Ether2
presents
Routing Without Routers and Switching Without Switches

Graham Campbell
Chief Science Officer
Ether2 Corp.

Abstract

DQSA (Distributed Queue Switch Architecture) is a switching technology that allows the basic circuit-switched synchronous infrastructure to efficiently carry packets. This paper describes historical events, the current communications milieu, and how Ether2 technology will enable the true convergence of voice, video and data.

Introduction

Given the remarkable expansion of the Internet, the average person would assume that the world appears to be on the verge of full convergence wherein all voice, video and data are delivered over a common communications infrastructure. But this is not the case and convergence, as desirable as it is, is not likely to be achieved with current technology. To explain why, we must go back to 1876, the year that Alexander Graham Bell ushered in a new era of personal communications with the telephone. Of equal importance to the phone it self, was the switch Bell implemented that enabled any two of those few first customers to talk with each other. That first switch, manually controlled by an operator, can be described in modern day terms as a circuit switch, i.e., the operator was asked beforehand to physically establish the circuit and once it was established, the latency or time for information to traverse the circuit never varied. This feature meant that the switch could also be termed "synchronous".

This type of switch remained unchallenged up to the 1960s, through the introduction of electronic amplification, long distance calling, direct dialing, and even to the introduction of the first digitized T1 circuits. The latter led to the term STM (Synchronous Transfer Method) to describe the vast infrastructure that was developed to support the world's telecommunications.

The computer arrives... Just as with people, computers had the need to "talk" to each other. However, there was a problem: even the electronically controlled synchronous switches could still take what appeared to be an eternity compared with the time it would take to transmit the actual message. The alternative was to establish a circuit and leave it connected. This worked when there were two points that generated sufficient traffic, but most often it resulted in an under-utilized circuit. STM could be described as both economically and performance challenged when it came to supporting computer communications.

A solution to this problem occurred to Leonard Kleinrock, Paul Baran and Donald Davis, who individually conceived the idea of packaging the information in packets and to then send these packets, each of which contained a destination address, to a "packet" switch that would look at the address and then forward the packet out of the port that would carry it to the desired destination. This type of "packet switching" would enable one computer to transmit a minimal amount of information to another computer without having to first establish a circuit. The circuits that interconnected the packet switches would be "permanent", and would be better utilized since they could now transport traffic belonging to multiple users. The switch itself would be better utilized since it would no longer have under-utilized circuits. ARPANET was eventually established as a full scale test of the concept, and proved so successful that it progressed into today's Internet.

By the late 1970s, packet switching showed so much potential that communications research establishments around the world, and the telephone companies that owned them, collectively decided that all communications could be supported by using that technology, i.e., convergence. The specific technology selected was ATM (Asynchronous Transfer Method). All traffic including voice, video and data would be segmented into 48-byte chunks and transmitted in 53-byte cells. It would be universal in scope.

Some 20 years later most people, including those technically trained, think of cash machines when the term ATM is used, an indication that ATM was not sufficient to the task. The problem was that while packet switching is good at transmitting files and intermittent data, it did **not** prove satisfactory for carrying real-time data, e.g., voice, video, and data that had stringent QoS (Quality of Service) requirements. The very efficiency of the packet switch -- its ability to "collect" packets from various sources and direct them to sundry destinations -- is also the reason for the "asynchronous" in ATM. If two or more packets arrive simultaneously destined for the same output port, then all but one of the packets will be queued for later transmission. Thus, unlike a circuit switch, <u>the latency of a path between two points via a packet switch does not remain constant.</u> Furthermore, queuing theory predicted that the queues would grow to infinite length when the offered traffic reached 100% of capacity. Practice confirmed this, and so, it is accepted that in packet networks there will always be a loss of packets because of buffer overflow.

The telecom industry has slowly but surely abandoned the concept of universal ATM. Today, we find that the world's telecom infrastructure is almost totally circuit-switched as it was in the time of Bell with an infrastructure that includes fiber, optical switches and SONET rings, with operating speeds in the multi-gigabit per second range. But what everyone is aware of today, is that virtually all information is placed in packets before being sent to a destination. So, how does this coincide with the above statement that "circuit-switching is economically and performance challenged with respect to packets"? Cisco and other router manufacturers came to the rescue.

Virtually every business in the world, and an increasing number of residences, now have routers that connect to even bigger routers that are in turn connected to other routers. All these routers are connected using the existing telecommunications **synchronous** infrastructure, ranging from the local copper loops to cross-country multi-gigabit fibers. The synchronous infrastructure inadequacies with respect to packets are addressed by superimposing thousands of packet-switched networks over the circuit-switched infrastructure, thus providing a packet interface to the user.

***The end result is that there are two distinct sets of networks*:** (1) Thousands of synchronous (circuit-switched) networks that encircle the globe; and (2) sitting on top of these networks, tens of thousands of asynchronous (packet-switched) networks that again encircle the globe, and through the standards of IP (Internet Protocol), provide worldwide packet communication.

The Problem

Packet networks on top of circuit-switched networks do not provide the seamless environment required for true convergence of voice, video and data. Major users still directly utilize synchronous circuits, in parallel with their packet networks, to satisfy communications requirements.

