Shared versus Switched

or

Reports of the Death of Shared Access are Greatly Exaggerated

Graham Campbell
Computer Science Department
Illinois Institute of Technology
campbell@charlie.iit.edu

Abstract

Indications are that a shared environment supported by DQSA (Distributed Queue Switch Architecture) is superior to the conventional switch/router environment for most communications. DQSA and its constituent MACs (DQRAP, XDQRAP, DQLAN, PDQRAP, and CDQ), developed at the Illinois Institute of Technology, represent a major development in shared access communications. A DQSA-based MAC (medium access control) overcomes the limitations of Ethernet, especially those with respect to distance and QoS, thus permitting shared access, normally restricted to LANs, to be applied to WANs. This paper provides an overview of the current LAN environment, presents a set of criteria describing the ideal MAC, describes the DQSA constituent protocols, and describes potential applications for these protocols. It closes with the argument that DQSA has the potential to satisfy most communications requirements without the use of STM circuit switches, packet routers or conventional ATM switches.

1 Introduction

1.1 Preamble

The May 18th, 1998 issue of Network World contains an interview with the inventor of Ethernet, Bob Metcalfe, on the occasion of the 25th anniversary of the invention of Ethernet. There was also a companion roundtable discussion with five industry leaders. The consensus of the leaders was that, as speeds increased, switched Ethernet would dominate and the shared access component of Ethernet would disappear leaving only the frame structure and addressing. The feeling was summed up in a statement by Dr. Eric Cooper, Chairman of Fore Systems: "As the historical baggage of shared access is left behind, there is no technological reason why Ethernet-compatible address and frame formats can't be used on higher and higher speed dedicated links."

We agree with the statement that Ethernet compatible address and frame formats will remain in use but argue that the use will still be in shared access. The difference is that the shared access will be the MACs that constitute DQSA (Distributed Queue Switch Architecture). Read on and learn about a communications technology that will introduce a fundamental change in the way packets are transported.

1.2 How It All Started

Communications networks utilize one of two physical environments: (a) shared medium or (b) switched medium. In a shared medium network a station communicates with another station by transmitting a message on a channel that is shared by all stations -- all stations receive the message but only the intended destination station(s) copy the message. In a switched network a station communicates with another station by utilizing zero or more switches that connect a series of point-to-point lines -- only the destination station(s) receives a copy of the message.

Both mechanisms have been in use since the advent of electrical-based communications. Samuel Morse's first telegraph line between Washington, D.C. and Baltimore was a point-to-point connection as were Alexander Graham Bells's first phone call to Mr. Watson and Marconi's first transmission of a wireless message. But very quickly economics dictated shared use of lines for the above in many scenarios: in wireless a specific frequency was shared by multiple users; in the nascent telephone industry, competitors of the original Bell Telephone Co. introduced the ubiquitous party line that allowed multiple customers to share a common line attached to a central switch. In the latter case the shared/switched system has been replaced, at least in the U.S., by connecting each subscriber via a point-to-point line to a central switch so that a direct connection can be established as required between any two subscribers.

The fundamental problem of utilizing a shared medium for communications is the allocation of the medium to a specific user. There are two basic methods: deterministic and non-deterministic. The deterministic method ensures that each station receives an equal opportunity to use the medium: the common methods are to either use a master station to poll each station in turn or to distribute the responsibility by having a token pass in turn to each station, possession of the token granting temporary control of the medium. The deterministic method can provide guarantees with respect to access time and intervals between opportunities to transmit but overall utilization is good only when: (a) every station is always ready to transmit; or (b) the message transmitted by a single station is very long relative to the time it takes to canvas all the stations.

Non-deterministic methods utilize contention-based mechanisms and can make no guarantees about performance that has proved to be satisfactory only when the offered traffic is light relative to the capacity of the medium.

In general a shared medium network, using either deterministic or non-deterministic access methods, provides satisfactory performance only when the length (time to transmit) of the message is much longer than the time it takes for the message to travel from source to destination. The desired condition is

© 1998d

that when the first bit of a message has reached the destination, more than 50% of the bits in the message still remain to be transmitted. The ratio of the time it takes for a bit to travel between two stations to the time it takes to place the message in its entirety on the line is very useful and is often given the name “a” (little “a”) = Tp/Tx where Tp is the propagation delay and Tx is the length of the message. The units can be bits or time but must be consistent. Stallings [1] provides a good discussion of the fundamentals.

1.3 The Arrival of Computers

The party lines of the telephone system and the shared use of a frequency in wireless worked well because the time it took for users to sort out who next used the line was short relative to the time the user occupied the line, i.e., the equivalent of our little “a” was very small. However the arrival of the computer and the subsequent desire to network computers reintroduced the question: to switch or to share? Initially all networking of computers utilized point-to-point links but the bursty nature of computer traffic that often resulted in low utilization of the links spawned active research into the development of mechanisms, termed MACs (medium access control protocols), that would permit multiple computers to utilize a common channel.

By the mid 1970s the theory of operation of the deterministic (token ring, token bus) and non-deterministic (Aloha, Slotted Aloha, and CSMA (Carrier Sense Multiple Access) MACs was reasonably well understood. Very early on, Aloha, not sensitive to the value of “a”, was and continues to be used in satellite networks. CSMA with collision detection (CSMA/CD) was introduced soon after as Ethernet and to this day dominates the MACs used in LANs. All are described by Stallings [1].

1.4 The Standards

Given the subject of this paper it is well to spend some time discussing the IEEE 802 standards since collectively they make a fundamental statement about the status of shared access.

Ethernet™ was invented in 1973 by Bob Metcalfe while at Xerox but it was in 1979 while in a discussion with Gordon Bell of DEC that he had the idea of establishing a standard as a mechanism that would allow DEC to contribute to the development of LAN technology without having to reinvent the wheel [15]. Xerox was agreeable, Intel joined the effort and in 1980 the IEEE established the LAN/MAN Committee to define standards for the emerging LAN market. Three projects were established: IEEE 802.1- Station Management; IEEE 802.2 - Logical Link Control; and IEEE 802.3 - CSMA/CD. IEEE 802.1 provides standards for the administration of the LAN, currently it is defining a standard for implementing VLANs (Virtual LANs) that span several physical LANs interconnected via switches and/or bridges; IEEE 802.2 implements a mechanism that supports virtual channels and provides for acknowledgments for the IEEE 802.3 datagram service. In practice most implementations do not utilize IEEE 802.2 and instead use Ethernet. The sole difference between IEEE 802.3 and Ethernet is that the contents of a 16 bit field in the frame contains either a length signifying an IEEE 802.3 frame or a type signifying an Ethernet frame. Systems can run in either mode.

Almost immediately there were concerns about the limitations of IEEE 802.3, specifically its lack of guaranteed service and distance limitations. Some potential users, including General Motors, were looking to LANs for factory automation and concluded that a protocol that could provide guaranteed service over a greater distance was the answer. The IEEE 802 Committee established IEEE 802.4 -- Token Bus. IBM also had reservations about Ethernet’s lack of guaranteed service and also with the fact that it could not scale to higher speeds while maintaining current distance and packet size and so IEEE 802.5 -- Token Ring came into being. The reader will note that if you have a protocol it appears advisable to have a patron that is a heavy hitter. Project IEEE 802.6 was established to address the MAN (metropolitan area networks) part of the IEEE 802 official designation, specifically targeting the cable TV plant as a platform to deliver digital data to homes. Unfortunately there was close to zero interest on the part of the cable companies in this concept in the early 1980s so IEEE 802.6 drifted for awhile and then settled on making a standard out of a new shared access method that had originated in Australia -- DQDB (distributed queue dual bus).

Once again the lack of what we now call QoS (Quality of Service) was noted in Ethernet and so IEEE 802.9 -- Integrated Services came into being. It defined a standard that “stitched” some 6 Mbps of synchronous channels, suitable for voice, on top of Ethernet. The dissatisfaction with Ethernet continued with the establishment of IEEE 802.11 whose goal was a standard for wireless applications. IEEE 802.12 -- Demand Priority, was the result of a schism in 802.3 between a group that argued that Ethernet could run at 100 Mbps by simply reducing the distance covered from 2.5 Km to 250 meters. The other group argued that only a deterministic access method would be suitable for 100 Mbps speeds. The IEEE 802 Executive settled the matter by disregarding their charter of one problem, one solution, by agreeing with both camps and authorizing 802.3 to develop a 100 Mbps version of CSMA/CD and establishing project IEEE 802.12 to develop a deterministic shared access method, i.e., Demand Priority. In the mid 1990s there finally emerged an interest in using the cable TV plant to carry data. IEEE 802.6 had been established to define a standard for this area but it was felt they had forfeited their franchise with DQDB so a new project, IEEE 802.14 - Cable TV, was established and is currently at work on a standard.

