

Broadband Access over Cable for Next-Generation Services: A Distributed Switch Architecture

Subra Dravida, Dev Gupta, Sanjiv Nanda, Kiran Rege, Jerome Strombosky, and Manas Tandon
Narad Networks

ABSTRACT

The hybrid fiber coax architecture deployed by the cable service providers has been successful in capturing a substantial piece of the residential broadband access market. In the United States over five million homes connect to the Internet using DOCSIS¹ cable modems. In this article we describe an evolution path to enhance the HFC plant to provide, initially, Gigabit Ethernet (and eventually multi-Gigabit Ethernet) on the trunk and feeder portions, and 100 Mb/s Ethernet on the subscriber drops. This next-generation HFC network will enable cable service providers to address the vast and underserved small and medium-sized business market, as well as offer emerging applications and services to the residential market.

INTRODUCTION

Data transmission and networking technologies have witnessed tremendous growth over the past decade. However, much of this development and growth has been in the core networks where high-capacity routers and ultra-high-capacity optical links have created a truly broadband infrastructure. The so-called last mile — the access link connecting the user to the backhaul infrastructure — remains a bottleneck in terms of the bandwidth and service quality it affords the end user. The last mile problem has impeded the growth of truly broadband services and applications.

The inadequacy of the access link is particularly felt when one attempts to envision the applications and services likely to become popular in the future. Interactive video applications, interactive gaming, video telephony, videoconferencing, remote storage, virtual DVD, and high-speed virtual private networks (VPNs) between geographically separated office locations or between homes and office locations for telecommuters are just a few such applications. It is likely that once the bandwidth and quality barrier of the access link is removed, new and unforeseen applications will

emerge and attain widespread popularity. Many of these emerging applications cannot be supported with the asymmetric bandwidth available on today's access technologies. Moreover, they require guaranteed quality of service (QoS) in terms of packet delays and throughput. Thus, the capability to provide symmetric high-quality high-bandwidth bearer services will be a key requirement for access media of the future.

It is possible to provide such a high-quality high-bandwidth access medium by taking fiber to the home. However, from the service provider's perspective, such a solution is inordinately expensive and inefficient. Available technologies, such as asynchronous digital subscriber line (ADSL) and DOCSIS (cable modem) that make use of the existing infrastructure, are asymmetric and very limited in the bandwidth they provide. The virtually ubiquitous hybrid fiber coax (HFC) plant, however, represents a vast underutilized resource that can be exploited to provide a high-quality high-bandwidth access medium to practically the entire community of households and small business users. A large portion of the usable coaxial cable spectrum is as yet unused. By developing key components that can be inserted as required, the HFC infrastructure can be turned into a truly broadband QoS-enabled access medium. Our solution offers an incremental upgrade path to convert the HFC plant into a high-quality high-bandwidth access medium capable of delivering next-generation services.

THE HFC INFRASTRUCTURE AND DOCSIS

The cable plant, whose original objective was to deliver analog TV signals to homes obstructed from line-of-sight reception, is typically organized in a *trunk and branch* topology. Over time, however, cable service offerings have grown, adding numerous commercial and paid programming services, telephony, and asymmetrical data

¹ DOCSIS: Data over Cable System Interface Specification [1].

