

Performance Analysis of the Demand Priority LAN

P. Martini, J. Ottensmeyer
University of Paderborn
Dept. of Math. and Comp. Science, 33095 Paderborn, Germany
fax.: +49-5251-603516, e-mail: {martini, otty}@uni-paderborn.de

Abstract

The Demand Priority LAN is an upcoming standard of IEEE's LAN/MAN Standards Committee (still better known as IEEE Project 802). It aims at allowing for a smooth migration from both ISO/IEC 8802-3 and ISO/IEC 8802-5 LANs to a higher speed LAN. Additionally, the new protocol aims at supporting interactive multimedia applications by providing a low latency service for high-priority traffic. This paper is based on Draft 6.0 of the upcoming standard and the results of discussions during the meeting of IEEE 802.12 in Nov. 1994. It starts with a brief description of basic characteristics of the network topology and the medium access control protocol. In the main part, it presents simulation results for normal priority traffic in scenarios with different variable bit rate high priority loads. It addresses important implementation issues and clearly shows the advantages of using service strategies different from those included in the current version of the standard.

Keywords

Demand Priority LAN, FDDI, Interactive Multimedia Communication

1 INTRODUCTION

In February 1980, the IEEE Computer Society held the first meeting of a "Local Network Standards Committee" which became known as "Project 802". After a recent change in the organizational structures which gave this committee more involvement in the standards sponsorship and approval process, the committee was renamed to "LAN/MAN Standards Committee (LMSC)". Today, the LMSC is organized in 14 groups which are numbered from 802.0 which is the "Sponsor Executive Committee" through 802.14 which is a new group working on a Cable-TV protocol. IEEE 802.12 is the group working on a central-repeater-controlled LAN using the demand priority access method. The creation of 802.12 was initiated by massive interest in a "higher speed Ethernet" which was firstly noted by the LMSC in autumn 1992. This massive interest finally resulted in the approval of two different activities: "Project 802.3u" working on a standard for a CSMA/CD protocol operating at 100 Mbit/s over a distance of at most 250 m and "Project 802.12" working on a new demand priority protocol also supporting multimedia applications. Draft 6.0 dated Nov. 23, 1994, states the goals of the demand priority protocol as follows:

- a) Provide a minimum data rate of 100 Mbit/s
- b) The overall goal of this protocol is to provide a higher speed LAN with deterministic

- access, priority, and optional filtering.
- c) Support either ISO/IEC 8802-3 or ISO/IEC 8802-5 frame format and MAC service interface to LLC.
- d) Support a cascaded star topology over twisted pair and fiber optic generic building wiring.
- e) Allow topologies of 2.5 km and greater with 3 levels of cascading
- f) Provide a physical layer bit error rate of less than 10^{-8}
- g) Provide fair access and bounded latency
- h) Provide two levels of priority, normal and high priority
- i) Provide a low latency service through high priority for support of multimedia applications over extended networks
- j) Support an option for filtering of individually addressed frames at the repeater to enhance privacy
- l) Support network management to monitor network performance, isolate faults, and control network configuration
- m) Enable low cost implementation and high levels of integration
- k) Provide for robust operation by testing the Physical Layer connection before allowing an end node to enter the network and by removing disruptive nodes.

The demand priority medium access protocol has been defined both to be bit-rate independent and to allow flexibility in the network topology. Basically, it is a priority-based, round-robin protocol where transmission is controlled by central network controllers (called "repeaters").

In this paper, we study the performance of a single-level 802.12 network with both normal and high priority requests for transmission. But before doing so, we discuss the most important details of the Demand Priority LAN. Additional information on the protocol may be found in papers such as [Gri94], [Spr94a] and in the draft standard [IEEE94].

2 TOPOLOGY AND PROTOCOL OF THE DEMAND PRIORITY LAN

For many reasons there is a trend towards star wired systems where end nodes are connected to hubs by point-to-point twisted pair or optical fiber. Even in the FDDI token ring when operated as a LAN the workstations are interconnected by FDDI Concentrators. Therefore, with smooth migration from existing LANs being a major goal, IEEE 802.12 decided also to use a star wired topology.

