shielded Twisted Pair and optical fibre.
- Priorities to support real time multimedia and other delay-critical traffic. Priorities can be utilised to give guarantees of bandwidth and delay to streams of multimedia data sent from users' end nodes.
- Cascading beyond the immediate wiring closet.
- Compatibility with IEEE 802.3 and IEEE 802.5 frame format, allowing transparent upgrading of current widely installed LANs.
- Low cost and low end node system complexity.
- Privacy as a security feature.

This paper describes the operation of the Demand Priority protocol, and positions it relative to other technologies.

2 Motivation behind Developing Demand Priority

In this section we summarise the motivation behind the development of Demand Priority, partly setting the context with reference to other LAN protocols.

2.1 Support of Category 3 and Bundled Cables

The large majority of the installed cabling base is star-wired Category 3 cable. Cables run from wiring closets to end nodes at the users' desks, with hubs and cross connects located in the wiring closets. Some of the cabling is 4-pair all the way from the wiring closet to the desk, while other installations have 25-pair or 50-pair bundles on part of the path.

1. Allowing multiple hubs to be connected in a single LAN.

A clear user requirement for 100 Mb/s solutions has become apparent, but ring technologies (such as FDDI) appear unable to support operation over Category 3 cable, in part because of the full duplex nature of ring operation. See Figure 1.

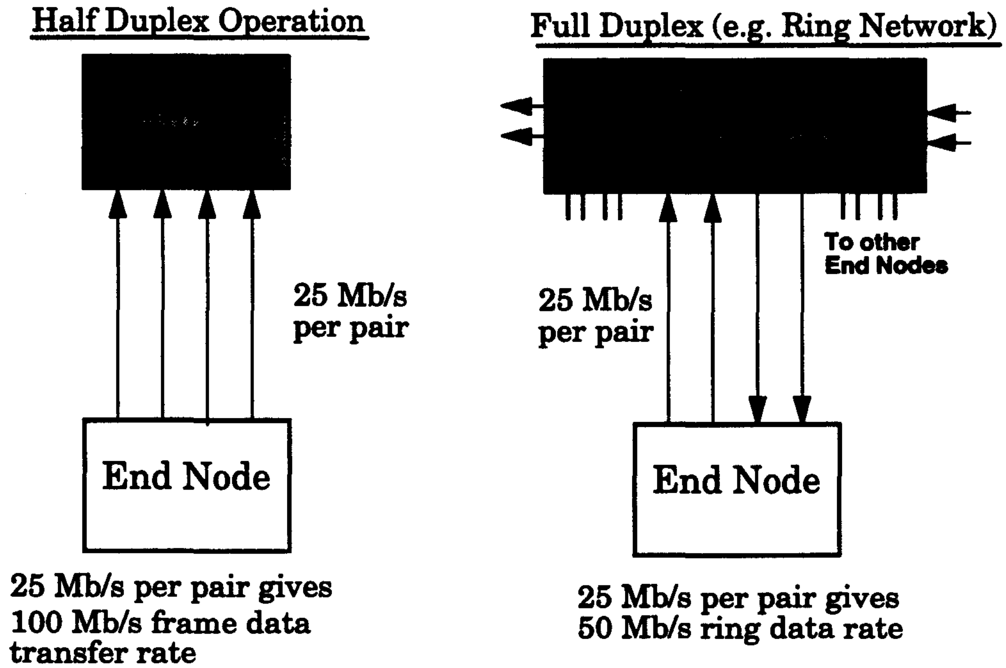

Figure 1 The advantage of Half Duplex Link Operation as opposed to Full Duplex Link Operation. Assumes 25 Mb/s transmission on each Category 3 cable physical layer.

A key idea in Demand Priority is to use Half Duplex operation over 4-pair Category 3 cable, thus achieving 100 Mb/s transmission with a data rate of just 25 Mb/s per pair. By contrast, a ring LAN protocol with a data rate per pair of 25 Mb/s allows a ring speed and therefore a frame data transfer rate of just 50 Mb/s.