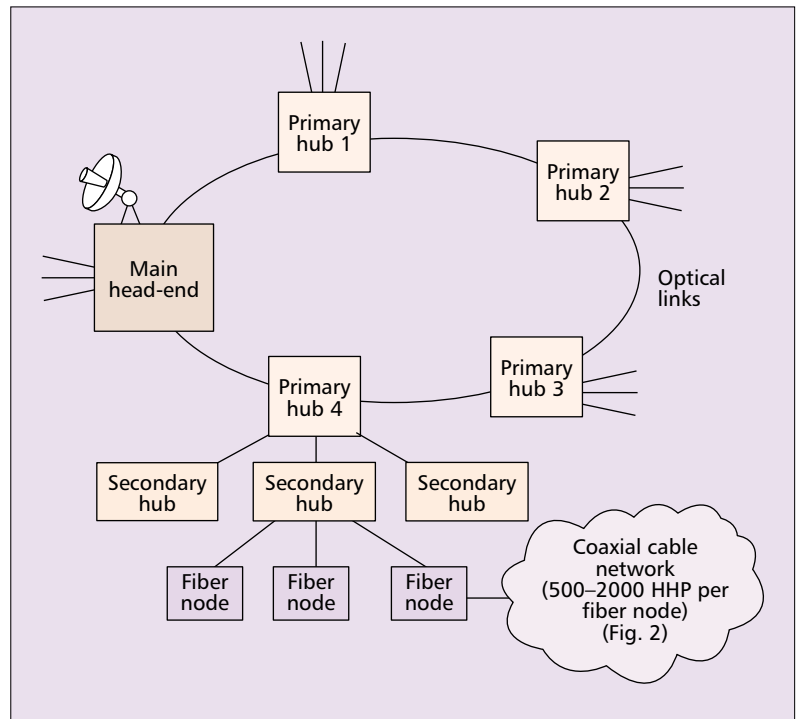
communication capability based on the DOCSIS standard [1]. In order to improve overall signal quality and facilitate growth in the data communication sector, cable companies have been gradually upgrading their plant to HFC, which often involves partitioning the old branch-and-tree structure and running fiber optic links from the head-end to fiber nodes strategically located within the old plant. The fiber nodes, in turn, inject higher-quality RF signals deeper into the coaxial infrastructure. The evolution of the cable plant to HFC, and current and future Cable Labs initiatives were covered in a special issue of *IEEE Communications Magazine* [2].

THE HFC INFRASTRUCTURE

Figure 1 shows a high-level schematic of a typical HFC plant designed to deliver broadcast television and DOCSIS-based data services to the home user [3]. As seen in this figure, a typical large metropolitan area with 0.5–1 million homes is served by one or two master head-ends. The master head-ends may be connected to several primary hubs with a multifiber optical ring network. Each primary hub is connected to a few secondary hubs, which, in turn, are connected to fiber nodes using optical links. Each fiber node serves 500–2000 homes (households passed, HHP), each secondary hub about 20,000 homes. The optical fiber connections end at fiber nodes [3]. From there on, the connections are over coaxial cable. The head-end typically has satellite and land-based microwave antennas to receive broadcast television signals from the various television network channels. It also has equipment to introduce local advertising and manage interactive pay-per-view services. In addition, if the HFC plant has been upgraded to provide DOCSIS-based two-way data services to home users, the head-end has routers and access servers that are needed to enable home users to access the Internet.

In an HFC plant upgraded for DOCSIS-based data services, special devices called cable modem termination systems (CMTS) are located at the primary or secondary hubs. A CMTS is a termination point for the physical and medium access control (MAC) layer communication protocols to the cable modems (in user homes). The downstream transmissions from the CMTS are combined with the broadcast television signals, and sent over optical links to the appropriate fiber nodes. In the upstream direction, DOCSIS signals are carried over optical links from the fiber nodes to the hubs where they terminate at the CMTS. The CMTS extracts the DOCSIS signals, and, after physical and MAC layer processing, forwards the packets to the router.