2.1 Network Topology

According to the draft standard of IEEE 802.12 each end node (DTE) is interconnected to a hub by Category 3 UTP cabling (possibly bundled), by Category 5 UTP cabling, by Shielded Twisted Pair or by optical fiber. These hubs called 'repeaters' are the network controllers of the Demand Priority LAN. They are configured with n down-link ports where one down-link port is required for each connected DTE. Optionally, a repeater can also have one up-link port. A layered network may be achieved by connecting upper-level repeaters to lower-level repeaters through down-link ports as shown in Figure 1.

The DTEs may be any kind of end stations (PCs, workstations, peripherals, file servers, LAN analyzers, ...). They may also be bridges or routers to other LANs. As far as the links are concerned, the draft standard specifies 4-pair cabling in case of UTP category 3 and 2-pair cabling for UTP category 5 (in both cases link segments from 0 meters to at least 100 meters). Bundled cable may only be used for repeater-DTE links where the down-link port is not operating in the so-called "promiscuous-mode", i.e. where the DTE is not receiving all frames. Alternatively, links may be two shielded twisted pairs (again link segments from 0 meters to at least 100 meters) or two optical fibers (link segments from 1 meter to at least 2 km for 1300 nm transceivers).

Each down-link port connected to a DTE may be configured either to privacy mode (frames are sent to the DTE only when addressed to it) or to promiscuous mode (all frames are

forwarded to such a port). Down-link ports connected to repeaters, bridges, or network analyzers are always configured for promiscuous transmission.

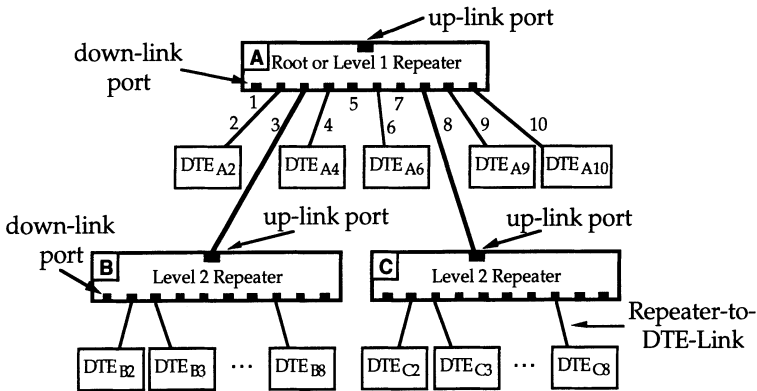


Figure 1 A Multi-Level Cascaded Network.

Since privacy was one of the major goals of the Demand Priority LAN, promiscuous-mode requests may be over-ridden by the network manager. In contrast to draft 3.0 of the standard a revised version now makes the support for privacy mode optional.

2.2 MAC Services

The MAC sublayer provides services to the local LLC (in case of an end node), to the internal relay entity (in case of a bridge), and to the repeater's local LME entity, cf. Figure 2. The repeater medium access control (RMAC) sublayer also included in this figure is the primary control center of the network or subnetwork. More information about the RMAC may be found in section 2.4 of this paper.

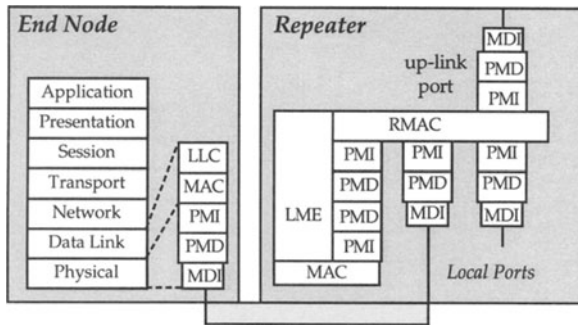


Figure 2 Relation to the OSI Reference Model and to the IEEE LAN Model.

The 802.12 draft standard defines two service primitives for the MAC-LLC interface (MA-UNITDATA.request and MA-UNITDATA.indication) with Destination Address (DA), Source Address (SA), msdu and priority as parameters.