Bundled cables are in widespread use in Category 3 cabling installations. In Demand Priority, Near End Cross Talk (NEXT) is minimised by (a) sending unicast frames only to their intended destination, and (b) in the case of hubs connecting to bundles, fully receiving broadcast frames before sending them on to the destination end nodes. Therefore, frame data arriving at the hub in a bundle experiences NEXT from at most just 4 disturbing pairs. In this way, Demand Priority combines with the Physical Layer to support bundled cables, whereas in CSMA/CD and ring networks, data arriving at the hub will experience NEXT from too many disturbers for successful operation.

2.2 Minimising End Node Adaptor Card System Complexity

One reason why ring solutions are relatively expensive, is that they are generally implemented with a processor on the end node adaptor card to carry out the functions particularly associated with rings - reconfiguration, ring recovery etc. (In the case of FDDI, these functions are defined in SMT, the Station Management protocol.)

Demand Priority allows a simpler implementation of end nodes. A processor per end node adaptor card is *not* required, minimising implementation cost.

2.3 Two Priority Levels

Demand Priority supports 2 priority levels, which can be used to support real time multimedia transmission (e.g. video and audio) in the following way. The real time multimedia data is sent at the High priority, while less time critical computer data is sent at Normal (the lower) priority. The High priority multimedia traffic will take precedence over the computer data traffic, and if it is ensured that the total High priority traffic does not exceed say 90% of the available bandwidth minus overheads, the High priority traffic will experience very low delays. One mechanism to do this is the Target Transmission Time algorithm, described in [3], [4] and [5].

In fact, the network delay and bandwidth allocated to a High priority multimedia application can both be *guaranteed*, so Demand Priority will be able to support a much wider range of the emerging multimedia applications than a non-deterministic single priority network such as 10 Mb/s or 100 Mb/s CSMA/CD. A Demand Priority network will be able to support the in-building portion of a high quality video[1] call across the state or country, owing to its capacity to guarantee bandwidth and delay irrespective of the computer data load.

2.4 Network Diameter

Another drawback of speeding up CSMA/CD to 100 Mb/s is that the network diameter (the furthest distance between two end nodes, via cabling and intermediate hubs) would be restricted[2]. Hub to hub connection in different wiring closets would have to use a bridge or router i.e. there would be no cascading outside of a single wiring closet. The 100 Mb/s MAC protocol would also be restricted to short hub to end node connections.

This is unattractive, given the cost and limited throughput of bridges and routers.

By contrast, Demand Priority allows hub to end node links to be up to 2 km (in the case of optical fibre links), and hubs can be cascaded to a diameter of several kilometres. Thus Demand Priority networks can use all of the current 10Base-T compliant cabling.

1. With MPEG coding, good quality video can be sent with a data rate in the 0.5 to 1 Mb/s. The crucial requirement for video and audio coding schemes is the guarantee of the bandwidth and bounded delay across the networks it transverses.

2. In speeding up 10 Mb/s by a factor of 10, the diameter expressed in terms of propagation delay has to be reduced by a factor of 10.

Finally, the table below summarises some of the differences between the respective LAN technologies.

	Supports Multimedia with guaranteed delay and bandwidth	End Node adaptor card complexity	Half duplex links enable Category 3 cable support	Bundled cables supported	Connects multiple hubs outside single wiring closet without bridging / routing	Upgrades users from current ethernet / IEEE 802.3 without software changes
FDDI	Yes	Station Management expensive	No	No	Yes	No
CSMA/CD @ 100 Mb/s	No	Low	Yes	No	No	Yes
ATM LAN	Yes	Segmentation, reassembly of frames required	No	No	Yes	No
Demand Priority	Yes	Low	Yes	Yes	Yes	Yes

ATM LANs are generally considered to be switched technologies. Interestingly enough, Demand Priority can also potentially encompass switch technologies, with end nodes connecting across a link operating the Demand Priority Link Protocol (described in Section 3.1) to the switch.

A potential advantage of a Demand Priority based switch is that it can avoid congestion by simply not responding to any Requests from end nodes, while its internal buffers are (nearly) full. ATM switches have no standardised flow control over connected end nodes.

