Low-Cost Implementation of Ethernet-based Services over CATV Interconnect Networks

Robert W. Harris
ADC Broadband Communications

Michael Coden
ADC Codenoll

Abstract

A low cost solution for implementing Fast Ethernet data transport over local backbone and long distance regional CATV networks is presented. The approach allows switching and transmission of standard Ethernet and Fast Ethernet data in its native mode over any size ring network with no limitation to distance. Furthermore, this method is independent of both data format and protocol allowing service providers to deliver fully transparent Ethernet-based services to their customers. This approach, in effect, extends the functionality and low cost of the local area network (LAN) into the Wide Area Network (WAN).

This paper describes the theory of operation of a new "ring-based" Fast Ethernet switch and transport solution known as EtherRing™. Implementation of EtherRing in cable data modem applications is presented as an alternative to current data transport solutions. EtherRing makes use of the data networking concept known as "Route Once-Switch Many" and is extended into the CATV network environment as a method for centralizing network routers. This solution, when used in conjunction with EtherRing transport technology, is shown to provide a highly scaleable and efficient approach to data modem service delivery resulting in network cost savings between 50% to 80% over alternative data transport methods. Additional CATV data transport applications are presented including; advertisement insertion transport, set-top box access and control transport and general purpose LAN service delivery.

INTRODUCTION

Computer networks have become commonplace in large and small businesses, universities, government facilities and other organizations. With the advent of Internet access, advertisement insertion, video on demand, IP-video, IP-telephony and other emerging CATV applications, computer networking principles and techniques are fast becoming an integral part of the CATV network. By definition, computer networks allow a number of users or devices to share data and resources such as file servers, data storage systems, printers, switches, routers, modems and other peripherals. The CATV industry will need to interface to these and other IP-based devices such as cable modems, MPEG encoders and network management systems over their networks.

During the past 20 twenty years, a number of methods for connecting computer devices within a network have been devised. The IEEE 802 and ANSI committees have developed a number of standards used for various computer networks. These same standards and practices will be followed as CATV networks deploy data-based services.

Computer networks are often classified in two categories - the Local Area Network (LAN) and the Wide Area Network (WAN). The LAN is characterized as a shared medium where all devices share the network bandwidth. The shared medium can consist of coax, fiber, twisted pair (Category 3 and 5) or any combination of the three. LAN bandwidths typically range from 10 Mb/s for Ethernet to 100 Mb/s for Fast Ethernet and FDDI (Fiber Distributed Data Interface).

The LAN is typically used to transport data between network devices over relatively short distances such as within a building or groups of buildings as in a campus. The maximum distance between devices on a LAN is generally (but not always) limited to several km due to the protocols used within the LAN.

The WAN is used to connect two or more LANs separated by much larger distances - several km to thousands of km. WAN network devices have the function of re-packetizing LAN data packets and provide routing protocols to determine the destination of the transmitted data. The WAN can employ standard telecommunication trunking such as T1 circuits, fractional T1, DS3, ATM, etc. The WAN circuit may be a private link or supported on the publicly switched network. The main advantage of WAN functionality is its support of long distance transport. This, however, comes at the expense of the use of complex protocols, potential traffic congestion and high costs.

There are three basic types of protocols used for transmitting data in LAN networks: Ethernet, Token Ring (or FDDI) and encapsulation. Each of these network protocols have specific advantages and disadvantages and are described below. Following this is a discussion on packet switching and its advantages over routing.

This paper will then introduce a new technology concept, called EtherRing™, presented as a solution for significantly reducing the cost and complexity of data transport within the CATV network. Following this discussion, the concept known as Route Once - Switch Many is presented as a method for consolidating network routers at a single headend location leading to additional savings in network costs. Several CATV network applications incorporating EtherRing are also presented.

DISCUSSION OF CURRENT NETWORK PROTOCOLS

To understand the functional capabilities and limitations of existing network protocols, a brief description of three popular protocol solutions is provided below.

Ethernet

Ethernet is basically a broadcast protocol where its main advantage lies in it simplicity. This allows Ethernet to be implemented with less costly hardware and software. The main drawback with conventional Ethernet is that there are limitations on the physical distance that the network can cover. 10Base-T is limited to 4.5 km while 100Base-TX is limited to 1 km.

Despite its distance limitations, Ethernet has become the most common protocol for LANs due in large part to its low cost, ease of use, low complexity and support of relatively high bandwidths (10 Mb/s for 10Base-T and 100 Mb/s for 100Base-TX, commonly referred to as Fast Ethernet). According to one market report, [1] more than 80% of all networked devices are connected via Ethernet.

For the purpose of this paper, the term Ethernet includes the entire class of Carrier Sense Multiple Access with Collision Detection (CSMA/CD) protocols covered by the family of computer industry standards known as IEEE 802.3. A number of good tutorials covering on the operation and functionality of Ethernet are available from a variety of sources [1,2,3].