The 802.12 network supports the integration of real-time services by providing two levels of service. *Normal* and *high* priority transmission shall be configurable in the MAC-driver

software. In IEEE 802.12 networks using the token ring format, the priority corresponds to the priority level in the 802.5 frame: the draft standard recommends to use *normal* priority for 802.5 levels "000" through "100" and *high* for 802.5 priority levels "101" through "111". The 802.3 protocol does not support different levels of service. Therefore, in 802.12 networks using the 802.3 format the priority is an additional service parameter which would be mapped to "not available" in a 802.3 network. The default value for MA-UNITDATA.request without priority indication is *normal*. The priority is encoded in the 802.12 frame start delimiter.

2.3 MAC Frame Formats

A special 802.12 frame format is used for training frame construction during link initialization only. Otherwise, 802.12 LANs use either the 802.3 frame format or the 802.5 frame format. In addition to the MAC overhead (at least 18 byte) according to Figure 3, the IEEE 802.12 frames include preambles, start and end delimiters (260 bit for fiber optic networks). It is not intended to support both 802.3 and 802.5 formats at the same time.




Figure 3a The 802.3 MAC Frame Format. **Figure 3b** The 802.5 MAC Frame Format.

2.4 Polling and Data Transmission

In a Demand Priority LAN as specified in draft 6.0 of the draft standard, at any given time there may be at most one DTE transmitting a frame. This frame is forwarded to all repeaters. It is also forwarded to all DTEs explicitly addressed and to all DTEs connected to a port operating in promiscuous mode. If, for example, DTE C8 in Figure 1 transmits to DTE B3, then the frame is transmitted from C8 to repeater C, from there (via repeater A) to repeater B which forwards the frame to the final destination B3. The frame is also forwarded to any port operating in promiscuous mode.

Our performance analysis presented in section 3 only studies a single-level network, i.e. a network where all DTEs are connected to one repeater. Therefore, the description of polling and data transmission presented in this section only addresses this case. A full description of the protocol between repeaters is not within the scope of this paper.

2.4.1 Selecting a Port in a Single-Level Network

In a single-level network the demand priority access reduces to a local priority-based round-robin polling: DTEs ready for transmission send *normal* or *high* priority requests to the network controller, i.e. the repeater. Active DTEs not requesting transmission send Idle_Up control signals. The repeater keeps track of transmission requests by *regularly* checking each port. The draft standard does not require a specific method of checking for transmission requests. However, it requires each port to be checked at least once each frame time.

Normal and *high* priority requests are served in separate cycles. If the repeater is in the process of servicing a sequence of normal priority frames and a high priority request is received then it suspends service of the normal priority traffic to serve the high. However, it does not interrupt *normal* priority transmission in progress. When there are no more *high* priority requests pending, the *normal* priority cycle is resumed with the port following the last *normal*

priority transmission.

For both priority levels, ports are served according to a single frame service strategy, i.e. no DTE is allowed to follow itself in the transmission cycle if there is at least one DTE with a pending request at the same (or higher) priority.

At the same time that the repeater selects the next DTE with a request pending, it alerts all other DTEs of a possible incoming frame by sending a specific control signal called *Incoming*. The DTE selected by the repeater starts transmitting its frame immediately after detecting the *Grant* signal sent by the repeater. The repeater decodes the destination address of the frame currently transmitted and forwards it to the addressed DTE and any promiscuous port as soon as possible. Store-and-forward operation is only required for broadcast and multicast with 4-UTP link bundled cable. For the fiber optic networks studied in section 3 of this paper frame transmission can be overlapped with frame reception.




Figure 4 Frame Transmission Space/Time Diagram.

Figure 4 shows the data flowing through the repeater. The gap between outgoing frames is determined by two components:

- One component is the time that elapses from selecting the next sender until the first bit is received. For a 100 m UTP link this is less than 1 μ s, for a 2 km fiber link it is 20 μ s.
- The other component is the "retransmit delay" which is the sum of the time that elapses until the destination address is completely received and the time needed for matching the destination address to the corresponding port.

Obviously, the gap between outgoing frames determines the efficiency of the network. It determines the difference between the data rate used on the links and the bandwidth really available to the end nodes.