A further congestion control issue for ATM is the segmenting of frames into slots. If 1500 bytes of data is lost in a switch through congestion, this could correspond to around 30 slots being lost, and if the cells are part of different frames, could result in the loss and consequent re-transmission by senders of 30 frames. By contrast, the loss of 1500 bytes of data in a frame-based switch running the Demand Priority Link Protocol would result in the retransmission of far fewer (possibly just 1) frames.

3 A Description of Demand Priority.

3.1 The Demand Priority Link Protocol

Demand Priority uses a small number of control signals, four control signals sent from the Hub to the End Node, and four sent in the opposite direction.

In a 4-pair UTP link, the control signals are implemented as low frequency tones, which do not significantly contribute to cross-talk. In both directions, each of two pairs transmits one of two possible low frequency tones, giving four control signals in each direction.

The link has 3 states plus transitions between these states :

1. The Hub is sending a frame at 100 Mb/s to the End Node
2. The End Node is sending a frame at 100 Mb/s to the Hub
3. The Hub and End Node are exchanging control signals.

This is shown below for half duplex links.

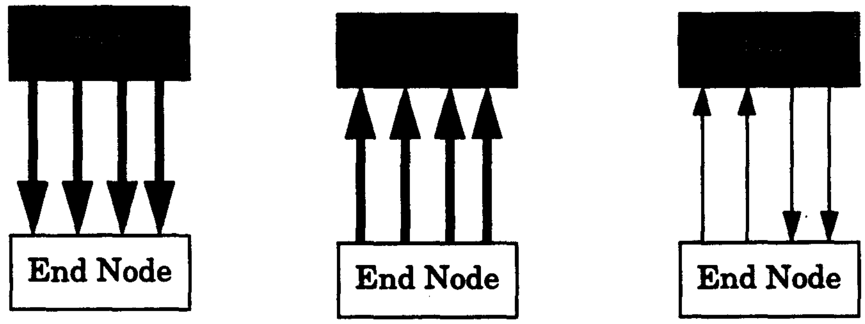

1. 25 Mb/s data on 4 pairs
2. 25 Mb/s data on 4 pairs
3. Control signal tones in both directions

Figure 2 The Link Protocol States in Half Duplex Operation

For reference, the 3 link states are shown for a dual simplex 100 Mb/s optical fibre link.

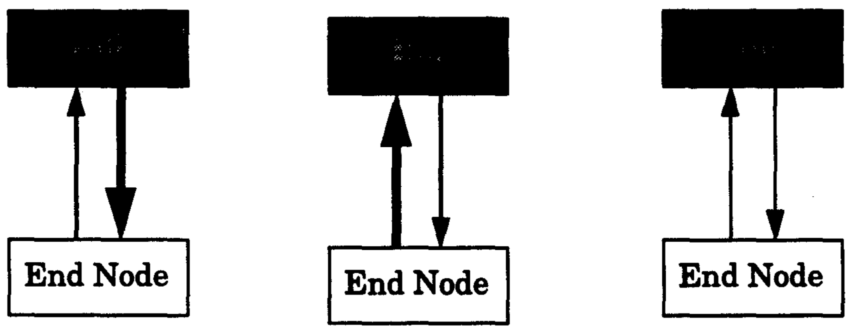

1. 100 Mb/s data from Hub to End Node
2. 100 Mb/s data from End Node to Hub
3. Control signals in both directions

Figure 3 The Link Protocol States in Dual Simplex Operation

We now step through the Demand Priority Link Protocol on half duplex links, the dual simplex case being essentially very similar.

Figure 4 shows how an End Node is selected by a Hub, and how it sends a frame to the Hub.