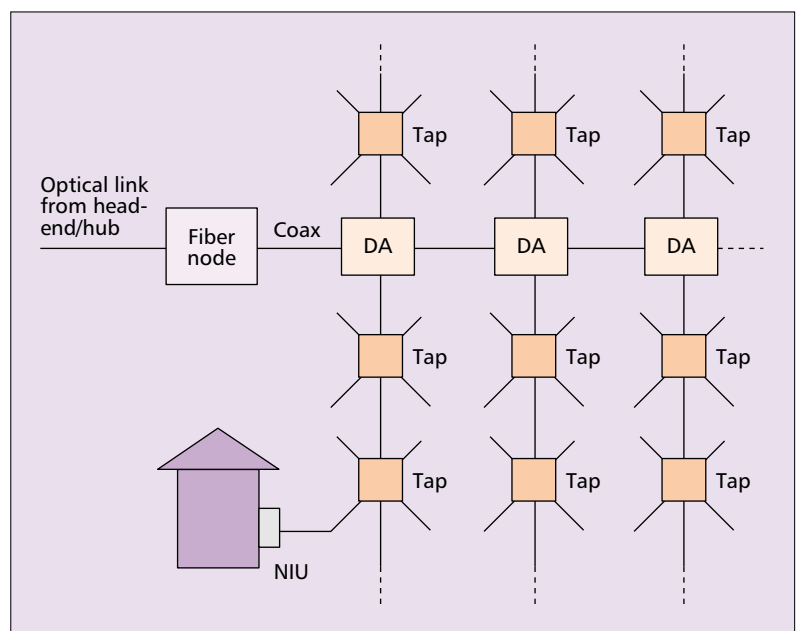
The final segment of the HFC network — between the fiber node and the home — is over coaxial cable. Figure 2 shows a schematic of this segment. The fiber node demodulates the optical signals it receives from the corresponding hub to extract the downstream signals and translates them to the appropriate frequency band for transmission over the coaxial cable medium. These signals are carried over the *main trunks* that emerge from the fiber node. On a main trunk, distribution amplifiers (DAs) are located at places where *feeder* cables need to be taken out to serve user communities.



■ **Figure 1.** A high-level view of the fiber portion of the HFC plant in a metropolitan area.

Feeder cables have *taps* located at places where cable *drops* are taken out to serve individual homes. Line extending amplifiers (LE) can be inserted into feeder cables or main trunks to provide signal amplification on longer cable segments. A cable drop connecting the HFC plant to a home terminates at a network interface unit (NIU), typically located at the side of the house. Cable segments, possibly with splitters, are run from the NIU to different parts of the house.

In HFC plants upgraded for bidirectional data services (including DOCSIS), the upstream signals generated by home users are carried to the fiber



■ **Figure 2.** A schematic of the access segment of the HFC plant.

The DOCSIS standard for bi-directional data services over the HFC medium has been developed by a consortium of North American cable operators. It uses portions of the HFC spectrum that is left over after accommodating broadcast television services.

node over the same coaxial cable segments that carry the corresponding downstream signals. To make this possible, the taps are equipped with directional couplers. Also, at DAs and LEs, diplexers are deployed to separate the upstream and downstream signals before they are amplified for further transmission. A key feature of this architecture, which limits the capability of the current HFC plant to carry symmetric broadband services, is that the DAs and LEs do not regenerate the upstream (or downstream) signals. They merely act as analog amplifiers. At the fiber node, the upstream signals are modulated on to the optical fiber and transmitted to the corresponding hub.

DOCSIS

The DOCSIS standard [1] for bidirectional data services over the HFC medium has been developed by a consortium of North American cable operators. It uses portions of the HFC spectrum that is left over after accommodating broadcast television services. Specifically, DOCSIS uses a fraction of the frequency band between 5 and 42 MHz for upstream transmissions, and between 550 and 750 MHz for downstream transmissions. These bands are further divided into individual channels as follows. For downstream transmissions, the allocated spectrum is divided into 6 MHz channels. With 64- or 256-quadrature amplitude modulation (QAM), this yields a data rate of 30.342 Mb/s or 42.884 Mb/s per channel. For upstream transmissions, the allocated spectrum is divided into individual channels, each with a bandwidth of 200, 400, 800, 1600, or 3200 kHz [4]. While both quaternary phase shift keying (QPSK) and 16-QAM are permitted, the noise conditions on the cable segments are typically such that only QPSK modulation is possible. This works out to a data rate ranging from 320 kb/s for a 200 kHz channel to 5.120 Mb/s for a 3200 kHz channel. The aggregate bandwidth over all upstream channels is typically around 30 Mb/s. The aggregate downstream bandwidth, on the other hand, is an order of magnitude larger.