2.4.2 Buffer Control and Address Lookup in the RMAC

The Repeater Medium Access Control (RMAC), cf. Figure 2, controls network access and frame sequencing. The RMAC makes sure that *high* priority request are served first, that no sending station follows itself if other requests are pending and that transmission sequences follow the network port order. Many functions performed by the RMAC are identical to those performed by the MAC sublayer of a DTE. However, instead of forwarding the frames received to upper layers, the RMAC returns the frame to the physical sublayer for retransmission. The RMAC also checks for FCS violations. However, in contrast to MAC sublayers in DTEs it does not discard corrupted frames. Obviously, this would require store-and-forward operation which is not acceptable except for special cases.

Within the RMAC, the most important functions are performed by the Buffer Control and Address Lookup Logic (BCALL). The BCALL connects the RMAC buffer to the input source currently selected for transmission, ensures that the maximum frame size of 4520 byte is not exceeded and controls the frame retransmission. Thus, it controls an ADDRESS_MATCH register in such a way that each one-bit component (there is one per port) is set to True, if and

only if the frame is to be transmitted through the corresponding port. For the fiber optic networks studied in section 3, frame forwarding can start immediately after setting the ADDRESS_MATCH register. Thus, the retransmit delay shown in Figure 4 basically is determined by the time needed for the address lookup. Section 3.3 studies the impact of this parameter on the network performance. As opposed to version 3.0 of the draft standard, version 6.0 now specifies a maximum time of 4.5 μ s allowed for this operation.

3 PERFORMANCE ANALYSIS OF DEMAND PRIORITY

A performance analysis of a newly proposed LAN does not make much sense if it does not include a comparison to existing standards and products. The obvious alternative to Demand Priority is the FDDI multiple token ring. FDDI is the first high speed local area network operating at 100 Mbit/s which is based on an international standard. The main advantages of the Demand Priority LAN supposedly will be lower costs, use of the existing cabling and maintaining the frame format of Ethernet. Points like these cannot be included in performance studies. Instead, we compare the performance of FDDI to the fiber optic version of the Demand Priority LAN. Coming close to FDDI means a good result since the Demand Priority LAN has been tailored to the requirements of totally different environments.

A description of the well-known FDDI is not within the scope of this paper. The interested reader is referred publications such as [Jai94].

3.1 The Load Model

For our studies of the Demand Priority LAN a load model is required which includes both multimedia and data traffic. As far as the high priority (HP) traffic (interactive multimedia) is concerned, we stay close to the VBR load model which was already used in [Mar94] for a comparison of real-time communication strategies in DQDB. This load model is an on/off source generating fixed length bursts (16 ms) with isochronous frame arrival, cf. Figure 5. The bursts are separated by idle periods of exponentially distributed length where the average length is 16 ms, too. Thus, the peak rate is two times the mean rate.




Figure 5 The VBR Load Model.

For each scenario studied, the average high priority load imposed to the MAC sublayer is 0 Mbit/s, 2 Mbit/s, 4 Mbit/s, 6 Mbit/s, and 8 Mbit/s, respectively. We choose the frame sizes as follows:

average HP-load [Mbit/s]	0	2	4	6	8
frame size [byte]	-	512	1024	1536	2048

The reason for choosing different frame sizes for different loads is to keep frame spacing (1000 μ s) identical within a given scenario. Of course, choosing frame sizes as multiples of

512 byte is somewhat arbitrary. Therefore, in section 3.4.2 the impact of this assumption is studied by changing the frame sizes to multiples of 48 byte resulting in a frame spacing of $\approx 94 \mu\text{s}$.

Normal priority (NP) traffic (file transfer, telnet or e-mail) is assumed as poisson arrivals of fixed size frames. We also choose 512 byte for this size. And we also study the impact of this assumption in section 3.4.2.

It should be noted that the total load imposed to the network is higher because it additionally includes preambles, delimiters and 802.3 (or 802.5) MAC frame overhead for all normal and high priority frames. For the FDDI simulations included for comparison, we of course add the FDDI overhead instead. This is all taken care of in our simulations.