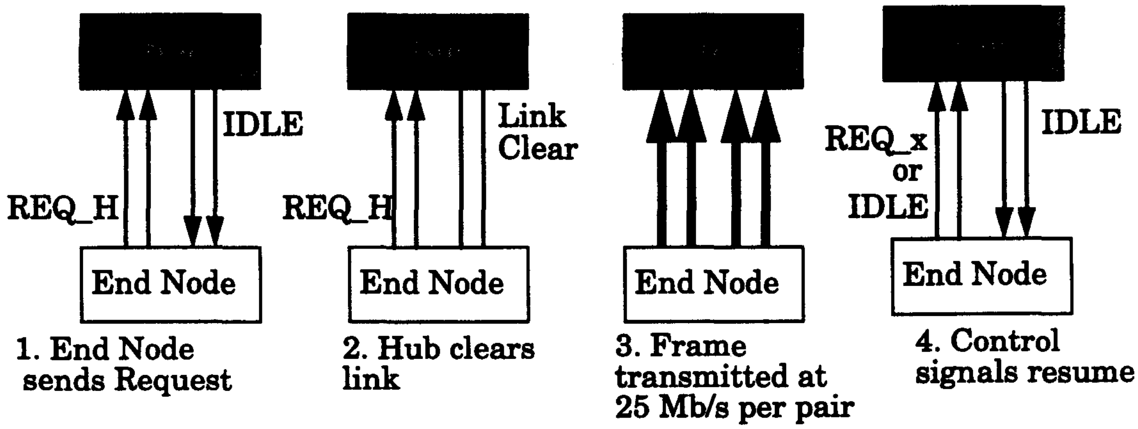

Figure 4 Link Protocol Sequence of the End Node transmitting a Frame

(1) When an End Node has a frame to send, it sends up a control signal REQ_H (High Priority Request) or REQ_N (Normal Priority Request) to the Hub.

(2) When the Hub decides to select this particular End Node, it stops actively sending the control signal on 2 pairs, in effect it *clears the link* to the End Node.

(3) When the End Node senses that there is no signal on the link, it realises that it has been selected to transmit by the Hub. It can transmit immediately on 4 pairs, because the Hub is no longer transmitting actively on any of the pairs.

(4) When the frame is finished, the Hub will send IDLE, and the End Node will send REQ_H or REQ_N (depending on the priority level of further frames to be transmitted) or IDLE (if there are no such frames).

We now look at the Link Protocol when the End Node is receiving a frame from the Hub.

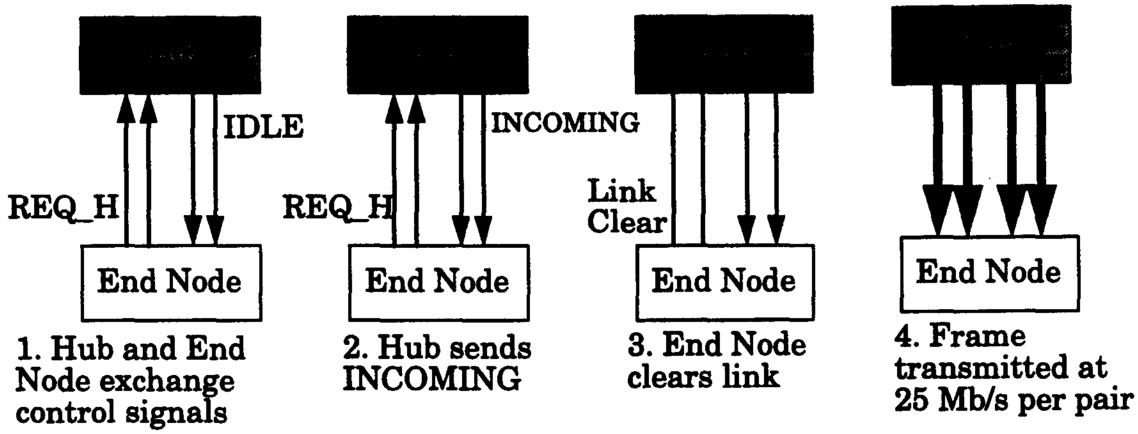

Figure 5 Link Protocol Sequence of the End Node receiving a Frame

(1) The Hub and End Node are exchanging control signals.

(2) The Hub is expecting to receive a frame from some other End Node. It sends INCOMING to the End Node in Figure 5.

(3) A short time later[1], the End Node clears the link.

(4) The link is now clear, so the Hub is able to start sending the frame on all 4 pairs, giving a 100 Mb/s frame data rate.