Why? The inadequacies that prevented ATM from providing universal service are still present in today's packet networks. Consider Voice-over-IP (VoIP)… It is expanding rapidly and will eventually dominate phone service, yet it is now, and will always be, subject to serious quality problems. VoIP works well only when all intermediate routers on the path taken by the VoIP packets are not near capacity. For almost all technical systems, e.g., power, transport, etc., a major responsibility of management is to continually strive for the highest possible utilization, which includes short periods when utilization could reach 100%. In contrast, a major responsibility of the managers of router-based networks is to ensure that they remain under-utilized. If there is just a hint that the capacity of a router buffer might be exceeded, it is necessary to start **discarding** packets. This does not have much impact on file transfers but with real-time dependent traffic like voice, the results are disastrous.

Inefficiencies also exist in all wireless networks… Cell phone systems use duplex synchronous circuit established between a base station and a cell phone. Traffic flows in only one direction at a time, leaving the channel in the other direction unoccupied. There is additional unused bandwidth in the pauses between words, and research has shown that utilization is just over 40%. In wireless systems where "naked" packets are transmitted, as in IEEE 802.11 compliant systems, the lack of an efficient MAC (Media Access Control…variations of the original Ethernet are used) means that utilization is rarely over 40%. Additionally, it is possible to offer QoS (Quality of Service) only by lowering the utilization of the circuit.

The **ideal** infrastructure for the Internet, and for all communications, is one where 100% of the data capacity of the network is utilized when subject to surges in packet traffic beyond full capacity, but with <u>no loss of packets</u>. At the same time, maintaining satisfactory service for VoIP and other real time traffic is paramount.

The Solution

Ether2 offers the enabling technology for the above described "ideal" network in the form of DQSA (Distributed Queue Switch Architecture), a technology invented at the Illinois Institute of Technology by Professor Graham Campbell and his students. DQSA is based on a MAC that provides close to ideal performance, in that it permits an arbitrary number of users to efficiently share a communications channel regardless of the distance between the stations, the topology of the network, or the speed of transmission. A highly efficient MAC has obvious use in a LAN or wireless application, but what is not so obvious is that when implemented directly on a *synchronous* circuit-switched infrastructure, it satisfies the criteria mentioned above, i.e., <u>congestion-free packet transmission</u> intermixed with TDM-like channels.

DQSA technology is based on a family of highly efficient access methods that utilize the concept of a station transmitting a request for bandwidth that instead of being sent to a central controller is instead sent to all stations, including the sender, on the network. If the sending station subsequently receives the request intact the assumption is made by the sender that the bandwidth has been allocated, and the sender joins a distributed transmission queue to await its turn to transmit. All other stations have received the same request and so they all update their copy of the state of that same distributed transmission queue.

What happens if two or more stations request bandwidth simultaneously? This is at the heart of the efficiency of DQSA; collisions are resolved at a rate faster than the transmission capacity of the line so utilization of the output channel is 100%. The two or three mini-slots used to accomplish both the requests for bandwidth and the collision resolution consume between 3% and 15% of the available bandwidth. This is an extremely low price to pay to achieve full utilization and QoS, and to eliminate the need for routers.

Readers familiar with Ethernet will recognize the use of collisions to alert stations to the fact that they must "try again." However it is DQSA's unique method of "trying again" that overcomes all the inadequacies of the original Ethernet and opens up an entirely new way of switching for all communications. What was overlooked in the many years of searching for a "better" Ethernet is that a MAC that overcomes Ethernet's limitations does not result in only a better Ethernet, but can also provide the basis for an entirely new method of switching that eliminates congestion and provides QoS. The routers and switches are **removed**, not replaced, thus greatly reducing initial capital cost and ongoing maintenance costs. The advantages of Ether2 in specific environments are presented in a series of companion white papers.

Detailed descriptions of the algorithms plus simulations, etc., can be found in the documents at www.iit.edu/dqrap.

Potential Applications of Ether2 Technology

Ether2 technology is the basis for designing and building networks that address specific applications that include:

- High Performance Computing (HPC) as implemented using clusters of processors.
- Storage Area Networks.
- Server farms.
- Metropolitan Area Networks.
- Local Area Networks.
- Wide Area Networks.
- Wireless networks that will provide services presently addressed by IEEE 802.11 and IEEE 802.16.
- Cell phone service.
- Satellite services (LEO, MEO, and GEO).
- RFID.
- Backplane and bus structures.
- Virtually all other segments of communications.

In all these applications Ether2 will take advantage of an already existing synchronous infrastructure or will implement a synchronous backbone that will then allow the commingling of packets and TDM-like channels with efficiency not presently available from any network. The most startling potential implementation is that telecom carriers could provide a continent-wide service without the need for routers. This would offer a greatly improved packet service, along with TDM capability, using only their already installed synchronous infrastructure.

Conclusion

DQSA is an access method that allows an arbitrary number of users distributed over arbitrary distances to efficiently share a communications channel implemented on virtually any medium, topology and at any speed. This technology is the basis for a novel form of switching that will allow the reorganization of all networking that will in turn allow true convergence of voice, video and data. Routers and conventional packet switches will be phased out so that the basic circuit-switched infrastructure that served society so well for well over a hundred years can provide an even better level of service in the coming centuries.