Token Ring and FDDI

Token Ring and FDDI, as described in computer industry standards IEEE 802.5 and ANSI XT3.9, respectively, provide the distinct advantage that data can be transported over greater distances relative to conventional Ethernet. Further, Token Ring and FDDI provide virtually equal bandwidth throughout the network. The main disadvantage of both is their use of complex protocols. As a result, Token Ring and FDDI hardware and software costs are significantly higher relative to Ethernet. Further, the loss of a token - which determines what network device can send its data - can cause significant network transport delays as token recovery routines are performed leading to loss potential of critical data.

Encapsulation

Encapsulation protocols have been developed to allow Ethernet packets to be transported over the longer distances covered by the WAN. In such protocols, the entire Ethernet packet is placed within another type of packet with its own header including additional addressing information, protocol information, etc.

These protocols typically suffer from the problem that they may require special higher level protocol information to be included in the data field of the Ethernet packets for the purpose of directing routers in the network. This has the effect of limiting the types of data packets that can be handled and places a significant processing burden on both the network devices generating the packets and the routers used to transmit and receive the packets between the various Ethernet network segments.

These additional protocol elements and restrictions typically require expensive hardware and software be added to an otherwise inexpensive Ethernet network. Further, such protocols typically require the use of manually created address tables for the routers.

PACKET SWITCHING

A packet switch, often referred to as an Ethernet switch, or just a "switch", is a multiport bridge that simply forwards packets from a device connected on one port to a device connected on another port. The forwarding decision is based on the destination MAC (Media Access Control) address at the head of the packet. A switch will ignore a packet that is destined for a device located on the same port as the source device. Forwarding decisions of the switch are based on link-layer (layer 2 of the open systems interconnection, or OSI, model) information. Routers, on the other hand, forward packets based on network-layer (layer 3) information. Key here is that switches do not modify packets as they pass through, whereas routers must change the packet to include the MAC address of the router at the next-hop and may also increment a hop count field [4].

Each port on a switch may be connected directly to network devices such as servers, routers, printers, PCs, etc. or be connected to completely different LANs operating independently and simultaneously. Only those packets that need to pass from one LAN segment to another are forwarded by the switch. As a result, a multiport switch will often increase the overall bandwidth of a single shared LAN by many times.

Layer 2 switches provide a simple and elegant way to increase the aggregate bandwidth of the network and are often less expensive and faster (higher throughput) than routers. Switches are simpler because they operate at layer 2, do not modify packets and do not require complex routing protocols like Routing Information Protocol (RIP) or Open Shortest Path First (OSPF). Routers often use proprietary and complex routing protocols which typically have to segment and reassemble packets on the fly. Moreover, routers are generally protocol specific requiring different software for each protocol used and require constant maintenance of the routing tables as network parameters change. As a result, the cost to maintain routers on a monthly and yearly basis can be significant.

Switching, on the other hand, is entirely self-learning. This means that the switch automatically "learns" which port (or LAN segment associated with a port) each device is connected to - even if the device is moved to a different port or LAN segment. Each switch port records the source address of every packet, as it receives the packet, in a memory table for that port.

Further, when a packet is received at a port of the switch, the destination address of the packet is compared to the memory tables for the other ports of the switch. When a match is found for the destination address in the tables for one of the ports, the packet is switched to and sent out that port. With broadcast packets, the packet is broadcast to the other ports on the switch but never back to the original receiving port.

Likewise, "multicast" packets using specially reserved destination addresses will be broadcast to a selected group of devices. Therefore, the switch does not require special "management" of memory tables since the process is performed automatically. As a result, the switch is effectively maintenance free compared to router functionality.

Layer 2 switches offer key advantages over routers in other ways that affect network performance and cost. Because they operate at layer 2 (switching by MAC address) instead of routing by network address at layer 3, switches can operate at higher speeds contributing to lower latency and higher throughput. Further, routers are generally 3 to 10 times the price per port of layer 2 switches [5].

One reason why switches provide higher throughput, lower latency and are lower in cost than routers is because the switching functions are often implemented entirely in VLSI (very large scale integration) rather than performed by software running on an expensive high performance processor. Consequently, both the initial cost and the maintenance cost of layer 2 switching will always be less expensive than routing.

Given the above attributes of layer 2 switching, it's not surprising that organizations often deploy layer 2 switches to "front-end" their routers, offloading some of the traffic from the router and utilizing much less expensive switch solutions for switching traffic between LAN segments [5]. This concept, used often in the traditional LAN environment, is commonly referred to as "Route Once - Switch Many". In the CATV environment, an extension of this concept can be employed leading to a powerful and economical solution for centralizing most or even all network routers at a main server or headend location and distributing the data service(s) via layer 2 switch technology. Route Once - Switch Many as related to CATV-based data transport applications is discussed in further detail later in this paper.

However, before one can implement Route Once - Switch Many techniques within the CATV network, a number of significant challenges need to be overcome. This is because CATV networks and LANs have significant differences in network topology, network distances and necessary protocols and applications to support. Overcoming these challenges led directly to developing a new technique to deliver layer 2 functionality - with its low cost, low maintenance and higher throughput - over the wide area CATV network while maintaining complete compliance with data networking standards. The solution is known as EtherRing.