An overview of the IEEE 802 activity is available at [2]. The missing numbers: 7, 8, and 10, represent standards that are not in themselves access methods. The author doesn’t know what happened to 13 but can guess.

1.5 Comments on the Standards

We stated that the existence of the IEEE 802 standards in itself makes a fundamental statement about the state of shared access. To wit, the existence of eight standards that address basically the same problem indicates that none is satisfactory. Now if Token Bus had overtaken Ethernet and in turn been taken over by Token Ring and so on this would be a natural progression and would account for eight standards. However Ethernet has met all on the field of battle and emerged victorious. In fact the newer the standard the shorter its expected life cycle. Both IEEE 802.9 Integrated Services and IEEE 802.12 Demand Priority disappeared in the MAC pool leaving barely a ripple while IEEE 802.14 may not be deployed anywhere due to the selection by the cable operators of MCNS (Multimedia Cable Network System). IEEE 802.11 Wireless reflects a committee effort to define a standard for the wireless environment. It employs CSMA/CA i.e., CSMA with collision avoidance. The fact that this standard is just being completed is somewhat ironic in that the original work in the late 1960s and early 1970s with Aloha and the CSMA variations were conducted with the view that they would be deployed in the wireless arena. It is much too early to comment on the viability of IEEE 802.11 but the author has seen at least one ad for a wireless system which touted as one of its major features non-compliance with IEEE 802.11. Another possibly more disturbing event is that the HomeRF Working Group, now numbering some forty plus members, will soon release the specifications for SWAP (Shared Wireless Access Protocol) for use in the home environment [27]. SWAP can perhaps be treated as the offspring of both IEEE 802.9 and IEEE 802.11 since it uses the CSMA/CA of dot 11 with provision for TDM channels to carry synchronous voice as in dot 9. The author contends that this is yet another proof of the failure of the existing standards to satisfy anything more than very specific, temporary requirements.

2 What's Wrong and What's Right

Ethernet dominates the communications market for all applications excepting POTS (plain old telephone service) for distances up to a couple of kilometers. Fast Ethernet (100 Mbps) is restricted to 250 meters on fiber and approximately 100 meters on UTP but serves areas beyond this distance by means of switches. Beyond these distances, excepting for satellite and wireless networks, switching/routing rules all communications including POTS. But we are addressing the role of a shared medium in communications and whether a shared medium can be used in distances beyond Ethernet's realm. Let us review the situation: Ethernet rules its ever-decreasing (geographically speaking) domain despite being a three-time loser:

1. **Low Utilization:** Theoretical utilization is high but in practice Ethernet systems are operated in the 5% - 40% utilization range.
2. **Distance Sensitive:** Ethernet is limited to 2.5 Km at 10 Mbps decreasing by a factor of ten for every tenfold increase in speed.
3. **Lack of QoS:** Ethernet has no effective priority mechanism and thus cannot support any semblance of QoS. This problem is being addressed by IEEE 802.1 but the proposed solution has more application to switched Ethernet than to classic shared access Ethernet.

A pair of questions: "Ethernet obviously has serious limitations but just what are desirable features of the "Perfect MAC"? It is strange that in the years since research started on shared access that the boundary conditions of an ideal MAC have not been developed so that at least the flood of standards emerging from IEEE 802 could be rated. And the second question, "If a MAC existed that possessed or came close to the desirable features could that MAC dominate communications for LANs, MANs, WANs, etc., in the way that Ethernet dominates LANs?" We list a set of desirable features for a MAC in the next section and then describe the component protocols of DQSA that will in the future dominate communications.

2.1 Features of the Ideal MAC

1. **Immediate Access**: If the channel is not busy then transmit immediately. *This is the feature that along with a good backoff algorithm makes Ethernet what it is. We shall see that it possesses little else.*
2. **Full Channel Utilization**: System throughput is equal to the offered traffic up to the capacity of the channel whether one station or all stations are transmitting. *Ethernet can achieve full utilization if only one station transmits but terrible things happen when all stations have something to send. Deterministic protocols such as Token Ring provide almost full utilization when all stations are ready to transmit but when only one station is ready much of the capacity could be utilized in presenting not ready stations the opportunity to transmit.*
3. **Minimum Delay:** Packet transmission delay should be that of an ideal M/M/1 or M/D/1 queueing system for variable length or fixed length packets respectively when the offered traffic has a Poisson distribution.
4. **Predictable Delay:** Once a station is ready to transmit it should know how long it will be till it actually can transmit.
5. **Fair Access:** Transmission should be based on FIFO with priorities an option.
6. **Distributed Control:** Nodes manage requests/transmission independently, i.e., no central control required. *Ethernet offers this feature.*
7. **QoS (Quality of Service):** Available via priorities and/or via the equivalent of synchronous channels. *When an application desires say a 1 Mbps channel then that should be available either approximately or exactly according to the requirements of the station.*
8. **Topology Independent:** Our ideal MAC should operate on all topologies. *DQDB possesses many of the ideal characteristics in this list but restriction to a dual-bus topology has turned out to be a fatal weakness. Ethernet originally ran on a single bus but easily adapted to a tree-and-branch topology and then to the current hub and star topology.*

9. **Idle Nodes:** Should not utilize any network resources. *We have reached the fourth and final ideal characteristic that Ethernet possesses.*
10. **Distance Insensitive:** Performance should be independent of little "a": i.e., the distance can range from a few meters to tens of thousands of kilometers, e.g., continental or satellite networks, at any transmission rate.

The first five features describe the performance of a typical packet/cell switch, which is reasonable, since even though our stations are distributed across some arbitrary distance our goal is that the medium connecting them is utilized just as if all the stations were locally connected to a switch.

We indicated that Ethernet possesses just four of the ten desired characteristics of a MAC and yet it is "king of the hill" for all applications excepting voice in a 2.5 Km area. There is a question posed in the title "Shared versus Switched?" i.e., if a MAC possesses or came close on all ten of the desired characteristics, especially #7 (QoS) and #10 (distance insensitive) could it then support most of or all communications that, beyond the two km boundary, presently utilize switching/routing? Needless to say the reason this paper was written and that you are now reading it is that the DQSA family of MACs possesses or comes very close to all ten desired characteristics and so "Yes, Virginia, there is a Santa Claus -- DQSA will bring you all the feature your little heart desires."

3 DQSA: The Protocols

3.1 DQRAP: The seminal DQSA protocol

We have taken a long trip to get to the gist of the paper but here it is. DQRAP (Distributed Queueing Random Access Protocol) is the seminal protocol of DQSA. The author admired the simplicity of DQDB and launched a research effort, the goal of which was a MAC that provided the performance on other topologies that was provided by DQDB over a dual-bus topology. The details of DQDB are available in Stallings [4], but briefly, the performance of DQDB (which does satisfy seven of our ideal characteristics) is based upon stations sending requests via single bits to upstream stations, while at the same time keeping count (queue length) of the bits received from downstream stations. When ready to transmit, the station permits a number of empty slots equal to the current count (queue length) to pass before transmitting. Thus a ready station "joins a queue" of those stations already ready and transmits in turn. The performance is close to ideal but being limited to a dual bus is what caused DQDB to enter the MAC pool leaving barely a ripple.