3.2 The Configuration

In this paper, we study a single-level fiber optic Demand Priority LAN, i.e. a network where all DTEs are connected to one repeater. Obviously, this topology is the same as in a FDDI where all workstations are connected to the same concentrator.

We study a configuration with 10 DTEs each assumed to transmit 5 Mbit/s normal priority traffic (in case of FDDI using the asynchronous service). Five DTEs additionally send multimedia traffic at high priority (in case of FDDI using the synchronous service). The length of each cable connecting a DTE to the hub is assumed to be 100 m resulting in a $1 \mu\text{s}$ segment round-trip delay. For the Demand Priority LAN this is the time that elapses from selecting the next sender until the first bit of the frame transmitted reaches the hub. The minimum inter-frame gap of outgoing frames additionally includes the retransmit delay.

With such a small cable length which is typical for office or laboratory environments the station latency of FDDI becomes an important factor. Based on products commercially available, we choose 60 bit latency for the DTE and 60 bit for each concentrator port resulting in $22 \mu\text{s}$ total ring latency. Thus, $22 \mu\text{s}$ is the minimum token rotation time. The maximum token rotation time is less than two times T_{OPR} , the "operative Target Token Rotation Time". The choice $T_{\text{OPR}} = 5 \text{ ms}$ resulting in a maximum access delay of 10 ms (necessary for real-time multimedia traffic, cf. [Spr94b]) is a reasonable choice for the configuration studied.

The bandwidth allocation for the synchronous service of FDDI is done as specified in [SMT7.3] based on the rate during a period of $125 \mu\text{s}$. The actual allocation is scaled to the common T_{OPR} value by the factor $T_{\text{OPR}}/125 \mu\text{s}$. The allocation chosen is a peak rate allocation, i.e. there are deterministic guarantees for delay and bandwidth.

3.3 Impact of the Retransmit Delay

Firstly, we study the impact of the retransmit delay of the repeater. For the load model described in section 3.1 and the configuration described in section 3.2 the average normal priority load is 50 Mbit/s. With the frame sizes given in section 3.1, the high priority load varies from 0 Mbit/s to 40 Mbit/s.

The retransmit delay is set to 0, 5, 10 and 15 μs , respectively. The value of 0 μs is only used as a boundary which no implementation can achieve (except for repeaters operating all ports in promiscuous mode). A major vendor of workstations and networking equipment aims at 3.2 μs retransmit delay. However, for cheap implementations even 15 μs must be regarded as realistic.

There is no point in presenting detailed studies of high priority traffic delays in our scenario: Early arrivals at the destination do not improve the quality of real-time services; a maximum delay of 10 ms is not violated in any of our scenarios.

Figure 6 compares the mean access delay for the normal priority traffic for different retransmit delays to the delays obtained for FDDI. The dotted line separates simulations with frame loss due to overload from simulations without frame loss.

For the "all promiscuous case", i.e. without retransmit delay, the Demand Priority LAN yields results almost identical to FDDI. For light load, they are even better which is a well-known result from comparisons between central and distributed control in polling systems. For high load FDDI reduces the network control overhead by allowing stations to transmit more than one

frame after capturing the token. This raises the question, whether the MAC protocol of the Demand-Priority LAN should or even must be improved. This is addressed in section 3.4.

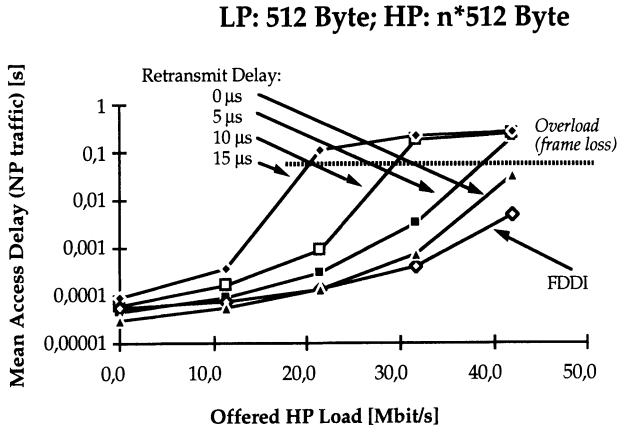


Figure 6 Results for different retransmit delays.