EtherRing - A NEW CONCEPT FOR ETHERNET TRANSPORT

Because of the clear advantage of Ethernet switching, there is a strong to desire to extend this functionality and flexibility into larger scale applications. However, transmission over the regional CATV network effectively means transmission over a WAN. As a result a number of challenges exist in delivering native-mode Ethernet over the WAN.

Challenges to Overcome

Where Ethernet topology is basically broadcast in nature, telecommunication networks, such as those used in CATV regional headend consolidation, are typically configured in a ring. In standard Ethernet if a packet is not used, or is a broadcast packet, it travels once through the LAN then fades away into the "Ether". On a ring, the packet will come back to the originating point and can continue circulating the ring indefinitely. Figure 1 summarizes the challenges of transporting native mode Ethernet over telecommunications networks.

Challenge	10BaseT Ethernet	100BaseTX Fast Ethernet	Telecom Systems
Topology	Broadcast	Broadcast	Ring
Maximum Time Delays	51.2 µs	5.12 µs	1,000's of µs
Maximum Distances	4.5 km	1 km	1,000's of km
Protocols	Any and All	Any and All	Any and All

Figure 1. *Challenges for Long Distance Ethernet Transport*

Standard Ethernet is unable to cover the long distances (potentially >1,000's of km) associated with ring telecommunication networks - this has been the job of the WAN. Because of the collision detection scheme in Ethernet, the maximum time an Ethernet packet can take to traverse an entire network is 5.12 µsec for 100Base-TX (51.2 µsec for 10Base-T). This maximum time restriction translates into a distance restriction (based on velocity of propagation through the medium) of 1 km for Fast Ethernet and 4.5 km for 10 Mb/s Ethernet.

The alternative currently in use is routing, which, although powerful, can have a number of limitations listed earlier such as; lower throughput, protocol dependence and higher hardware and maintenance costs relative to Ethernet switching. Ideally, a data transport system would be protocol independent so that the operator can connect Novell Netware LANs, TCP/IP LANs, Netbeui LANs, digital set-top box controllers, network management systems and any other device(s) from any manufacturer using any other protocol all on the same transport platform without requiring expensive software and routing table maintenance.

EtherRing Functional Overview

Through simple modifications of the Ethernet standard [6] (while still maintaining IEEE 802.3 compliance on the local ports) EtherRing allows native mode Ethernet packets to be transported via a ring rather than the typical star configuration.

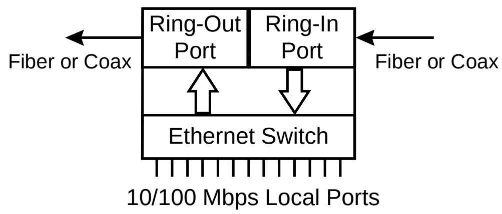

Figure 2. *Functional Block Diagram of EtherRing*

The EtherRing platform consists of three primary components; a Ring-In port, a Ring-Out port and an Ethernet switch. The Ethernet switch is just that - a standard 10Base-T/100Base-TX switch that possesses all of the normal features and functionality found in Ethernet switches. Connected to the switch are Ring-In and Ring-Out ports that may have either a fiber optic or coaxial interface.

Distance Solution. The maximum-distance / collision-time-domain problem is resolved by eliminating collisions altogether on special "Ring-In" and "Ring-Out" ports. By connecting only one Ring-Out port to each Ring-In port, (and vice versa) we ensure that each Ring-In port will only receive packets from one Ring-Out port (see figure 3). Therefore, there are no other possible transmissions to "collide" with the transmissions occurring (simultaneously all around the ring) on each connected pair of Ring-out and Ring-in ports.

The key point is that there can be no collisions on a Ring-In or Ring-Out port since the Ring-In port on any switch will only receive packets from the Ring-Out port of one and only one switch. Therefore, the only time limit on data transmission is that of the higher level software protocols, whatever that may be, but typically it ranges from half a second to several seconds.

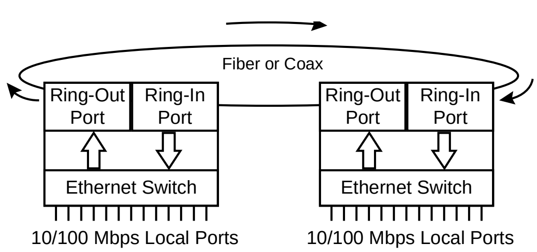

Figure 3. *EtherRing Distance Solution to Eliminate Transport Collision Domain via Ring-In/Ring-Out Ports*

An added benefit resulting from a Ring-In/Ring Out design is a doubling effect of data throughput. Because there is no collision domain on the ring, a full 100 Mb/s is available on *each* In/Out port. Consequently, an effective throughput of 200 Mb/s is available between any two devices.

Protocol Solution. The protocol and unique device software problems are entirely avoided by using MAC address switching. The MAC address is a unique 48-bit number that is built into the Ethernet hardware by every manufacturer, as the device is manufactured, and is completely protocol independent. Therefore, any protocol, such as TCP/IP, IPX, etc., will be switched correctly. Users and system manufacturers are able to install a wide variety of software products on the hardware without concern for incompatibility.