The dual bus topology allows requests to transmit on the downstream channel to be sent upstream on the other bus. There is no danger of colliding since all upstream stations can "see" the incoming request and thus wait, if ready, to transmit their own request. In a tree-and-branch topology, requests transmitted "upstream" can and will collide with requests arriving on other branches. The problem reduced to one of

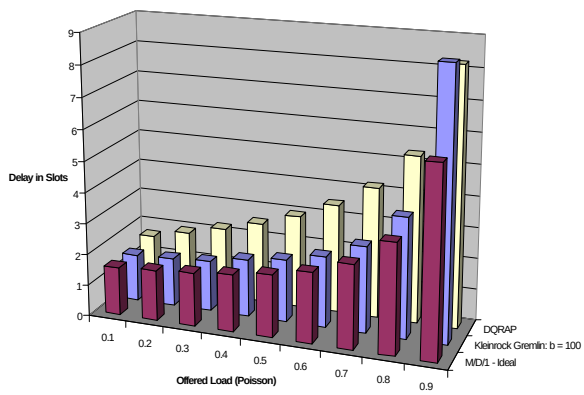

Figure 1. Delay Comparison: DQRAP, M/D/1 and Kleinrock Gremlin

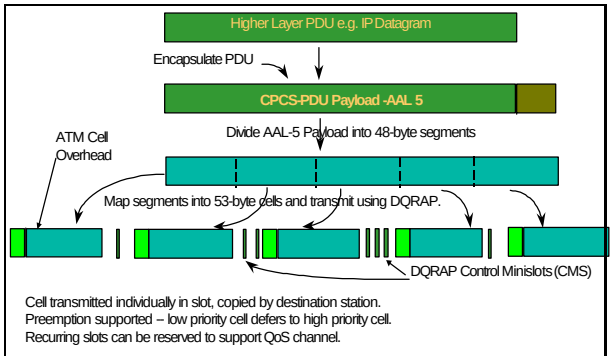

Figure 2. DQRAP Segmentation for ATM

resolving collisions of requests, equivalent to the problem of resolving collisions of data in a shared medium, a problem that had been the subject of essentially fruitless research since the early 1970s.

The author then did what any good professor would do: he assigned his then student Wenxin Xu to "work out the details". Some details! DQRAP is the result but there was only one "detail" taken from DQDB.

Briefly, in DQRAP, control minislots take the place of the request bit in DQDB. Ready stations select one of three CMS (control minislots) randomly in which to transmit a request for a slot. Figure 2 illustrates the relationship of the CMS and the dataslot. Global feedback indicating a success, i.e., no collision, causes all stations to increment their TQ (Transmission Queue) by one. If a collision occurs in a minislot, the Collision Resolution Queue (RQ) is incremented indicating that a group of stations has collided. The group at the head of the RQ is given exclusive use of the CMS to resolve their differences, new arrivals are blocked. While this contention resolving is ongoing the station at the head of the TQ transmits its data. The key to DQRAP is that using only three CMS the system will, on average, resolve collisions caused by the arrival of requests of multiplicity N in a time span of less than N

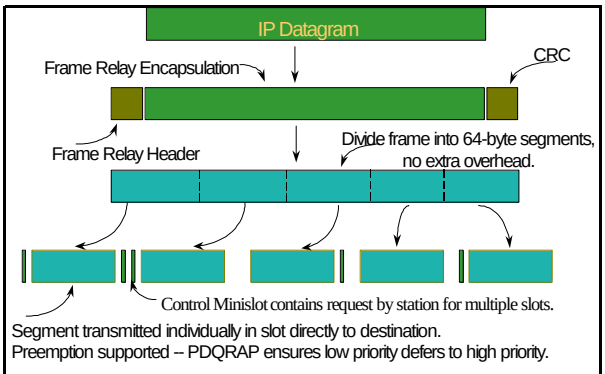

Figure 3. XDQRAP Segmentation for Frame Relay

dataslots thus ensuring 100% utilization of the dataslots.

Interestingly enough, even though the research was inspired by DQDB, the "transmission queue" component turned out to the only DQDB "detail" utilized. DQRAP utilizes components of MACs already extant that include CMS (control minislots), immediate access, tree splitting, and blocking. The tree protocol introduced by Capetanakis [25] can be used to explain why DQRAP works with three minislots. Capetanakis' first tree protocol has a throughput of 0.347, Chang shows that DQRAP using a single minislot has a throughput of 0.347, two minislots provide a throughput twice that at 0.694 and Bingo! three minislots passes the 100% barrier [22]. Thus resolving with as few as three minislots along with queues for both transmission of data and contention resolution are the keys to DQRAP.

The perfect shared access method will have performance equal to an M/D/1 queueing system. Kleinrock in an invited paper [28] states the same and analyzed such a system assuming that there was an all-knowing "gremlin" that could pass the word to a terminal on the total number of busy terminals, thus permitting that terminal to determine its position in the queue. He was realistic in recognizing that even a gremlin requires bandwidth to pass this information and so included this factor in his analysis. He plots delay for ratios, b, of message size to information of 10, 100 and 1000 against offered traffic. Figure 1 shows the delay for the ideal M/D/1 system, Kleinrock's Gremlin: b = 100, and DQRAP for offered traffic ranging up to 90% of capacity. DQRAP performance is impressive as compared with the perfect M/D/1 system and Kleinrock's theoretical system. It does appear that at 90% traffic DQRAP has a lower delay than Kleinrock, but it is not so. Kleinrock charges the gremlin overhead against capacity thus impacting the curves as system capacity is approached. DQRAP plots against available slots excluding the overhead of the minislots. Kleinrock also assumes past history is not available to the gremlin thus total information is sent each time. DQRAP maintains a past history with the TQ and RQ thus reducing the amount of information sent, especially at high loading. DQRAP minislot overhead in practice will range from 1% - 2% in a synchronous system to 15% - 20% in a satellite system.

The delay for DQRAP at offered traffic of 90% is 8.25 slots as compared with 6 slots delay for the M/D/1 system. Although not shown, the DQRAP delay at 95% load is 13.52 as compared with 11 slots for M/D/1. With DQRAP at 90% offered traffic, a new station makes an average of 1.6 requests before entering the TQ. The original DQRAP simulation and analytical results, other details, and the DQRAP algorithm are available in [5] and [13]. Kleinrock's material is in [28].

DQRAP utilizes a fixed size data payload. Therefore data being transmitted, e.g., Internet Protocol (IP) datagrams, must be sliced into fixed-size segments. If the ATM (Asynchronous Transfer Method) adaptation layer (AAL) is used for the SAR (segmentation and reassembly) the result is a distributed ATM switch. The conventional central switch is eliminated -- all switching is carried out on the transmission medium under the control of the distributed stations. The segmentation is illustrated in Figure 2.

The operation of DQRAP is totally dependent upon successful determination of the contents of a minislot so Miramica investigated the robustness of DQRAP. His research showed that DQRAP can suffer up to 10% misreading of minislots before performance is affected [14]. The reason is that while three minislots are used to resolve contention, mathematically speaking less than three minislots will do the job -- the difference, especially under heavy loads, allows mistakes to go unpunished. McPheters showed that the performance of DQRAP in the finite-model environment moved to that of the infinite model with as few as six stations [18]. Khasawneh compared the performance of DQRAP on a dual-bus topology with that of DQDB and showed that performances were comparable in that restricted topology [19].

Summing up, DQRAP represents a remarkable breakthrough in that it satisfies two key ideal characteristics: QoS and distance insensitivity, and satisfies or comes close to the other eight desirable characteristics. QoS is discussed in section 3.4, distance insensitivity is discussed in section 3.5.

3.2 XDQRAP (Extended DQRAP)

Each request in a minislot in DQRAP reserves a single slot. What if we permit a ready station to request a multiplicity of slots with one request? This obviously will increase the size and complexity of the minislot (error checking must now be included) but what it buys is the ability to slice a variable length IP datagram into a series of fixed-size segments and to transmit the segments without any further overhead. Another key benefit is that the number of minislots can be reduced to two. In section 3.1 we pointed out that the throughput of DQRAP with two minislots is 0.69, but that assumed one dataslot per minislot request. If stations on the average request anything greater than $1/.69 \approx 1.5$ dataslots then two minislots suffice.

The ATM SAR (segmentation and reassembly) stage requires that each segment include enough information so that the cells can be stitched back together into the original datagram. Let us assume that in XDQRAP an IP datagram is passed across the

Frame Relay interface and thence to XDQRAP. The HDLC frame is 1500 bytes. The SAR stage of XDQRAP requests 1500/64 = 24 slots. All stations receive the request and add 24 to their copy of the distributed TQ to make it 57 (this assumes that the current value was 33). Every station is now aware that after 33 slots have passed they should then treat the subsequent 20 slots as carrying a single datagram. Thus there is no requirement to identify the individual segments. Figure 3 illustrates the segmentation process.

A priority mechanism, described in Section 3.4, allows higher priority messages, for instance control messages or voice packets, to preempt the ongoing transmission of a longer packet. Figure 6 illustrates the simulated performance of XDQRAP and priorities where the offered load consists of a mix of 50% single-slot (64-octet) packets and 50% 30-slot (1920 octets) packets. This type of load is representative of an environment where there are many long packets representing ongoing file transfers with an equal number of short packets sent as acknowledgments, etc. Note that the average delay of the short packets, assigned high priority status, is almost constant at approximately 2.5 slots over offered loads ranging from 10% to 95%. Note also that the average delay of the longer packets at 95% offered load is approximately 300 slots, approximately ten times the packet length. This compares favorably with the ideal delay of an M/D/1 system at 95% load and is a result of both the efficiency of XDQRAP in that when the average packet length is several times the slot length then contention in the minislots is minimal and the fact that a packet of length 30 slots has a built-in latency of 15 slots just to access the system.