Figure 6 may also be used for a discussion of the total network capacity really available to the DTEs. Values above the dotted line are only based on frames reaching the destination. In addition, there is frame loss because a part of the bandwidth is wasted by the MAC protocol. For large retransmit delays, this obviously is a considerable part of the capacity: For 15 μs retransmit delay the 100 Mbit/s - Demand Priority is overloaded with less than 70 Mbit/s total traffic.

From Figure 6, it comes clear that the retransmit delay has a strong impact on the overall performance of the Demand Priority LAN. By not specifying an upper limit, even extremely bad implementations with exorbitant retransmit delays could have claimed conformance with the standard. After presenting those performance results to IEEE 802.12, a maximum retransmit delay of 4.5 μs was specified in the draft standard.

3.4 Impact of the Service Strategy

Our discussion in section 3.4 already initiated thoughts about possible improvements to the MAC protocol of the Demand Priority LAN. In its current version, the draft standard says: "No DTE is allowed to follow itself in the transmission sequence if other DTEs have requests pending at the same priority level." Additionally, normal priority requests yield to high priority requests. In this section, we compare the following service strategies specifying what to do after receiving a grant:

(a) *Single Frame*

Send one frame. This is the strategy specified in the current draft standard.

(b) *Single Frame and Exhaustive*

b1: grants for high priority requests -> exhaustive service

b2: grants for normal priority requests -> one frame only

Exhaustive service may be used for high priority traffic since bandwidth allocation and control may be assumed for this priority class. Work on the details of bandwidth allocation for high priority traffic is currently regarded as out of scope for IEEE 802.12. Discussions on this issue resulted in the establishment of a Multimedia Task Force within IEEE 802.1.

(c) *Max. Length and Exhaustive*

c1: grants for high priority requests -> exhaustive service

c2: grants for normal priority requests -> up to 4520 byte (maximum 802.12 frame)

As an example: with a normal priority frame size of 512 byte, a node may send up to 8 frames after receiving one grant.

In the following sections, we discuss the results for the different service strategies. Based on the results presented in section 3.3, we take a retransmit delay of 4.5 μ s for granted.

3.4.1 *Large High Priority Frames*

We now compare our simulation results of section 3.3 to the alternative service strategies (b) and (c). Figure 7 shows the results. It turns out that strategies (a) and (b) yield the same results. This fact may be explained by a closer look at the high priority load. The spacing between high priority frames of a given DTE is such that each remaining high priority DTE may send one frame during this time (see also Figure 9). Thus, no backlog of these frames occurs in the transmit queue: At the time a DTE gets the grant on its high priority request, there is always just one frame ready for transmission. In section 3.4.2. it comes clear that this reflects a property of the load model.

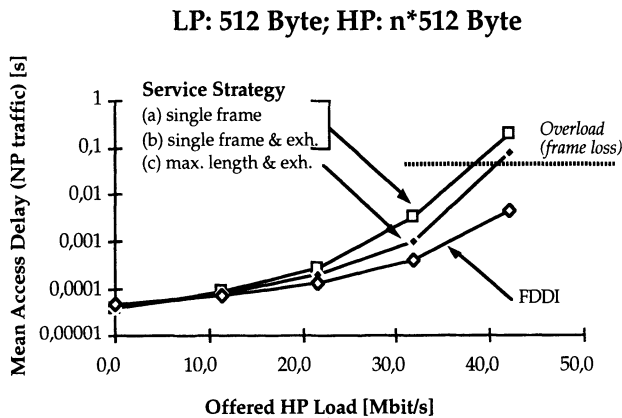


Figure 7 Impact of the Service Strategies for Large High Priority Frames.

Strategy (c) improves the performance of the system by approximately 4%. This is due to the fact that the control overhead is slightly reduced by allowing the DTE to send up to 8 normal priority frames per grant. Even though not shown explicitly in the figure, this also affects the efficiency of the network.