Self-learning Solution. As the packet enters a port of a standard Ethernet Switch, the switch reads the Destination Address and if it finds a matching address in its internal address table, switches the packet to the port on which that device is located. The switch learns the addresses by then reading the Source Address as each packet comes in each port and making entries in the address table that store that MAC address and associate it with the port on which the packet was received. This standard process is applied to the local Ethernet ports in the EtherRing switch and, with simple modification, on the Ring-In and Ring-Out ports. When a packet is received at the Ring-In port, the Source Address is read and entered into the address table. But in this case it is associated with the Ring-Out port instead of the Ring-In port at which it arrived. This teaches the switch that while all packets arrive at the Ring-In port, all devices on the ring must be reached by transmitting on the Ring-Out port. The result is a very simple, low cost and self-learning solution that is, again, maintenance free.

Topology Solution. To solve the problem of the Ring Topology, i.e. that a packet may come full circle around the ring and could continue going around the ring forever, another simple modification is performed on the standard switching technique. This "problem" is a normal part of Ethernet. A packet will not be switched off the ring by a Destination Address match at a local port if it is a broadcast packet, initial packet or packet that does not find a Destination Address match for any other reason at a local port. That packet will come full circle around the ring, back to the switch which has the originating device, and if not dealt with properly that packet would continue to travel around the ring indefinitely.

We solve the Ring Topology problem by performing packet filtering on the Source Address at the Ring-In port instead of the usual filtering on the Destination Address done by a standard Ethernet Switch and done at all the local ports of the EtherRing Switch. In a normal switch, and on the local ports of the EtherRing Switch, if the Destination Address of a packet is already in the table for that same port, we know that the packet is going from one device on a hub connected to the switch port to another device connected to a hub on the same switch port.

Therefore, we filter the packet, i.e. we don't let the packet into the switch because we "know" that the source and destination devices are both on the same port and there is no reason to use up switch bandwidth. In the case of the Ring-In port, if the Source Address of the packet entering the Ring-In port is the same as a source address that has already been entered into the table for a local port, then we know the packet has been full cycle around the ring and we filter it, i.e. we don't let the packet into the switch, thereby removing it from the ring. We know that we can catch all such packets, because the switches (unlike routers) create no packets or data. All packets have to be created originally by one of the Ethernet devices attached to a local port. Therefore, before the packet can get started on the ring, we have learned its Source Address, put that address in the table and associated it with the appropriate local port. When the packet comes around it gets caught.

NETWORK CONFIGURATIONS USING THE FAST EtherRing SWITCH

This section describes the basic building blocks used to configure data networks based on the EtherRing standard.

Stand-alone EtherRing

To build the most elementary EtherRing-based network, a series of Fast EtherRing switches are connected together to create a simple ring as shown in figure 3. Customers can then be served at each location via the local 10Base-T and 100Base-TX ports. The maximum distance between each EtherRing switch is determined only by the optical components being used. In this case, either 1310 nm or 1550 nm optics can be used with a loss budget up to 25 dB.

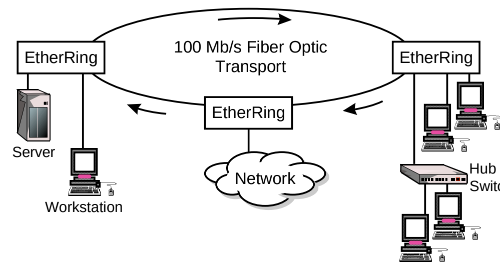

Figure 3. *EtherRing as a Stand-alone Solution*

Multiple EtherRing Elements within a High-speed Transport

The Fast EtherRing Switch has also been designed to interface directly to a high-speed multichannel digital fiber optic transport platform. This type of high-speed digital platform is used extensively [7,8] within the CATV industry for the purposes of distributing video, QAM and data services to local hub sites from a regionalized headend.

This design allows up to 16 separate and independent 100 Mb/s EtherRing channels to be multiplexed onto a 2.4 Gb/s transmission system (see figure 4). Further, using Dense Wavelength Division Multiplexing (DWDM) techniques, up to 144 separate and independent Fast EtherRing channels may be combined on a single fiber. This is accomplished by using 8 wavelengths within the 1550 nm window along with another wavelength at 1310 nm where each wavelength carries a 2.4 Gb/s aggregate data rate. This solution addresses the significant issue of scalability as network demand and subsequent traffic load increases.

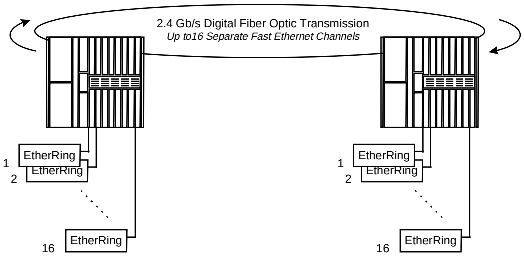

Figure 4. *16 EtherRing Channels as Part of a Higher Speed Multiplex*

The trunking method also allows multiple Fast EtherRing channels to be delivered to hub site locations where individual EtherRing channels are broken out as single channel tributaries (via direct fiber optics). This allows data delivery to remote hub sites that may not have the subscriber density to support the cost of a multichannel Ethernet distribution system - let alone support its own router. Further cost savings are realized as the EtherRing bandwidth is shared over the several sites. As shown in figure 5, EtherRing channel number 16 is not only distributed via fiber from the multichannel shelf but its bandwidth is also shared over several remote sites.