Obviously the shorter packets did not queue behind longer packets, a key bragging point of ATM but unlike ATM XDQRAP does not have to encapsulate each segment -- the segments can be sent "naked". Figure 3 illustrates the segmentation process when XDQRAP supports a distributed frame relay switch carrying IP traffic.

XDQRAP was proposed to IEEE 803.14 [26] and the three minislot feature of DQRAP was picked up by the IEEE 802.14 project and stitched on top of a tree protocol [16]. They would have come up with a much cleaner MAC with better performance if they had just used DQRAP or XDQRAP.

XDQRAP is the basis for the DQSA distributed Ethernet, Frame Relay, and IP switches. Simulation results are available in [6], the detailed algorithm is available in [17].

3.3 DQLAN: Variable Length Frames

DQLAN (Distributed Queue Local Area Network) is an implementation of DQRAP that is restricted to the same geographic area as an equivalent speed Ethernet system. Variable length frames are carried without segmentation. The performance at lighter offered loads is almost identical to that of Ethernet, i.e., immediate access -- the station finding the medium not in use transmits immediately. However, as the offered load increases, the number of collisions occurring in Ethernet increases to the point where operation is unreliable. DQLAN moves seamlessly from immediate access to a reservation system thus providing reliable operation up to the full capacity of the channel. DQLAN cannot provide a guaranteed bandwidth to the exactness provided by the slotted versions of DQSA, but implemented with PDQRAP it provides a "controlled" level of service that approaches that of guaranteed bandwidth. Table 1 shows the performance of DQLAN over a range of offered loads at 10 Mbps. Note that throughput tracks the offered load linearly and that at 90% load factor the delay is still reasonable. Details of the algorithm and simulation results comparing DQLAN and Ethernet are presented by Wu and Campbell [7]

DQLAN can be used in wired applications but its major use will be in the wireless environment where it can often be substituted for CSMA/CD or ISMA (Inhibit Sense Medium Access) implementations by simply modifying the code in a DSP. ISMA is the access method commonly employed in mes-

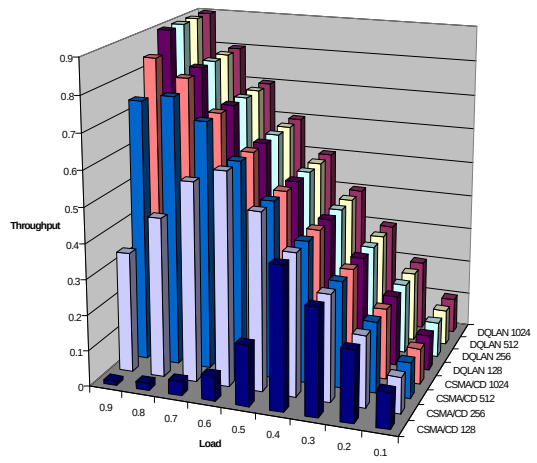

Figure 4 Throughput: DQLAN vs Ethernet

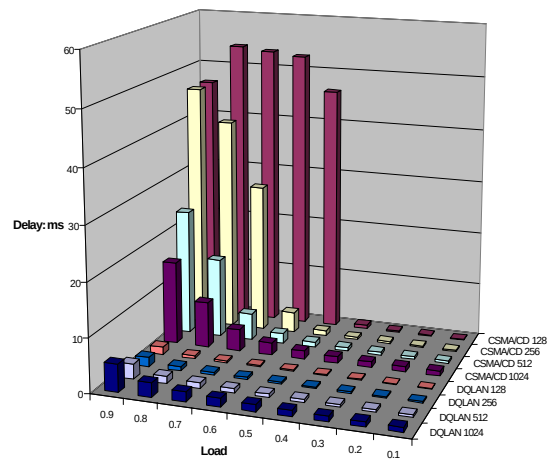

Figure 5 Delay: DQLAN vs Ethernet

sage systems used by parcel delivery services and has the same limitations with respect to throughput that affect CSMA/CD (Ethernet).

3.4 PDQRAP and QoS

The buzzword (or buzz acronym) nowadays is QoS (Quality of Service). In fact it would not matter how good DQSA was with respect to delay, throughput, etc., it would not fly if it could not support QoS. A simple definition of Qos would be "the ability to give 'em what they want!" We will be more precise and use the classes of service as defined by Microsoft for their NDIS (Network Device Interface Specifications), to be available in Windows NT Version 5.0 [24]. The three classes are:

1) Best Effort:
- Default flow
- Not typically requested by applications
- Low priority
- Typically borrows from other flows

2) Controlled Load
- Gets service equivalent to lightly loaded network
- Medium priority

3) Guaranteed Service
- Guaranteed delay bounds
- Highest priority

We address each in turn. DQSA without priorities provides a "best effort" that as shown in Figure 1 is superior to other MACs right up to full utilization.

"Controlled Load" is taken care of by PDQ-RAP, a terrible play on acronyms but PDQ-RAP is Priority DQRAP. In DQRAP a ready station transmits in a CMS and upon receipt of feedback indicating success it joins the transmission queue, TQ. All requests are treated on what is, based on slots, a FCFS (first come, first served) basis. In PDQRAP, a priority bit is added to the minislot and if a station's request is high priority the station sets the priority bit in the minislot when transmitting. All stations receive the CMS and if the priority bit is set then the high priority queue, HTQ, is bumped instead of the TQ. The two queues operate in parallel but stations on the TQ only transmit when HTQ is zero. Simulation and analysis shows that when a PDQRAP system is 90% loaded the average delay for high priority traffic is reduced to that of a lightly loaded system. This is ideal characteristic #5. This also satisfies the Class 2 service requirements described above for Controlled Load. This performance is graphically illustrated in the chart of Figure 6 that was referred to in the section on XDQRAP. The delay of the high priority traffic is nearly constant across all loads with a delay that is comparable to the delay encountered by normal traffic entering a lightly-loaded network, see Figure 1.

PDQRAP is not a self-standing access method in itself but is implemented in the other DQSA access methods. The number

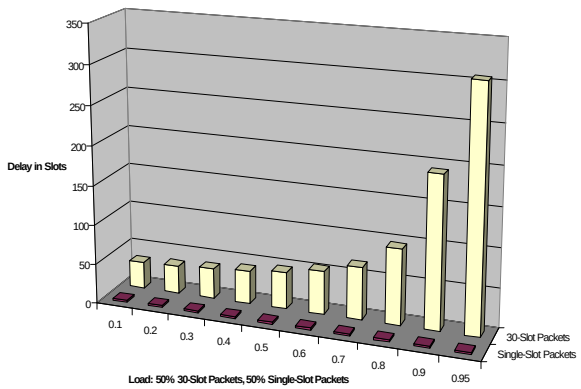

Figure 6. XDQRAP Delay: High and Normal Priority

of priority levels is easily extended by increasing the size of the priority field to accommodate 2^M priorities where M is the number of bits in the field.

"Guaranteed Service" is provided when DQSA is implemented in a slotted system thus allowing a network manager to allocate recurring slots to a requesting station. These recurring slots reflect a specific requested bandwidth and thus Microsoft Class 3 service is supported. This allocation of "real" bandwidth to a user is contrasted with the ATM approach of controlling the total traffic entering the network and then trusting that the limit in traffic plus priority for CBR cells at intermediate switches will satisfy performance requirements. Bandwidth guaranteed by DQSA is still there even when the other two classes of traffic swamp the network.

DQSA is unique in its ability to support the three Microsoft classes of service and thence CBR, VBR, and ABR in ATM.

PDQRAP is described by Lin and Campbell [10]. Wu and Campbell describe how CBR (constant bit rate) channels are supported in DQSA [11]. Lin and Campbell [12] describe what a great job DQRAP does for packet voice using only "best effort".

3.5 The Long and the Short of It: How Does DQSA Spread Out a Switch?

The basic operation of the DQSA protocols assumes that the feedback of the minislots arrives back at the stations before it is time to launch the next transmission but the claim is made that DQSA allows a common channel to be shared by hundreds of stations spread over thousands of kilometers. This is accomplished by the use of interleaving, first described by Massey [8]. Interleaving requires the implementation of multiple protocol engines that operate in parallel, with each "engine" operating on a slot, the number of slots being equal to the round trip distance plus one. This technique will not work with the carrier sense protocols, e.g., CSMA/CD, where close to instantaneous feedback is required but is applicable to most other protocols.