DOCSIS uses the entire coaxial cable segment served by a fiber node as a single multiple access medium. It specifies the physical, MAC, and data link layer protocols for communication between the cable modems at homes and the CMTS. DOCSIS exploits the broadcast nature of the HFC medium for downstream transmissions using time-division multiplexed (TDM) transmissions in one or more 6 MHz wide frequency-division multiplexed (FDM) channels. On the upstream, DOCSIS relies on multiple access techniques to share the medium among geographically separated users. Upstream transmissions are also TDM-based and at the CMTS receiver the entire cable plant subtended at a fiber node appears as a single shared bus, subject to packet collisions. Due to the large distance between the farthest cable modem and the CMTS receiver, random access schemes such as slotted Aloha and enhancements such as carrier sense multiple access (CSMA) or busy tone multiple access (BTMA) [5] are inherently inefficient. To improve efficiency, DOCSIS uses a combination of random and reserved access in the upstream direction. This enables DOCSIS to minimize the waste of transmission bandwidth while being able to exploit the statistical nature of the data traffic.

LIMITATIONS OF DOCSIS

DOCSIS provides a truly broadband communication medium in the downstream direction where the relatively noise-free nature of the transmission medium, coupled with the 200 MHz spectrum allocated for downstream transmissions, can be exploited to yield an aggregate bandwidth of several hundreds of megabits per second. With 500–2000 homes passed per fiber node, this works out to about 0.5–1 Mb/s per HHP. With such large bandwidth available in the access medium, several bandwidth-intensive applications can be made available to the home user. However, these applications have to necessarily be asymmetric in their bandwidth requirements since the bandwidth available in the upstream direction is significantly smaller.

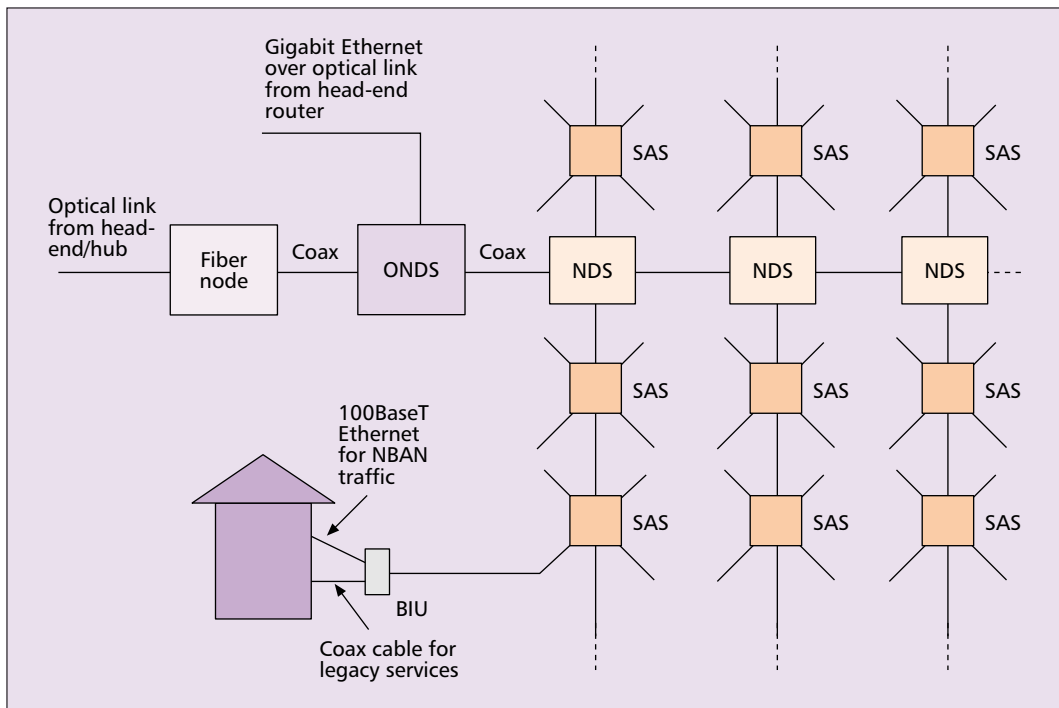
The major limitations of DOCSIS become immediately visible when one looks at its upstream bandwidth. The current DOCSIS [1] upstream is limited to a few tens of Mb/s by the following constraints:

- The DOCSIS upstream is assigned to a fraction of the 5–42 MHz return path of the cable spectrum. It shares this bandwidth with other legacy upstream applications, including set-top boxes and telephony.
- The return spectrum (5–42 MHz) is most subject to ingress noise. Many sources of ingress noise, including citizen band and amateur radio, shortwave radio, and even sunspot activity, cause ingress noise in this band.
- On the DOCSIS return path, signals from all cable terminations enter the cable plant and fan in toward the CMTS. Therefore, the CMTS receiver sees the noise aggregated from the entire plant subtended at the fiber node. This is referred to as a *noise funnel*.
- The achieved carrier-to-noise ratio (CNR) on the return path is limited by the transmit power at the cable modem as well as laser clipping and analog-to-digital (A/D) full-scale saturation limits at the fiber node. These factors limit the achievable spectral efficiency of upstream transmissions.
- The achieved CNR on the return path is further limited by the noise figure of the cascade of amplifiers from the cable modem to the fiber node. With a typical cascade of 3–6 DAs/LEs, the CNR is reduced by 5–8 dB.

An aggregate bandwidth of a few tens of megabits per second for 500–2000 HHP works out to an average of only a few tens of kilobits per second per HHP, which is clearly incapable of carrying symmetric high-bandwidth applications on a large scale.

Other limitations of DOCSIS that must be addressed by the next-generation HFC technology are as follows:

- Downstream transmissions in DOCSIS are broadcast to the entire HFC segment hanging off a fiber node, with cable modems filtering packets unless the MAC layer address matches their own. To prevent snooping by a rogue cable modem user of transmissions intended for other users, DOCSIS transmissions are protected by appropriate high-level encryption mechanisms.



■ **Figure 3.** A schematic of the access segment of the NBAN.

The coaxial cable itself has a rather large unused spectrum, stretching to 3 GHz and beyond. NBAN uses this spectrum to carry upstream and downstream transmissions to deliver symmetric bandwidth to the end-users.

- DOCSIS uses TDM transmissions where slot boundaries are set by a global clock residing in the CMTS. All cable modems need to synchronize with this clock so that their transmissions are correctly aligned in time when they arrive at the CMTS receiver. Since these cable modems are dispersed over distances that vary from a few hundred yards to a few miles, they need to implement ranging schemes to help them synchronize with the global clock. The ranging operations are wasteful of bandwidth and require guard period overheads to account for inaccuracies in range estimates.
- The DOCSIS frame interval places a minimum on the achievable delay jitter for real-time services, such as voice over IP (VoIP).

THE NEXT-GENERATION BROADBAND ACCESS NETWORK

The next-generation broadband access network (NBAN) overcomes these limitations of DOCSIS by making fundamental changes in the way packet data transport is carried out over the HFC medium *without* significantly altering the HFC plant itself.

NBAN UPGRADED PLANT

Coaxial cable itself has a rather large unused spectrum, stretching to 3 GHz and beyond. NBAN uses this spectrum to carry upstream and downstream transmissions to deliver symmetric bandwidth to end users. In a network completely upgraded to NBAN, the MSO replaces HFC devices such as taps, DAs, and LEs with intelligent devices with a store-and-forward switching capability (Fig. 3). DAs and LEs are replaced with network distribution switches (NDSs). Each

NDS subsumes the legacy spectrum functionality of the DA or LE it replaces. In addition, the NDS provides packet switching for NBAN data. Taps are replaced with subscriber access switches (SASs). In a typical implementation of NBAN, the main trunk (connecting to the fiber node) and feeder cables provide 1 Gb/s pipes,² whereas the subscriber drops (connecting SAS to the customer premises) provide 100 Mb/s links.