3.4.2 *Small High Priority Frames*

Now we change the size of the high priority frames from 512 byte to 48 byte each. With the same data rate this reduces the spacing of frames during busy periods by a factor of $512/48 \approx 10.7$ and increases the number of frames to be transmitted by the same factor. Thus, for single frame service, the total number of requests per second increases from 14,700 to 38,900. Figure 8 shows the differences between the service strategies for this scenario.

LP: 512 Byte; HP: n*48 Byte

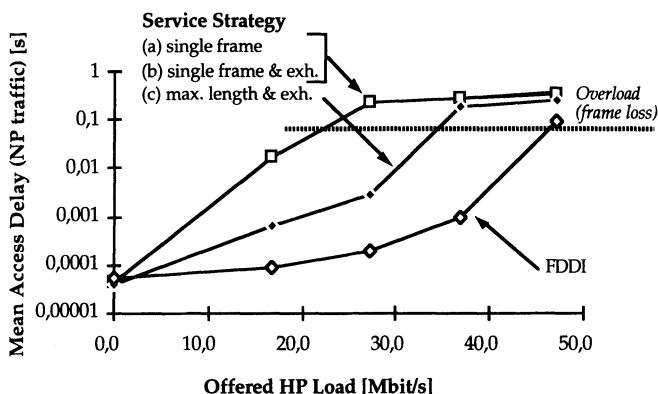


Figure 8 Impact of the Service Strategies for Small High Priority Frames.

Again, strategy (a) and (b) yield the same performance. Basically, the new frame size has not changed the way in which high priority frames are multiplexed. This is also illustrated in Figure 9. There is still no considerable backlog for high priority traffic.




Figure 9 Multiplexing of High Priority Frames.

Even though there is no service degradation for the high priority traffic, the normal priority traffic suffers badly from the additional control overhead in the network. Of course, additional frames are associated with additional preambles, delimiters and MAC frame overhead. This results in the "right shift" of the total high priority load for which simulations were run. In the scenario with all high priority DTEs generating an average load of 8 Mbit/s, the additional overhead takes the network into an overload situation even for FDDI.

For 6 Mbit/s each, there is no problem for FDDI but overload for all versions of Demand Priority studied. For 4 Mbit/s, only strategy (c) yields acceptable results for Demand Priority. The reason for not being as close to FDDI as for large high priority frames is the grant delay and the retransmit delay associated with each additional high priority request. Even for 2 Mbit/s strategy (c) improves the performance in terms of mean access delays for normal priority traffic by an order of magnitude.

Finally, we increase the normal priority frame size to 1024 byte. Running our simulations yields the results shown in Figure 9. In this scenario, the number of requests per second is slightly reduced to 47,000 due to a smaller number of normal priority frames. This does not change much for FDDI and strategies (a) and (c) for Demand Priority. However, for the first time we see different results for strategies (a) and (b). For 4 Mbit/s load per high priority sender, strategy (b) helps the Demand Priority LAN out of the overload situation. For 2 Mbit/s it also reduces delays for normal priority traffic.

LP: 1024 Byte; HP: n*48 Byte




Figure 9 Results for Large Normal and Small High Priority Frames.

The non-preemptive service for normal priority traffic in combination with an increased frame size delays the start of the next high priority cycle for at most $\approx 86 \mu\text{s}$ (for the transmission of a normal priority frame of 1024 byte, instead of $\approx 43 \mu\text{s}$ in case of a 512 byte frame). This is very close to the spacing of high priority frames during busy periods which is $\approx 94 \mu\text{s}$. Thus, there is a real chance for a backlog of high priority frames and a real chance of making exhaustive service different from single frame service. This reduces the amount of time wasted for exchanging requests and grants and improves the efficiency of the network. The additional bandwidth is available to normal priority traffic.

4 CONCLUSIONS

The results presented in this paper clearly show that the retransmit delay is a crucial parameter for the performance of the Demand Priority LAN. Hence an upper limit was specified in the version 6.0 of the draft standard after a no-vote from the authors. Otherwise, bad implementations could have jeopardized the success of this promising protocol for conventional and multimedia environments.