3.2 The Operation of Demand Priority on a Single Hub

We now see how links operating the Link Protocol combine to create the Demand Priority shared media access protocol.

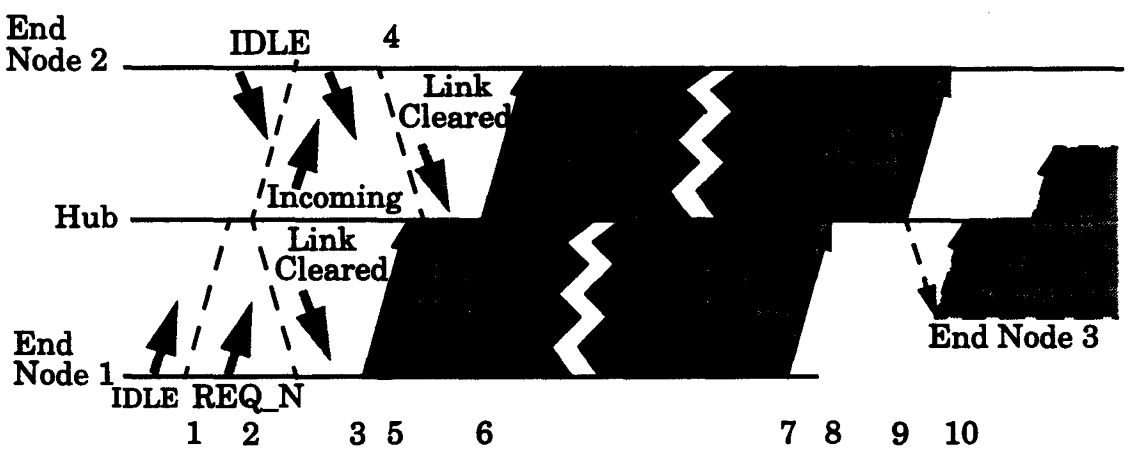

Figure 6 The Operation of Demand Priority on a Single Hub

1. The End Node ensures that it has sends the control signal (in this case REQ_H) long enough for it to be recognised at the Hub.

Initially, none of the End Nodes require to send any frames, so they are sending IDLE to the Hub. We now step through the sequence of events which illustrate the Demand Priority single hub operation.

(1) End Node 1 now has a frame to transmit across the network, so it starts sending REQ_N to the Hub.

(2) The Hub acquires the Request, and clears the link to End Node 1. Clearing the link will enable End Node 1 to transmit on all 4 pairs to the Hub. Simultaneously, the Hub sends INCOMING to all of the other connected End Nodes, including End Node 2.

(3) End Node 1 senses that the Hub has cleared the link, and starts transmitting the frame on all 4 pairs.

(4) End Node 2 has acquired INCOMING, and clears its link to the Hub, thus allowing the Hub to transmit on all 4 pairs to it.

(5) The Hub receives the start of the frame from End Node 1, and starts to store it in an internal buffer. It is waiting for the 48-bit Destination Address to arrive and a look-up to take place on the Destination Address

(6) The look-up on the Destination Address has now occurred, indicating that the frame should be sent to End Node 2. End Node 2 has already cleared the link to the Hub, and so the Hub can now transmit on all 4 pairs to End Node 2. Assuming the frame arriving from End Node 1 is unicast, and only needs to be sent on to End Node 2, the Hub will at this point replace INCOMING with IDLE to the other End Nodes. This will allow these End Nodes to transmit control signals on 2 pairs, and in particular send a REQ_H or REQ_N to the Hub, effectively informing it of their readiness to transmit a frame when the current one is finished.

(7) End Node 1 finally finishes sending its frame to the Hub.

(8) The end of frame arrives at the Hub. The Hub now starts sending Idle to End Node 1.

(9) The Hub now sends the end of the frame to End Node 2. If the Hub is receiving no Requests from any End Node, then it will send Idle to End Node 2. The network has now returned to the quiescent state.