Turning the HFC infrastructure into a network with store-and-forward switching elements has several benefits:

- NBAN consists of short point-to-point network segments between each pair of adjacent devices: NDS to SAS or SAS to SAS. Compared to current DOCSIS, this results in a high CNR due to shorter path loss and elimination of noise accumulation through amplifier cascades [6]. This enables NBAN to use bandwidth-efficient signal constellations even in the upstream direction. Moreover, this avoids the bandwidth reduction due to filter and connector cascades. As a consequence, NBAN is able to deliver on each cable segment attached to a fiber node a symmetric bandwidth of 1 Gb/s.
- On each point-to-point link, the receiver at each endpoint needs to synchronize only with the transmitter at the other end of the link. This can be done locally, eliminating the need to synchronize with a global clock and the corresponding ranging operations.
- The store-and-forward paradigm with short point-to-point links eliminates the need for random access, which is a source of inefficiency in DOCSIS.
- Only data destined to a particular customer premises broadband interface unit (BIU) is switched to the customer's drop at the SAS. This provides strong privacy of a sub-

² The data rate on the trunk and feeder cables can be increased to multiple gigabits per second in the future.

Switched Ethernets, with IP over Ethernet data transport, have become the norm for enterprise networks and are now being deployed in the MAN. This has given rise to the development of commercial off-the-shelf application and management software as well as a trained technical workforce.

- scriber's transmissions since other users' transmissions are not visible at the BIU.
- NBAN behaves like a switched Ethernet connecting home and business subscribers to the routers located at the hubs in the cable service provider's network. Switched Ethernets, with IP over Ethernet data transport, have become the norm for enterprise networks and are now being deployed in the MAN. This has given rise to the development of commercial off-the-shelf application and management software as well as a trained technical workforce. NBAN is consistent with and exploits the synergy with this "standard" networking technology.

NBAN DEVICES AND FUNCTIONAL DESCRIPTION

The BIU — The BIU is the interface between the NBAN and the subscriber LAN. A BIU has two subscriber-side interfaces: a coaxial cable interface for legacy services to the subscriber, and a 100BaseT Ethernet interface to carry NBAN traffic to and from the subscriber. On the NBAN side the BIU has a single coaxial interface that carries the aggregate of legacy traffic and NBAN traffic in both directions.

For upstream NBAN traffic, the BIU performs a learning bridge function, preventing the local traffic on the subscriber's home LAN from unnecessarily being transmitted over the NBAN. Also, to Ethernet frames generated by the subscriber, the BIU adds an NBAN label, inserting its own routing ID (RID) in the corresponding field, and setting the QoS and control fields appropriately (as discussed in the next section). The BIU also performs the QoS management function of policing upstream traffic before it can enter the NBAN.

Besides these transport functions, the BIU also implements features such as virtual private network (VPN) service via different tunneling mechanisms, traffic statistics collection for billing and performance monitoring, and so on. These features enhance the utility of the NBAN beyond that of a mere data transport vehicle.

NDSs and SASs — In terms of NBAN switching functionality, the NDSs and SASs are almost identical. They differ in that each SAS has four subscriber ports for carrying NBAN traffic at 100 Mb/s in addition to its two 1 Gb/s ports for the upstream and downstream sections of the feeder cable, whereas each NDS has one upstream port and four downstream ports, all at 1 Gb/s.

NDSs and SASs have the capability to separate legacy traffic from NBAN traffic at each ingress port and combine these traffic streams at the egress ports. The devices switch downstream NBAN frames to the appropriate port based on the RID field in the NBAN label attached to the frames. Downstream frames are switched in first-in first-out (FIFO) order. For upstream NBAN traffic, NDSs and SASs have multiple queues corresponding to different QoS classes. An incoming upstream frame is placed in the appropriate queue based on its QoS class. Frames in these queues are selected for trans-

mission further upstream according to a scheduling discipline that combines priorities with weighted round-robin scheduling.