As a final note, these types of high-speed digital trunking systems allow for complete opto-electronic and fiber path redundancy to achieve maximum reliability. And, the individual EtherRing units may also be configured with optical path and terminal redundancy when used as a stand-alone solution or when integrated into the multichannel digital trunk.

With the development of EtherRing, and its use in both stand-alone and multichannel trunking architectures, the low cost, low maintenance, high throughput, high bandwidth, native-mode Ethernet technology normally associated with the LAN environment is now extended into the WAN. Armed with this powerful solution, CATV operators are able to take advantage of cost saving techniques performed within a LAN. Consequently, the CATV network can now realize significant cost savings through the application of "Route Once - Switch Many" thereby consolidating network routers at a single location and allowing network bandwidth to be shared over multiple hub sites.

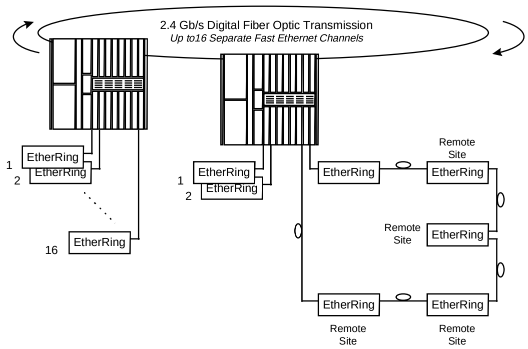

Figure 5. *Multichannel EtherRing Transport with Single Channel Fiber Tributaries*

ROUTE ONCE - SWITCH MANY

A common phase heard within the data networking industry is "switch when you can, route when you must." In other words, the deployment of layer 2 switches is the preferred mechanism for linking LANs, servers and workstations while deploying routers only when necessary. Prior to the advent of Ethernet switches, the generally accepted practice was to use routers to segment congested LAN segments or to link LANs to one another in a building or campus environment [5].

While standard Ethernet switches provide the level of sophistication (self-learning and protocol independence), these devices have been limited in distance due to the Ethernet definition. EtherRing effectively extends the functionality and flexibility associated with the LAN into the wide area network. Therefore, the Route Once - Switch Many technique, traditionally used only in the LAN, can now be extended into the WAN resulting in the consolidation of most, if not all, network routers in a single location. See figure 6a and 6b.

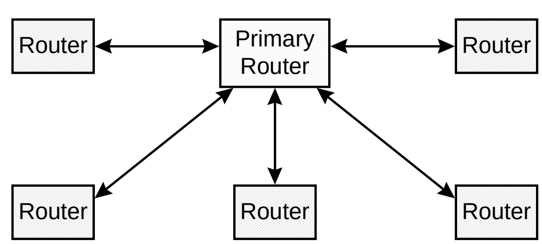

Figure 6a. *Routers Distributed using Current WAN Solutions*

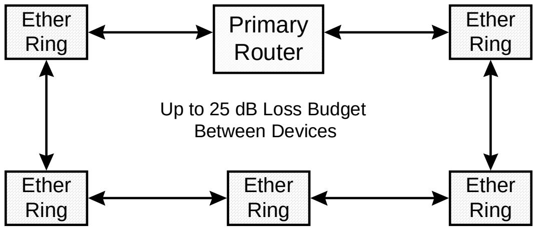

Figure 6b. *Routers Centralized using EtherRing for WAN Transport*

Several key outcomes of this approach follow. First, an initial cost savings is incurred directly through the replacement of higher cost routers with lower cost Ethernet transport leading to an overall lower initial installed cost. Bandwidth utilization is optimized as well by sharing the EtherRing bandwidth among several or more sites.

As network demand increases, the EtherRing network can be easily reconfigured at the headend to allow more or less bandwidth at each hub site. Further, with centralized routers and servers, all key personnel with the necessary expertise can be located at a single main headend site.

Current Data Transport Solutions over the WAN for Cable Modem Services

Data has traditionally been transported over the WAN using conventional routers along with some type network interface which may include; T1, DS3, OC-3c and others. In these cases the network circuits are dedicated point-to-point links. Depending on the type of network and network interface being used to transport the data service, an operator may have too much (OC-3c) or not enough (T1) bandwidth directed to the destination site.