Suppose instead that another End Node, for example End Node 3, has been sending REQ_H or REQ_N to the Hub. The resulting sequence is shown with dotted lines. At the same time as sending the of Frame to End Node 2, the Hub clears the link to End Node 3, in a similar way to (2). The Hub now sends INCOMING to all of the other End Nodes, including End Node 2 and End Node 1, as all of the End Nodes are at this point potential destinations of the new frame.

(10) End Node 3 senses that the Hub has cleared the link, and starts transmitting its frame. (c.f. action (3).)

The network now continues exactly as with the previous frame transferred, carrying out the equivalent of action (4) and the subsequent actions.

3.3 The Demand Priority MAC Support of various Media Types

Figure 7 shows diagrammatically the way in which Demand Priority supports different media types.

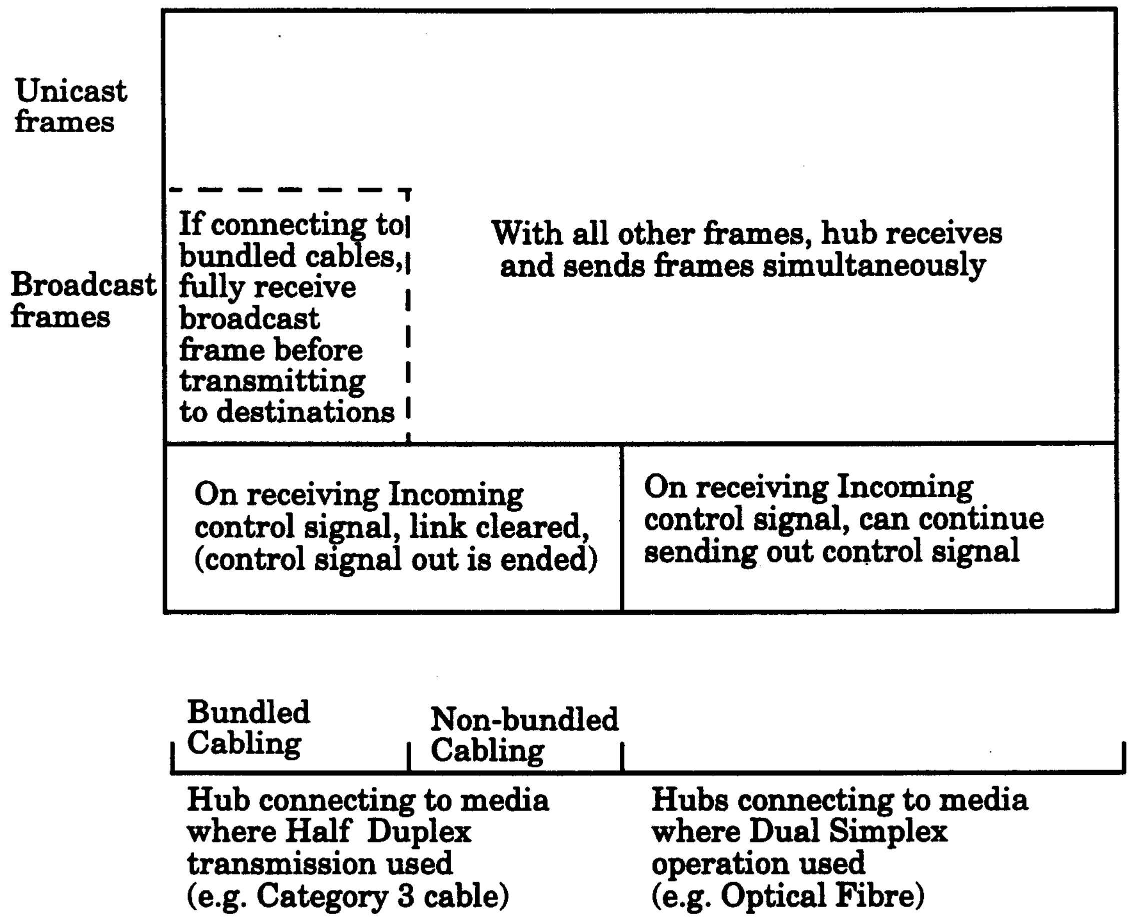

Figure 7 Demand Priority Support for various Media Types

All unicast frames are sent only to the destination End Node whose 48-bit Address matches the Destination Address of the frame which the Hub is receiving. (The Hub learns of the End Node's Address during a small 'Training' sequence of Hub to End Node frame exchanges when the link is initialised. Training sequences can be readily implemented in hardware.)