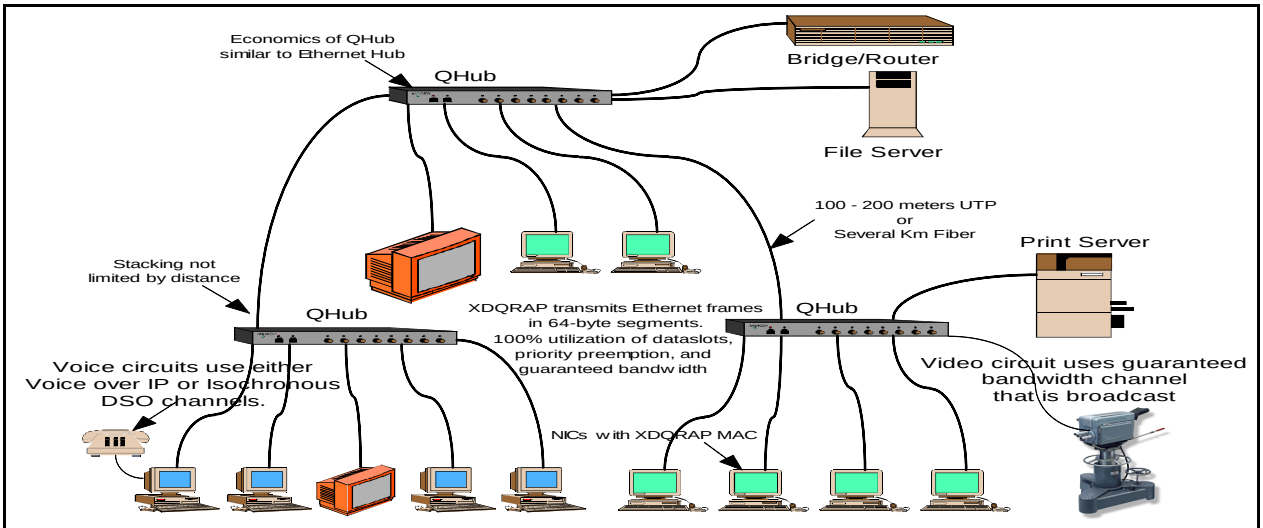

Figure 7. DQSA LAN Implementation

The interleaving technique developed for DQSA is highly optimized as compared with conventional interleaving. Wu developed a method that effectively reduced the number of parallel protocol engines from N to 1 thus making for efficient implementation. A DQSA DS1 system utilizing a 560-bit slot (64 bytes payload plus minislot overhead) results in a slot length of 363 µsecs. Referring to Figure 7, a WAN set up in the same manner as a LAN, (there will be more discussion of this network in a later section) the maximum distance from HQ in Chicago to the most distant stations is approximately 2000 miles. Using a propagation delay of 8 µsecs per mile the interleaving factor for this network operating at DS1 speed would be at least (8*2000*2/363) + 1 ≈ 100. The same network operating at DS3 speed would require an interleaving factor of about 2800. In both cases there could be added delay due to framing and switching in the underlying physical circuits that would require an increase in the interleaving factor. DQSA delay performance departs from the M/D/1 ideal as the distance increases. A rule of thumb is that the average access delay in a DQSA networking using interleaving is roughly 1.6 times the round trip delay at an offered load of 90%. DQSA interleaving is described by Wu and Campbell [9], Lee investigated the effect of interleaving on DQRAP and the benefits achieved using a global TQ as compared with separate TQs [20].

3.6 CDQ - Cascaded Distributed Queue Network

We take this opportunity to introduce what could be the final

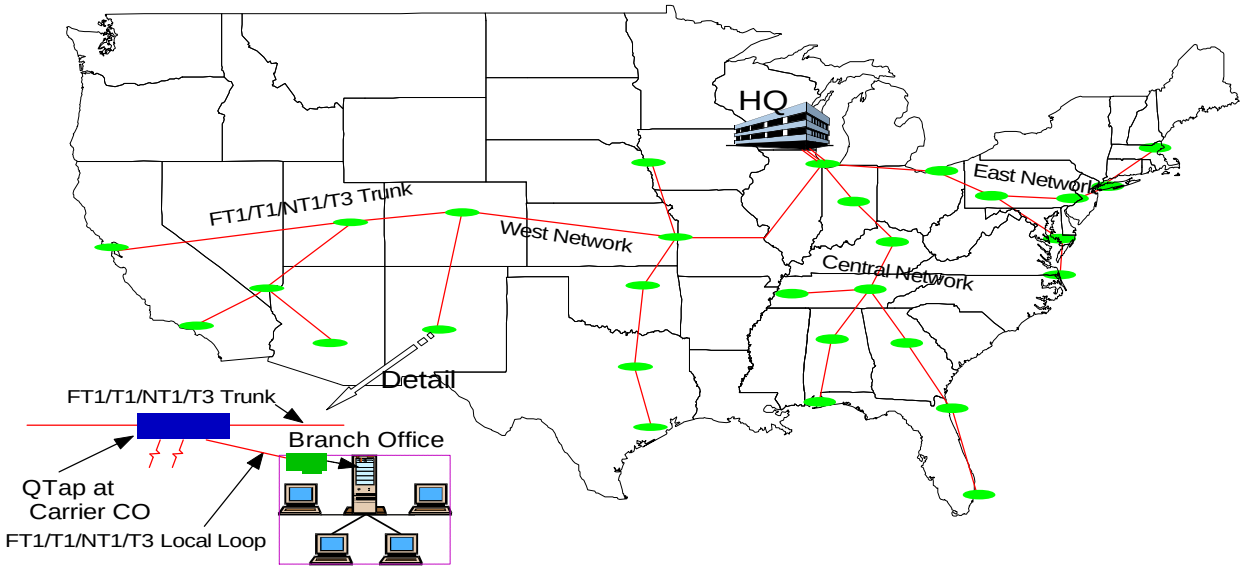

Figure 8. DQSA Network Functioning as Routerless WAN for Ethernet, IP, Frame Relay, or ATM

member of the DQSA family and one that fills a gap in DQSA. What is this gap? DQRAP and XDQRAP each allow a network of arbitrary size and speed to be organized using hub-and-spoke topology, as shown in Figures 7 and 8. But no one including the author suggests that "one big network" with a single hub can serve the entire country. Fedex, where business justifies, establishes satellite hubs such that parcels originating in and destined for the catchment area served by the satellite hub never leave that area. We introduce CDQ (Cascaded Distributed Queue Network) to provide this same function in a communications network. This permits what can be termed heterogeneous networks, e.g., the Internet, to be supported in a more efficient manner, i.e., a packet originating in Boston does not have to travel via Kansas City to reach New York.

The thinking behind CDQ was that the network would consist instead of a series of "satellite" DQRAP networks that are cascaded (thus the named CDQ). The idea was that when a request reached a hub in one segment, the request would be forwarded immediately with the hope that before the data reached the first hub, a position on the TQ would have been assigned thus permitting almost immediate onward transmission. The station at the hub would double as the end station for the next segment. Needless to say this approach did not work but once again a student, Chen-Hung Chang, worked out details that turned out to be far more complicated than originally envisaged.

The CDQ member of DQSA lands us with the grandaddy of all network problems, e.g., the potential for congestion at the nodes joining the segments. However, having the networks distributed across the segment rather than feeding them directly to a switch/router opens up new approaches to flow control -- wonderful things can be accomplished when total traffic consists only of "distant" traffic, i.e., from the next segment arriving via a single port plus "local" traffic that can be throttled very quickly.

The first stage of the research is complete but not yet published but here are some interim results. Extensive simulations were carried out of a network similar to that shown in Figure 9 excepting that there were twenty segments. Each segment had "a" = 100^{-} thus the network corresponded approximately to a network carrying ATM cells, plus overhead, from New York to San Francisco operating at something less than DS3 speed. Two types of simulations were carried out, one assumed a mix of traffic that added up to 90% average Poisson type load (at the origin) on a segment made up of 30% through traffic from N.Y. to S.F., 30% local to each segment, and 30% passing through multiple segments. The results showed no cells lost, delay of approximately two slots at each hub, maximum queue length of approximately 70 at each hub, and an overall transit time of approximately 1.3 times the actual propagation delay, e.g., some 35 ms for a cross-country trip. Other than propagation delay the major delay is access to the initial segment which as pointed out earlier is approximately 1.6 times the round-trip time for the segment at 90% loading.