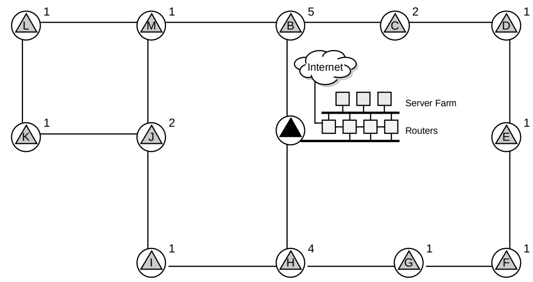

Figure 7. *Regional Headend Interconnect with Required Number of CMTS Units*

Figure 7 shows a typical regional headend interconnect. Based on subscriber count, homes passed, demographics, etc., each hub site is allocated a certain number of Cable Modem Termination Systems (CMTS). This value is labeled near the hub site symbol. In this example, each CMTS requires a 30 Mb/s data channel. The CMTS performs the digital to RF conversion (and vice versa) for access to and from the HFC network.

While some sites are expected to support multiple CMTS units, others will require only one. In fact, none of the multi-CMTS hub sites will likely require and deploy all CMTS units as data modem services are initially being implemented.

Figure 8 shows a possible data transport solution using current methods. The data signal for each CMTS unit is carried in a single OC-3c channel. As a result, an OC-48-based network is used to support multiple OC-3c circuits. Distributed routing must also be used because of the transport of encapsulated data over the network requiring subsequent re-packetization and assembly of the data channel prior to hand-off to the CMTS unit.

The network shown in figure 8 has an initial installed cost of approximately $1.1M. Note that the annual maintenance costs associated with supporting the functions of each router is not included.

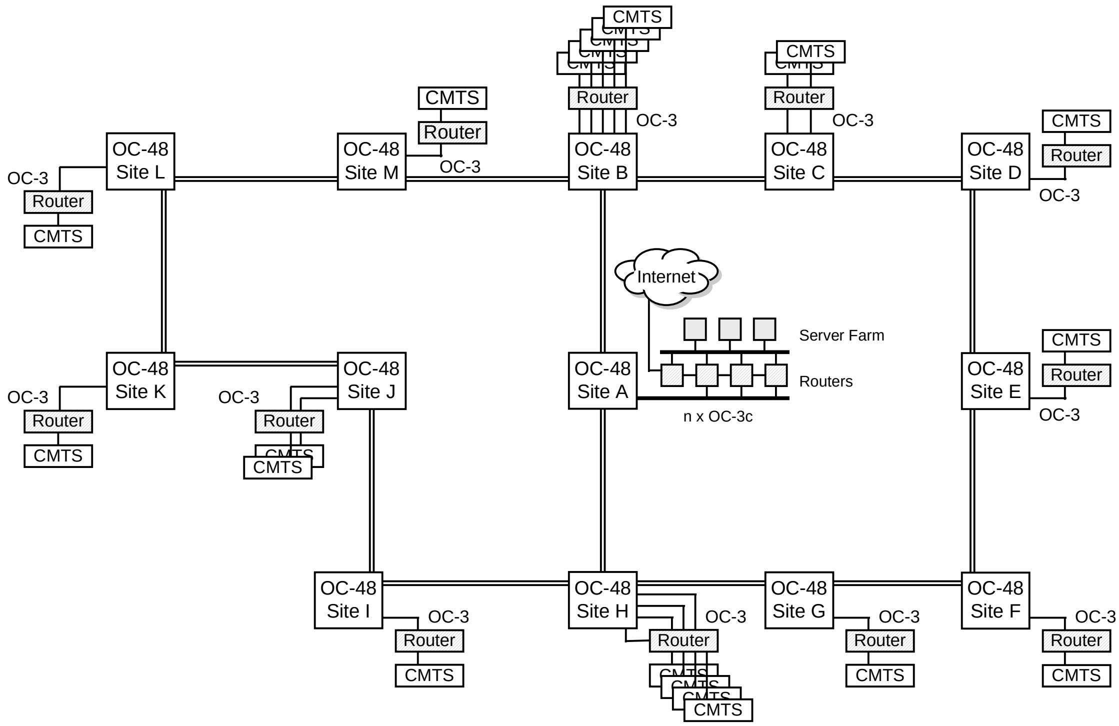

Figure 8. *Current Approach to Data Delivery over the Regional CATV Network*

As this design requires each hub be interconnected via OC-48, an extraordinary waste of network bandwidth results. For example, 21 CMTS units are served - each by a dedicated OC-3c on the OC-48 network. This results in a total of 630 Mb/s delivered on a 2.4 Gb/s platform.

Route Once-Switch Many as Applied to Cable Modem Data Transport

When EtherRing is applied to the same network (figure 9) dramatic cost savings and operational efficiencies are achieved. First, all routers previously located at the distribution hub sites have been eliminated as router functionality in now centralized within the main headend. In place of routers, Fast EtherRing switches are used. This leads to a direct hardware cost savings. Rather than an OC-n transport, a combination high-speed digital trunk along with single channel tributaries are employed. The EtherRing-based network as shown in figure 9 has an approximate cost of $500k.

This implementation takes advantage of both the statistical nature of Ethernet as well as the increased throughput (200 Mb/s) of EtherRing. Consequently, hub sites initially requiring only one or two CMTS units can now share the same EtherRing bandwidth over several sites resulting in a significant cost savings. Further, this provides an optimal solution for network scaleability where initial service demand may be low.