The sending of frames only to their intended destination is a privacy feature enhancing network security.

If a Hub is connected via bundled cables, broadcast frames are fully received before they are transmitted to the destination End Nodes. This minimises cross-talk (see Section 2.1). With unicast frames, the Hub receives and sends frames simultaneously, as shown in Figure 6.

In Hubs which do not connect to bundled cables, all frames are received and sent simultaneously, as shown in Figure 6.

Half duplex operation is defined over Category 3 cable to maximise the frame data transmission rate (see Section 2.1). On half duplex links only, the End Node clears the link on receiving the INCOMING control signal.

On 4-pair Category 3 cable, the control signals are implemented as low frequency (0.9375 MHz and 1.875 MHz) tones.

3.4 The Demand Priority Access Contention

At the point where a Hub clears the link to solicit the next frame from a End Node, it may be receiving Requests from several Nodes. It has to decide which End Node it will next solicit a frame from.

Demand Priority uses a 2-priority Round Robin.

In a single priority Round Robin, an N-port Hub which has just serviced a Request from a End Node connected to a particular port will select the End Node connected to the next port, looking at increasing port number, at which it is receiving a Request. (Increasing port number is cyclic; this implies that after examining Port N, the Hub goes back to examine Port 1.)

For the 2-priority Round Robin, the Hub services High priority Requests before any Normal Priority Requests are serviced.

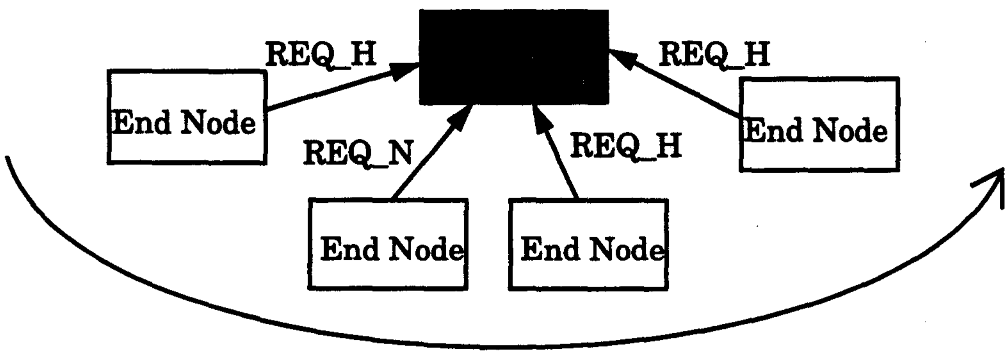

Figure 8 The Hub uses a Cyclic order of servicing Requests from End Nodes. High priority Requests take precedence over Requests at Normal priority.

If a Normal priority Round Robin cycle is interrupted by the Hub servicing newly arrived High priority Requests, then when the latter have been serviced, the Normal priority Round Robin is resumed where it was interrupted. The 2 priority Round Robin ensures that at each priority level, all of the Nodes attached to the hub have an equal opportunity of transmitting frames, ensuring fairness within each priority level.

4 Demand Priority Cascading of Hubs

Demand Priority Hubs can be interconnected such that there is a single Demand Priority shared media access domain across the multiple Hubs and their connected End Nodes.