The second set of simulations included extensive overloading of the system by introducing "bumps" in traffic on several intermediate segments that brought total offered traffic to well over 100% of capacity. TRAFC, the flow control program developed by Chang ensured that (a) priority traffic was not affected and (b) congestion problems are avoided, i.e., there is no cell loss other than that due to BER [22].

More research remains to be done on CDQ but the results so far confirm our belief that CDQ can successfully interconnect DQRAP networks and, using ATM cells to transport data, enable DQSA to satisfy most communications requirements. There is a discussion of how CDQ could be deployed in Section 4.3.

4 DQSA and the Real World

Let's now look at how DQSA might fare in real world applications. Potential applications range from parallel bus arbitration, e.g., PCI, to LANs, to MANs, to WANs, satellite networks, wireless networks, PCS networks, etc. We shall confine the discussion to LANs, WANs, a discussion of CDQ as applied to WANs, and a few other odds and ends.

4.1 DQSA in LANs and MANs

Figure 7 illustrates a DQSA LAN. At first glance it could be a typical Ethernet LAN but there are startling differences. Let's enumerate:

4.1.1 Distance: A connecting link can range up to several Km if fiber is used. There is in fact no upper limit on the distance of a link, the distance is limited only by the power budget, meaning that the configuration shown could, with the use of repeaters, represent a MAN as well as a LAN.

4.1.2 Speed: The Ethernet-like simplicity of the physical access method means that at whatever speed a bit can be generated for reading from/writing to a medium it is likely that the same chip logic can be used to implement the four-state FSM and binary counters that constitute XDQRAP. Thus a gigabit DQSA LAN with a radius of several kilometers could serve an entire campus.

4.1.3 QoS: A typical DQSA LAN will achieve over 90% utilization of the available bandwidth, i.e., the minislot overhead will typically be less than 10%. As described in previous sections, the three classes of service are easily implemented, i.e., best effort, controlled load, and guaranteed bandwidth. The diagram shows both a phone and a TV camera connected. In both cases the data can be compressed and packetized and then transmitted using either controlled load or guaranteed bandwidth. Both types of signals can also be transmitted in a basic, non-encapsulated format utilizing TDM-like slots. The latter case shows the versatility of DQSA in that bandwidth can be dynamically allocated so that a DQSA LAN could function as a conventional switch, e.g., in a 100 Mbps system a DS3 channel plus multiple DS1 channels could be supported

along with random traffic. Also the shared medium makes it trivially simple to multicast or broadcast the TV signal in any format to multiple receivers.

In section 1.1 we agreed with the statement that the Ethernet frame would be around for a long time but argued that its use would continue to be in the shared medium -- the DQSA environment. Our statement also applies to any enhancements to the standard that aim at implementing some type of QoS in Ethernet, an almost impossible task for shared Ethernet but possible for switched Ethernet. An example is the recently adopted IEEE 802.3ac extending the length of the standard IEEE 802.3 frame to accommodate amongst other features a priority capability. As described above DQSA will be able to utilize the extension in a much better fashion than can Ethernet.

4.1.4 Network Management: The shared nature of the medium means that the simplicity of management that was the hallmark of early Ethernet systems, is now possible even with say a gigabit DQSA LAN connecting hundreds of stations spread over kilometers. CDQ backbones can be utilized where it is desired to physically localize packets.

4.1.5 Economics: The complexity and projected cost of both the DQSA NIC and the QHub will be comparable to equivalent Ethernet components. The same physical media used to support Ethernet will be used in a DQSA LAN. The standard IEEE 802.3/Ethernet frame will be used and so little or no change is required to upper layer protocols. Any changes required will be desirable in that they will reflect support for a QoS not now possible with Ethernet.

4.1.6 Summing Up: DQSA achieves 100% utilization of available data slots, typically over 90% of total bandwidth. However, if great bandwidth utilization was important then when Token Ring arrived on the scene it would have immediately displaced Ethernet. The marketplace suggests that simplicity and the low-cost that simplicity usually brings to the table are the key factors. A rudimentary QoS, where required, is achieved with Ethernet, much to the chagrin of Token Ring advocates, by throwing bandwidth at the problem. But as pointed out above, a DQSA LAN provides simplicity along with its accompanying economic benefits as well as a QoS that even Token Ring cannot provide. And great utilization and distance insensitivity comes with the rations.

A proof-of-concept 10 Mbps DQSA LAN is in operation and confirms both theory and simulation results.

4.2 DQSA in Wide Area Networks

DQSA in WANs introduces a paradigm shift in communications since the use of shared access has not even been contemplated since the retirement in the 1980s of multi-point WANs utilizing 2.4 Kbps - 9.6 Kbps analog modems. We could now enumerate the benefits of a DQSA WAN in a manner similar to those presented in Section 4.1 about DQSA LANs. Instead we simply refer the reader to that section because a DQSA WAN is simply a DQSA LAN writ large. In the main we will concentrate our efforts on describing how a DQSA WAN can be implemented.

A DQSA WAN can be implemented as a distributed Ether-switch, distributed IP switch, distributed Frame Relay switch or a distributed ATM switch. The difference lies in the segmentation process as illustrated in Figures 2 and 3. Figure 8 illustrates a DQSA WAN consisting of three "spokes" emanating from a "hub" in Chicago. Each spoke contains subspokes that are joined to the main spoke via a simple QTap. The spokes are circuits that can range in speed from DS0 to DS3 and higher speeds.

A branch office is connected via a local loop to the spoke via the same "QTap" as described above and shown in the inset detail in Figure 8. The QTap fulfills the same function as the QHub in a DQSA LAN. It basically consists of an "and" gate to merge two digital signal paths -- what complexity there is in a QTap is devoted to the necessity of synchronizing the underlying DSx frames before merging. It is still probably more than an order of magnitude simpler than a router.

The inset detail in Figure 8 illustrates amongst other things that a single NIC is used to interface to the network. This NIC, effectively as simple and thus as inexpensive as a DQSA LAN NIC or Ethernet LAN NIC, provides in conjunction with the DQSA NICs at all the other sites all network control. Nothing else is required. Practicality dictates that there will be a network manager residing at HQ to monitor the allocation of guaranteed bandwidth, gather traffic measurements, etc. The entire network, spokes and sub-spokes, can be a single speed, e.g., DS1, or the main spokes could be say DS3 while the local loops or some of the sub-spokes can be lower speed, e.g., DS1.

Implementing DQSA over STM circuits actually offers a great advantage with respect to utilization -- the synchronization available in STM at the bit level means that two bits suffice for a minislot, e.g., a distributed ATM switch using DQRAP requires only a single byte (8 bits) to carry three minislots plus priority. This means a utilization of 424/(424+8) = 98.1%. The only disadvantage to implementing DQSA over STM, e.g., DS1, is that whereas a single and-gate is all that is required to merge two DS1 circuits, synchronization of the underlying DS1 frames is required before merging. Implementing DQSA on "dark" circuits increases the bandwidth used by minislots and dataslots in that guardbands, preamble, etc., are required, but the "joining" of two circuits is now accomplished with only an "or" operation. Thus a network implemented on dark fiber could be built using only passive taps and splitters, if the power budget is satisfied. This suggests that all the silicon devoted to creating and maintaining synchronous timing plus the DSx and OCx framing could be exchanged for bandwidth.