OTHER CATV TRANSPORT APPLICATIONS USING EtherRing

Since it's based on the Ethernet standard, an EtherRing solution may also be extended to other data transport uses within the CATV environment. These applications may include both high bandwidth as well as low bandwidth usage. A brief review of these applications follow.

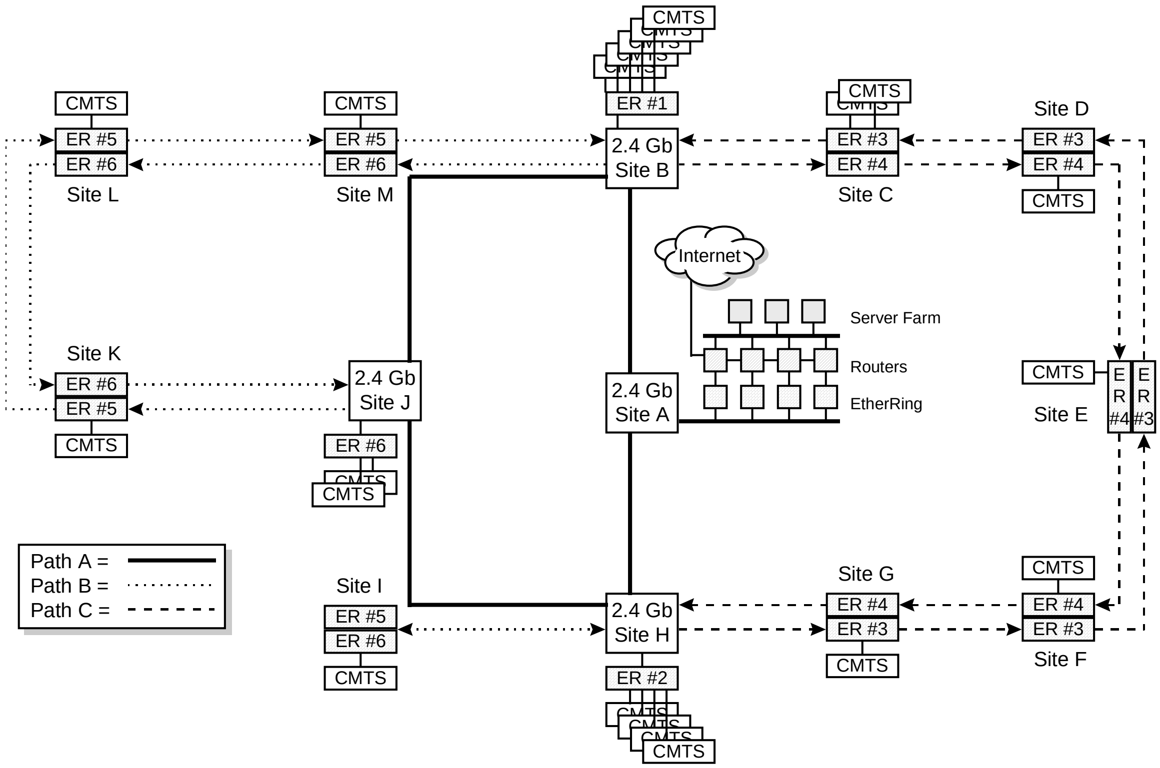

Figure 9. *EtherRing Approach to Data Delivery over the Regional CATV Network*

Advertisement Insertion

Advertisement insertion systems store MPEG compressed video files on a primary server. These files are typically delivered to local servers located within Hub sites over low-speed circuit connections such as T1. Multiple Hub sites require multiple leased circuits costing thousands of dollars per year in usage charges. Further, T1 speeds (1.544 Mb/s) are a potential bottleneck when multiple large (~10 Mb/s) files are transferred.

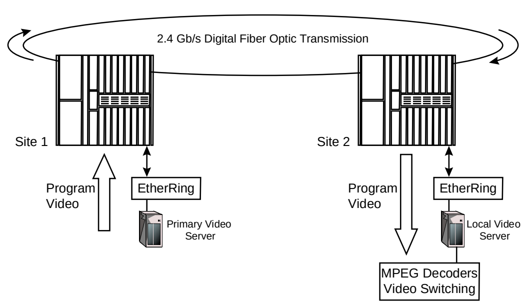

Figure 10. *Advertisement Insertion using EtherRing Transport*

A solution based on EtherRing (see figure 10) offers a simple, low-cost, high bandwidth method for transferring MPEG data files to local Hub site servers. Program channels slated for ad insertion (CNN, MTV, etc.,) may be transported in *uncompressed* format on the same platform as the compressed MPEG files allowing further network cost and operational efficiencies. A native mode Ethernet-based delivery system also offers a significant advantage not offered in current solutions - the delivery of an acknowledgment back to the Headend as the advertisement is inserted. This acknowledgment is logged into a central database thereby giving the operator an accurate record of all ad insertion activity.