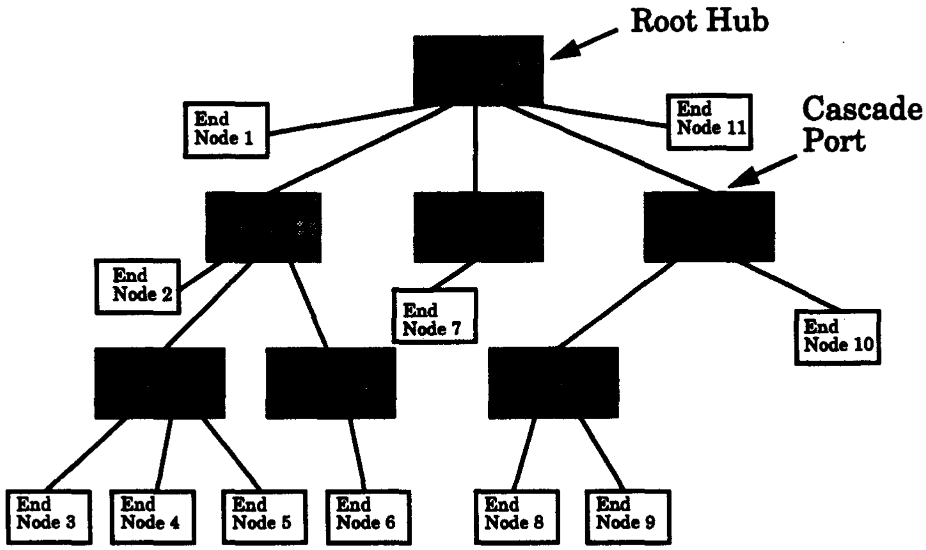

Figure 9 Cascading of Demand Priority Hubs

The Demand Priority Hubs each have a port acting as a Cascade port, which is connected to the Hub in the next level up. Each of the Hubs are identical. A Hub knows that it is the Root Hub by the fact that it has no communication through its Cascade Port.

If all of the Nodes in Figure 9 have frames to transmit at the same priority level, then the order of transmitting frames is the same sequence as the End Node naming - End Node 1, then End Node 2, etc. This ensures that fairness within each priority level is maintained. Also, a High priority Request will rapidly take precedence over Normal priority traffic anywhere in the network.

In brief, cascading is implemented by Hubs selecting Hubs below them, but realising that control will remain with the lower Hub while the lower Hub proceeds through *its* Round Robin cycle before passing control back to the original Hub. A Hub receiving a frame repeats it to all Hubs connected to it, the only exception being that if a Hub receives a frame from another Hub, it does not repeat the frame back to the original Hub. Thus all frames pass through each Hub once. Hubs take the decision whether to transmit frames on to End Nodes, based on address match.

Requests received from lower Hubs are reflected up to the Root Hub. A Hub (e.g. B in Figure 9) will be aware of the priority level of the traffic generated by Nodes connecting to some Hub (e.g. E) below it. If this traffic is Normal priority, and the former Hub (e.g. B) receives a High priority request on another port, then it sends a control signal PRE-EMPT to the lower hub. This forces the lower hub to suspend its Normal priority Round Robin cycle, and pass up control, allowing the upper Hub to initiate a High priority Round Robin cycle. This ensures that High priority traffic pre-empts Normal priority traffic, and can be given fast service.

5 Conclusions

Demand Priority, the IEEE 802.12 MAC is a forward looking hub-based protocol, with its support of the new networked Multimedia applications by a second priority level and the consequent capacity to guarantee bandwidth to applications.

It also supports the very large base of IEEE 802.3 10Base-T users by allowing them to upgrade to 100 Mb/s without changing any host software, and without having to change their currently installed Category 3 cabling. In a similar way, current users of IEEE 802.5 Token Ring can also upgrade to 100 Mb/s, without making software or cabling changes.

Demand Priority also minimises the complexity of the end node adaptor card.

References

[1] A. Albrecht, J. Curcio, D. Dove, S. Goody, M. Spratt, P. Thaler, "A Generic Communications Architecture", HP Memo, 1991.

[2] A. Albrecht, D. Cunningham, "IEEE 802 Tutorial on 100Base-VG", IEEE 802 Plenary, July 1993.

[3] J. Grinham, M. Spratt, "IEEE 802 Tutorial on Multimedia and 100Base-VG", IEEE 802 Tutorial, March 1993.

[4] G. Watson, "Monitoring Bandwidth on 100Base-VG", Presentation to IEEE 802.12, September 1993.

[5] J. Grinham, M. Spratt, "IEEE 802.12 Demand Priority and Multimedia", Proceedings of the 4th International Workshop on Network Operating Systems Support for Audio and Video, 1993.