The comparison of service strategies discloses that the single frame service enforced by the current draft standard yields bad results in case of small high priority and large normal priority frames. The load models assumed in the scenarios studied in figure 9 can no way be called unrealistic: 48 byte is a typical frame size for voice traffic, aggregation of voice "channels" can easily result in a spacing of even less than the $94 \mu\text{s}$ assumed. As far as file transfer or paging is concerned, frames should be at least 1024 byte.

The alternative service strategies proposed and studied in this paper improve the efficiency of the Demand Priority LAN, in some cases significantly. Based on early results, the strategy "max. length and exhaustive", i.e. up to 4520 byte per grant for normal priority and exhaustive service for high priority, was already proposed to the IEEE 802.12 committee. In the Working Group Letter Ballot (August 1994) a study item on a "burst mode option", allowing stations to transmit more than one frame, was included in the draft standard.

It should be noted that an exhaustive service (for any priority level) can lock-up the network. Therefore, the maximum amount of data admitted must be limited. In our studies it turned out that on each high priority grant no more than 4000 byte have been sent. This implies to use the maximum length service strategy (assumed for normal-priority service) also for high priority service.

The Demand Priority LAN aims at being a low cost, high speed network. Nevertheless, for light load the performance was shown to be (slightly) better than the performance of FDDI. In this case, it comes close to the performance of a random access network. It has one more point in common with random access networks: Inactive stations such as workstations of employees attending meetings do not change the performance at all. This issue will be considered in a subsequent paper. The same holds for phases without communication requests. For heavy load, Demand Priority does not quite reach the performance of FDDI. However, in this case it outperforms any random access network.

As far as support of real-time multimedia applications is concerned, there is no point in comparing Demand Priority to random access networks like the proposal of project IEEE 802.3u. When compared to ATM LANs it should be noted that there still is no standardized flow/congestion control in ATM. In contrast to this, the request/grant mechanism of the Demand Priority LAN is nothing but a specific flow control strategy between the DTE and the repeater.

5 REFERENCES

- [Gri94] J. Grinham, "Multimedia applications and the IEEE 802.12 demand priority", Proceedings of *EFOC&N '94*, Heidelberg, June 1994
- [IEEE94] IEEE 802.12: "Demand Priority Access Method and Physical Layer Specifications", Draft Standard, Rev. 6, November 1994
- [Jai94] R. Jain, "*FDDI Handbook. High-Speed Networking using Fiber and Other Media*", Addison Wesley 1994
- [MaOt95] P. Martini, J. Ottensmeyer, "Real-Time Communication in the Demand-Priority LAN - The Effects on Normal Priority Traffic", to be published by IEEE in the Proceedings of the "20th Annual Conference on Local Computer Networks", Minneapolis, Oct. 1995
- [Mar94] P. Martini, "Connection Oriented Data Service in DQDB", in *Computer Networks and ISDN Systems*, Vol. 26, Numbers 6-8, March 1994
- [Ott94] J. Ottensmeyer, "Simulation Results of the Demand Priority Protocol", Presentation at the IEEE 802 Plenary, July 1994
- [SMT7.3] ISO WD 9314-6: "FDDI - Station Management (SMT) Rev. 7.3a", June 1994
- [Spr94a] M. Spratt, et. al., "IEEE 802.12 - The local area network upgrade to 100 Mbit/s and beyond", Proceedings of *EFOC&N '94*, Heidelberg, June 1994
- [Spr94b] M. Spratt, "LAN Bandwidth Allocation using the Target Transmission Time (TTT)", Contribution to the March 1994 meeting of the IEEE 802.1 Working Group on Multimedia

6 BIOGRAPHY

Peter Martini is professor of computer science at the University of Paderborn, Germany, since 1990. His current research interests include computer networks, Local and Metropolitan Area Networks, performance evaluation, modeling and simulation techniques and workload characterization.

Jörg Ottensmeyer received his Diploma in Computer Science in 1992 from the University of Paderborn. Since then he has been a member of the Networking Group headed by Prof. Martini. Currently, he is active in the area of real-time communication in LANs and MANs.