Like the DAs and LEs it replaces, the NDS provides downstream and upstream signal amplification for legacy transmissions. The SAS can also be equipped to provide amplification for legacy signals. To avoid introducing too many actives and impacting end-of-line signal quality, the active circuitry in the legacy band of the SAS can be bypassed.

SASs also have a switch bypass feature for NBAN traffic, which is critical to the reliability of NBAN. This feature allows the system to put a malfunctioning SAS in a bypass mode, limiting service disruption to only those subscribers whose BIUs are directly connected to the malfunctioning device.

ONDS — The optical NDS (ONDS) is the NBAN's interface to the WAN. On the WAN side, it has two sets of interfaces: four coaxial cable interfaces that connect it to a fiber node in the HFC network for legacy traffic and an optical link that connects it to a router via a Gigabit Ethernet interface for NBAN traffic. On the downstream side, an ONDS has four coaxial cable interfaces to carry the aggregate upstream and downstream traffic.

The ONDS maintains a table associating MAC addresses of the end devices that are downstream from the ONDS and the RIDs of the BIUs through which these end devices access the NBAN. This association is learned and maintained in the same way as source learning is performed by Ethernet switches.


An ONDS terminates the Gigabit Ethernet interface to the WAN. To frames flowing in the downstream direction, the ONDS attaches an NBAN label with the routing ID field filled with the RID of the corresponding BIU. The frames are then switched to the appropriate downstream port where the NBAN output signal is combined with the legacy downstream transmissions.

In the upstream direction, the ONDS separates legacy upstream transmissions from NBAN transmissions. The former are passed on to the fiber node over the coax interface. For NBAN transmissions the ONDS removes the NBAN labels and passes the Ethernet frames on to the router via the Gigabit Ethernet interface.

NBAN MANAGEMENT SERVERS

Other key components of the NBAN are the management servers, which are responsible for critical management functions. Among them are the DHCP and topology server, the connection admission control (CAC) server, and the network management system (NMS). In a typical implementation of NBAN, these servers are connected via a LAN to a router in the service provider's network.

The Topology Server — The topology server is responsible for maintaining topological data for NBAN and making RID assignments to the various devices therein. During system bootup, the NBAN devices generate special DHCP messages that are carried upstream, with each device on



■ **Figure 4.** NBAN frequency allocation plan.

the way attaching its identifier in the DHCP extension field, until the message is delivered to the topology server. The topology server can then construct a complete picture of the network topology and assign RIDs to the various NBAN devices in a suitable manner. The NBAN devices are informed of their assigned RIDs via DHCP Response messages. The topology and RID information is maintained in a special database that can be accessed by other NBAN management servers.

The CAC Server — The CAC server is the centerpiece of the QoS management scheme for the NBAN. It acts as a resource manager to ensure that the system has adequate resources to deliver the contracted QoS to connections that have been established with QoS guarantees. It keeps track of bandwidth utilization along different segments of the network and grants a connection request with QoS guarantees only if the resulting bandwidth utilization along critical segments of the network is within certain permissible limits. In order to perform this function, the CAC server needs to interact with call agents (e.g., VoIP signaling servers), which communicate with it during call setup and termination phases. In addition, the CAC server communicates with the BIUs to update their policing parameters for regulating the inflow of subscriber traffic as new connections are established and old ones torn down.

The Network Management System — The NMS is responsible for provisioning and maintaining the network. It communicates with the other management servers and NBAN devices to carry out provisioning and maintenance activities. It also interfaces to other systems (e.g., those processing service orders or trouble tickets). Specifically, the NMS maintains configuration data on all the elements in the network, downloads configuration parameters and software images to these elements when required, collects performance statistics reported by these elements, and processes them for reporting to human operators. The NMS also receives traps reporting faults or alarms generated by different network devices, and processes them to help the maintenance crew isolate faults and their locations. All in all, the NMS is key to turning NBAN into a reliable and flexible platform for next-generation services.