SONET (synchronous optical network) is an attempt to move away from conventional switches towards the economy and simplicity of shared access, i.e., fiber in a ring topology. SONET interfaces to the STM (synchronous transfer method) world employing OCx whose 125 microsecond frames result in a granularity of thousands of bytes, not conducive to effi-

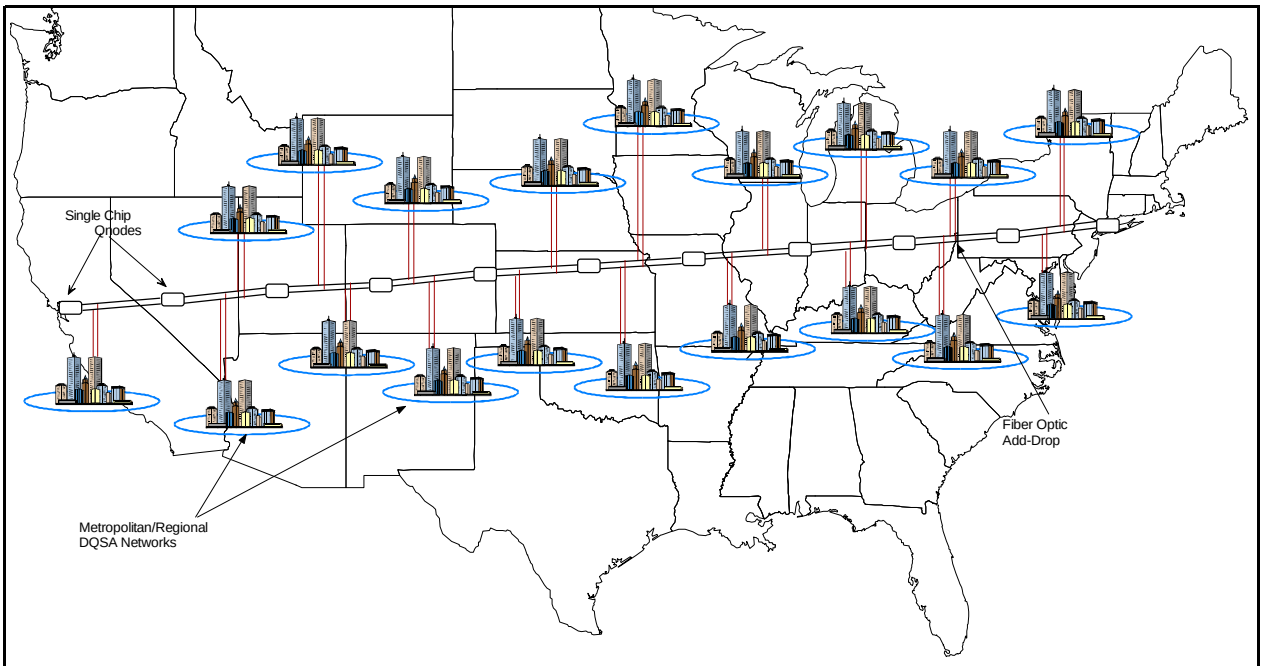

Figure 9. CDQ Cross-Country Distributed ATM Network

ciency with respect to shared access. In contrast a fiber-optic ring, possibly called ASONET (Asynchronous Optical Network), based on DQSA could switch individual ATM cells via add-drop optical taps in and out of the ring at gigabit speeds. Chang has simulated CDQ in a ring topology [22].

DQSA WANs can be implemented by the same corporations/ organizations that utilize existing WANs. The implementations could be private networks using leased circuits or public networks offering ATM or Frame Relay service. There is an ongoing move away from private networks to the utilization of commercial Frame Relay/ATM networks or the Internet, in both cases possibly using VPNs (virtual private networks). The illustration of Figure 8 shows a network that could be established using private leased circuits, or it could also represent a network established by one of the carriers, e.g., AT&T, using circuits dynamically assigned to this particular customer.

A DQSA WAN can provide a level of service with respect to delay, throughput, and guaranteed bandwidth not possible with existing networks. A carrier could actually provide a customer with a network consisting only of the carrier's circuits, i.e., no routers, and possibly charge a premium because of the level of the service guarantee. Guarantees with respect to uptime can be based solely on the reliability of the underlying circuits.

A proof-of-concept T1 DQSA system is operating and meets all expectations with respect to performance.

4.3 DQSA as a Backbone

CDQ lends itself to a backbone operation, whether it be for a single building, a campus, or continent-wide as shown in Figure 9 where conventional DQSA networks are connected as sub-spokes to the main spoke connecting two Qnodes. The number of networks/attachments is only limited by the capacity of the segment. The connecting points could be as simple as fiber-optic add-drop units if there is no speed change. Access to the network via lower-speed networks would utilize small buffers. Egress could present the classic problem of too much arriving at high speed for a lower speed sink but even this problem will be taken care of by the CDQ flow control mechanism. A CDQ network will operate as a distributed ATM switch with performance superior to conventional ATM networks in that upwards of 90% utilization can be expected without dropping cells excepting due to the BER of the line. Aside from a very good "best effort" performance and even better "controlled load" performance using PDQ priorities, guaranteed bandwidth channels can be established across the country that will support isochronous requirements.

A CDQ backbone could operate in the gigabit/s range and the segments can be formed into a ring. As with any of the DQSA protocols, if the power budget permits then an entire fiber network, excepting for the Qnodes, could be passive.

CDQ does make DQSA an architecture for all seasons.

4.4 DQSA and the Wireless World

All the good stuff said about DQSA applies to the wireless world, but we concentrate on just one facet. The availability of an efficient shared access mechanism opens up the possibility for true convergence of data and voice, i.e., packet telephony. An effective integrated system requires that voice be carried in

packets, but in most current voice systems the lack of an effective MAC has meant that some type of TDM (time division multiplexing) must be used thus making it very difficult to integrate voice and data. Implementing ATM in the wireless world is an obvious solution, but again the existing access methods are so inefficient that ATM cells must be "bundled" to achieve any sort of efficiency thus obviating much of the advantage of ATM. A packet telephony system based on DQSA is the solution.

We point out that a DQSA wireless system utilizing DQRAP/ATM or XDQRAP/IP provides all the goodies described in Section 4.1. Yes, there is the overhead problem for the minislots, a non-trivial problem in the wireless world with preamble, sync, and other requirements but this overhead will consume less than 20% of the bandwidth, yielding at least twice the throughput of any other system when the traffic is bursty. Lin demonstrates the performance of DQSA when subjected to voice traffic in [12].

4.5 Odds-and-Ends

An efficient shared access method quite literally opens up the entire world of communications but we must cut it off somewhere. We confine ourself to listing a few more potential applications along with a brief note on each.

4.5.1 DQSA and the Parallel Bus: All the bus arbitration schemes, e.g., ISA, EISA, PCI, etc., are shared access networks writ small. A short study of applying DQSA to a parallel bus arbitration scheme carried out by Wu indicates that it could have better performance characteristics than any system extant [21]. The main benefit is that it is truly distributed, e.g., a master arbitrator or controller is not required. One hundred percent of the bus bandwidth is available either being accessed randomly in single cycles or in blocks. As with all DQSA both priorities and reserved channels are available.

4.5.2 DQSA and VOD Mass Storage: VOD (video-on-demand) has been looming on the horizon for some time and perhaps one of these years it will be arrive. A DQSA network will be the ideal way to deliver a movie-on-demand whether it be in a single building or a metro area. However our interest here is in the storage mechanism itself. Low-cost storage devices, e.g., DVD, make feasible massive terastores with hundreds of DVD players where each player has a separate micro-controller that interfaces to a DQSA bus. Each controller will accept and respond to requests for some 30 - 40 hours of compressed video. The viewer after making a selection will interact directly with a controller and be permitted to fast-forward, freeze, pause, etc., everything that can be done on a VCR. DQSA is the only shared access method that would support a distributed system such as is described here.

4.5.3 DQSA and Cable TV: A natural for DQSA and yet it is on the sidelines except for the previously mentioned application of the concept of three minislots to the IEEE 802.14 standard. However the fact that two "standards", MCNS and IEEE 802.14, exist before there has been any significant deployment of any cable access indicates that neither is optimal. The author did discover that one should not attend a Standards dance unless one has a partner with clout.

4.5.4 DQSA and xDSL: One of the more exciting applications for DQSA is that it will permit clusters of DSL (digital subscriber loop) lines to be treated as a single network. DSL allows up to several megabits/s to be transmitted over the ordinary telephone wires serving homes and businesses. The power of DQSA means that true integrated services with packet telephony could be offered over the ubiquitous copper pairs. The DQSA hub could in turn be connected to a DQSA backbone and thence to other neighbourhoods and communities thus establishing a network for all communications services that bypasses the conventional switched system.

4.5.5 DQSA and RFID: The efficiency of DQSA in resolving contention means that it is ideal for RFID (radio frequency identification) systems. We present an example that also shows the efficiency of the QCR (queueing contention resolution) subsystem.