LAN Interconnects - For Internal and Revenue Generating Networks

Hub sites often serve as local CATV offices that may support technical, dispatch and CSR staffs. Subsequently, office computers, servers, printers, etc., may be part of an internal LAN that is connected to the main office via a telco WAN connection (56 Kb/s circuit, fractional-T1, etc.). An EtherRing solution (figure 3) allows multiple internal LANs to be interconnected to the central office LAN - completely by-passing the telco and the associated monthly access charges. The same approach may be extended to revenue generating applications for businesses, schools, libraries, etc.

Set-Top Box Access & Control and Network Management Transport

Set-top box manufactures are currently modifying their addressable network interface controllers (ANIC) to interface directly to an IEEE 802.3 10Base-T connection. Likewise, network management platforms are migrating to network control over Ethernet connections. These applications are not as bandwidth intensive as cable modem, data and ad insertion transport. Subsequently, their Ethernet data may be directed to spare ports within an EtherRing network supporting another application.

FUTURE DIRECTIONS FOR EtherRing

Since EtherRing technology is based on industry standard definitions and practices, future products based on this technology will make use of inevitable enhancements and developments within the Ethernet standard. Of particular note are the emerging developments in Virtual LAN (VLAN), Broadcast/Multicast IP, quality of service and Gigabit Ethernet. These developments will extend the functionality and versatility of Ethernet-based transport systems while still operating within the defined standards allowing complete backwards compatibility.

CONCLUSION

This paper has described a method for delivering the functionality, low cost and flexibility (normally associated with the LAN) into the WAN through the use of a new technique known as EtherRing. Consequently, native mode Ethernet and layer 2 packet switching may be employed throughout a larger network. Several key advantages of Ethernet packet switching over routing were presented. These advantages include lower cost, higher throughput, lower latency, independence of network and application protocols and, MAC address switching with automatic self-learning.

When applied to cable data modem transport applications, EtherRing allows the Route Once - Switch Many concept to be employed over the large scale regional CATV interconnect. This permits the centralization of network routers as well as the sharing of Ethernet bandwidth over multiple hub sites resulting in several key outcomes:

- *Lower network cost*
- *Increased operational efficiencies (centralize key technical staff in a common location)*
- *Reduction of network complexity within the hub site*
- *Easier network administration (through elimination of router maintenance functions at each hub site)*
- *Increased bandwidth efficiency and utilization*
- *Simple network re-configuration for increased service demand*
- *Scaleable for service growth*
- *Support of multiple applications based on the Ethernet standard*

EtherRing™ is a registered trademark of ADC/Codenoll

REFERENCES

(1) Bay Networks White Paper, "Networking: A Primer. 1996. http://www.baynetworks.com/Products/Papers/wp-primer.html

(2) http://www.lantronix.com/htmlfiles/mrkg/catalog/et.htm

(3) http://pclt.cis.yale.edu/pclt/comm/ether.htm

(4) Molle M., Watson G., "100Base-T / IEEE 802.12 / Packet Switching", IEEE Communications Magazine, Vol. 34 No. 8 pp 64-71.

(5) Passmore L., Bransky J., "Route Once, Switch Many" White Paper from Decisys, July, 1997.

(6) Coden M., "Circuits and Methods for a Ring Network", United Stated States Patent Application, August 1997.

(7) Harris R.W., "Digital Backbone Network Applications for Inter-City and Intra-City Regional CATV Networks" NCTA Technical Papers, May 1994.

(8) Holobinko J., "An Analysis of Interoperability and Cost Factors for Regional Digital Backbone Networks", NCTA Technical Papers, May 1994.

Robert W. Harris
Market Program Manager, Digital Transmission
ADC Broadband Communications
999 Research Parkway
Meriden, CT 06450
203.630.5732

Mr. Harris serves as Marketing Program Manager for the Broadband Communications Division of ADC responsible for new market, product and business development of digital transmission solutions for video and data applications. Mr. Harris is an 18 year veteran of the CATV and telecommunications industry. During the 1980's he held positions in construction, field service and system engineering with Viacom and Storer Cable. He later worked as a LAN/WAN Network Engineer at NASA's Langley Research Center. Since 1990, he has held engineering and product management positions with C-COR Electronics, Nortel and ADC working with analog and digital fiber transmission, computer network and MPEG based products. He completed his undergraduate training in physics at Old Dominion University and graduate training in engineering science and business at Penn State University and the University of Phoenix, respectively.

Michael Coden
President
ADC Codenoll
200 Corporate Boulevard South
Yonkers, New York 10701
914.965.6300

Michael Coden is President of ADC Codenoll (a subsidiary of ADC Telecommunications Inc.) which he founded in 1980. ADC Codenoll develops, manufactures, installs and services LAN hardware and software products for Ethernet, Switched Ethernet, Fast Ethernet and FDDI LANs. Mr. Coden has also served in various engineering, management and executive level positions for Exxon Enterprises, Maher Terminals Inc., Digital Equipment Corporation and Hewlett-Packard. Mr. Coden holds a Bachelors Degree in Electrical Engineering form MIT, an MBA from Columbia University and a Masters Degree in Mathematics from the Courant Institute of Mathematical Sciences at New York University. He has obtained nine patents in LAN technology and fiber optics, some of which are used in IEEE and ISO international standards.