Assume there are 10,000 items in a warehouse, each with its RFID implementing DQRAP. Assume that a relatively low-speed, low cost, radio system utilizes a time slot of 60 milliseconds which includes outbound beacon and feedback plus the inbound minislots and dataslot. When inventory is to be taken, the master unit issues a beacon to start the process and then repeats it every 60 ms. All 10,000 items respond to the first beacon. But after the master controller sends out the feedback along with the second beacon the DQRAP contention resolution process kicks in and the number of responses drops to approximately 3333 items in the second round. In succeeding rounds the number of responses will drop by approximately two-thirds each time: 1111, 370, 123, 41, 14, 5, and then 2. Thus after nine cycles, approximately 540 ms, the first couple of items will be successful in the minislots and thence move to the data transmission queue. The remaining 9,998 items will then resolve their contention in parallel with the ongoing data transmission and join the distributed data transmission queue, The entire inventory will be complete in little more than 10,000 x 60 ms = 600 seconds.

O'Connell investigated the contention resolution capability of DQRAP by simulation of situations where up to 100,000 arrivals contend in a single arrival period [23].

4.5.6 Wrap-up: We will stop listing potential applications at this point since by now we hope that the message is clear: once shared-access is viewed as a viable alternative then virtually all communications applications are subject to rethinking with respect to how best they can be implemented.

5 Conclusions

The rise of the Internet has been dramatic even by the standards of an industry that over the past thirty years has become accustomed to individual segments literally exploding. The growth will not slow down since more and more services are

moving to the Internet, the most significant of which are voice and video. Naysayers opine that the Internet will never have the capacity nor capability with respect to quality to make major inroads on conventional voice traffic or to provide reliable video, but the process is under way. The current and expected growth prompted one eminent personage to state publicly what many felt in private, that the Internet would suffer a meltdown. This has not happened as yet in the main because silicon and bandwidth are being thrown at the Internet in sufficient quantities to keep ahead of the curve. However these are brute force methods and will not be able to keep pace with demand forever -- the time frame of the Internet.

The author argues that the cost of bandwidth is dropping faster than the cost of silicon switching, thus the optimum strategy for staying ahead of the demand curve is to, where possible, substitute bandwidth for silicon. The shared access approach to communications is effectively just that, trading bandwidth for silicon. We also argue that the major benefit of this approach will accrue not from less silicon but from the lesser network management requirements.

The ability of DQSA to support isochronous traffic means that DQSA distributed switches could support both random access and conventional DSx circuits during any transition period. DQSA allows single ATM cells to be switched at any speed over any distance on either synchronous circuits that utilize DSx or OCx framing or "dark" circuits where each transmission is self-contained in that preamble, sync, etc., are required. STM circuits simplify the implementation of DQSA because the underlying synchronization allows minislots to be only a few bits long and permits back-to-back transmissions by different stations. The author estimates that implementing DQSA on "dark" circuits increases overhead by about 10% to satisfy guardband, preamble and sync requirements but, the payoff is elimination of out-of-band STM overhead. The suggestion was made in section 4.2 but here is the question: "Is it worth giving up say 10% of the bandwidth if it means that DSx and OCx framing could be eliminated?" The two main existing communications mechanisms, STM switched-circuits and routed/switched packets, are exemplified by the products of Lucent and Cisco. These two spend much time eyeing each other warily and attempting, as voice becomes interchangeable with data, to poach in each other's territory. Perhaps each should instead be looking over their respective shoulders.

The performance of contention-based shared access methods, until now, has been poor but there is no law of nature that states that performance which approaches that of an ideal system is not possible. Summing up:

- **A set of ten criteria is presented that describes the ideal shared access system.**
- **A set of shared access protocols that constitute the DQSA family is presented. It is further shown that the DQSA protocols satisfy or come close to satisfying the ten ideal criteria.**
- **It is shown how the DQSA protocols can be deployed in a number of communication applications, a number sufficient to suggest that shared access can play a major role in all aspects of communications.**

We conclude by saying that, contrary to reports, shared access is alive and well and will be marching briskly into the new millennium.

References

[1] W.Stallings, *Data and Computer Communications, Fifth Edition*, Prentice Hall, 1997.

[2] IEEE 802 Web Site. www.stdsbbs.ieee.org/groups/802/overview.html

[3] Specifications available at *www.cablemodem.com*

[4] W. Stallings, *Local and Metropolitan Area Networks, Fifth Edition*, Prentice Hall, 1997.

[5] W. Xu and G. Campbell "DQRAP - A Distributed Queueing Random Access Protocol for a Broadcast Channel", presented at SIGCOMM '93, San Francisco, September 14, 1993. *Computer Communication Review,* Vol 23, No. 4, Oct 1993, pp. 270-278.

[6] C.T. Wu and G. Campbell, "Extended DQRAP (XDQRAP): A Cable TV Protocol Functioning as a Distributed Switch", *Proceedings of 1st International Workshop on Community Networking*, July 1994, San Francisco.

[7] C.T. Wu and G. Campbell, "DQLAN - A DQRAP Based LAN Protocol", *Proceedings of the 1st Workshop on High-Speed Network Computing, 9th Int'l Parallel Processing Symposium*, Santa Barbara, CA, April 1995.

[8] J. L. Massey, "Collision resolution algorithm and random access communications," in *Multiuser Communication Systems*, G. Longo Ed. N.Y.: Springer-Verlag, 1981. pp. 73-137.

[9] C.T. Wu and G. Campbell, "Interleaved DQRAP With Global TQ", *DQRAP Research Group Report 94-4*, 1994.

[10] H. J. Lin and G. Campbell, "PDQRAP - Prioritized Distributed Queueing Random Access Protocol", *Proc. of 19th Conference on Local Computer Networks*, Oct. 1994, pp 82 - 91.

[11] C. T. Wu and G. Campbell "CBR Channels on a DQRAP-based HFC Network", *SPIE '95 (Photonics East)*, Philadelphia, PA Oct. 1995.

[12] H.J. Lin and G. Campbell "Using DQRAP (Distributed Queueing Random Access Protocol) for Local Wireless Communications." *Proceedings of Wireless '93*, July 14, 1993, pp. 625-635.

[13] M. Zhang and G. Campbell, "A Performance Analysis of Distributed Queueing Random Access Protocol", *DQRAP*

Research Group Report 93-1, 1993.

[14] M. Miramica and G. Campbell, "Robustness Analysis of the DQRAP Protocol", *DQRAP Research Group Report 93-6,* 1993.

[15] Paul Desmond, "Happy Birthday Ethernet", *Network World*, May 18, 1998, p 40.

[16] C. Bisdikian, B. McNeil, R. Norman, R. Zeisz, "MLAP: A MAC Level Access Protocol for the HFC 802.14 Network", *IEEE Communications Magazine,* pp. 114-121, March 1996

[17] C. T. Wu and G. Campbell, "The Extended DQRAP (XDQRAP) Algorithm", *DQRAP Research Group Report 94-3,* 1994.

[18] M. McPheters, "Approaching Reality with a Finite Network Node Model Analysis of DQRAP (Distributed Queueing Random Access Protocol)", *Ph.D. Dissertation*, Illinois Institute of Technology, 1993.

[19] R. Khasawneh, "A DQRAP Based Residential-Metropolitan Area Network Utilizing Cable TV Facilities", Ph.D. Dissertation, Illinois Institute of Technology, 1993.

[20] Byung. Lee, "Interleaving and DQRAP (Distributed Queueing Random Access Protocol) Based Networks Without Immediate Feedback", Ph.D. Dissertation, IIT, 1996.

[21] C.T. Wu and G. Campbell, "DQRAP as a Parallel Bus Arbitrator", Unpublished, IIT, 1994.

[22] C.H. Chang, "A Cascaded Distributed Queue Network", *Ph.D. Dissertation*, IIT, 1998 (Not Published)

[23] W. O'Connell, "Can Shared Access Networks Adequately Support High-Performance Computing? Re-Visiting MAC Protocols", *Proceedings of the 1st Workshop on High-Speed Network Computing, 9th Int'l Parallel Processing Symposium*, Santa Barbara, CA., April 1995.

[24] Microsoft Professional Developers Conference, San Diego, Sep 22-26, 1997.

[25] J. I. Capetanakis, "Tree Algorithm for Packet Broadcast Channels", *IEEE Transactions on Information Theory*, Vol. IT-25, no. 5, pp. 505-515, Sep. 1979.

[26] F. Koperda and B. Lin, "Providing CBR Service over XDQRAP", IEEE 802.14/95-147, Scientific Atlanta, 1995.

[27] Information available at www.homeRF.org.

[28] L. Kleinrock, "On Queueing Problems in Random-Access Communications", *IEEE Trans. Inf. Theory*, Vol IT-31, No. 2, March 1985, pp 166-175

NB: All DQRAP Research Group Reports plus simulation programs, etc., and this paper are available at *www.iit.edu/~dqrap*.