NBAN TRANSPORT AND SWITCHING

NBAN CHANNELS AND SPECTRUM

The NBAN defines four channels, referred to as channels 1, 2, 3, and 4, for data transport. These channels can be used flexibly for upstream and downstream transmissions. Of these, channels 1 and 2 are designed for carrying 100 Mb/s Fast Ethernet payloads, whereas channels 3 and 4 are designed for Gigabit Ethernet payloads.

The frequency band allocations of these channels are shown in Fig. 4. The NBAN broadband allocation resides above 860 MHz, already used for legacy applications. The NBAN uses the 900–1100 MHz regime for 100 Mb/s operation (i.e., channels 1 and 2) and the 1.2–2.4 GHz regime for 1 Gb/s operation (i.e., channels 3 and 4). The frequency allocations of the NBAN channels have been carefully selected to avoid interference with legacy transmissions.

PHY AND MAC LAYERS

The NBAN PHY (NPHY) and NBAN MAC (NMAC) layer design supports 100 Mb/s or 1 Gb/s Ethernet transfer rates, using 16-QAM modulation at symbol rates of 31 or 311 Msymbols/s. The NPHY framing structure for 1 Gb/s operation is shown in Fig. 5. The frame synchronization (FS) period is of duration 1 μ s and provides framing and carrier recovery. The symbol synchronization (SS) period is 400 ns and is used to recover symbol timing. Carrier, frame, and symbol synchronization retraining is updated every 10 μ s. A four-byte flow control and pointer (F/P) field precedes each 252-byte block of payload. The concatenation of 252-byte payload segments offers a continuous byte stream to the NMAC.

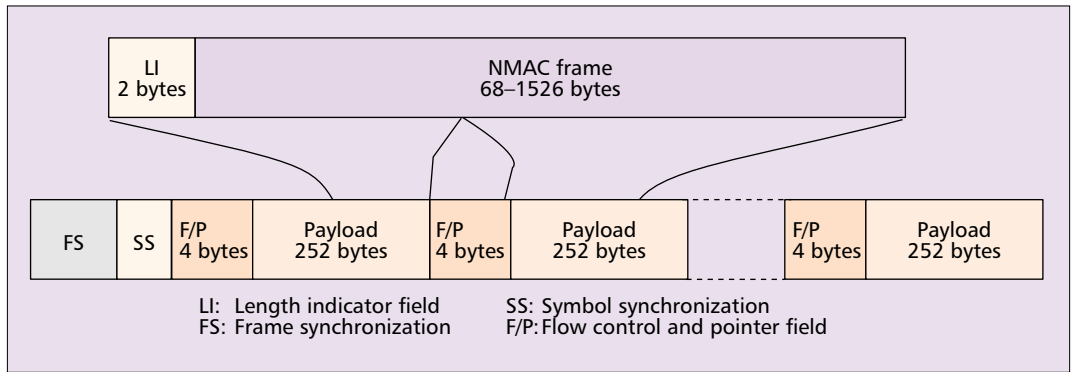
Within each NPHY frame, which begins with an FS period, there are five instances of F/P bytes and payloads (approximately, one every 2 μ s). The 4-byte F/P field contains 1 byte of pointer field, 4 bits to carry flow control signals, and 1 byte of cyclic redundancy check (CRC). Twelve unused bits of the F/P field are reserved for future use.

A length indicator (LI) is attached to the NMAC frame as shown in Fig. 5. The LI is 2 bytes (11 bits to indicate the length and 5 bits of CRC.) The frame length can vary from 68 to 1526 bytes inclusive of the NBAN label. The LI field indicates the length of the corresponding LI-encapsulated NMAC frame in bytes. The length includes the 2 bytes used for the LI field itself.

The LI-encapsulated NMAC frames are carried in the byte stream comprising concatenated

The CAC server is the centerpiece of the QoS management scheme for NBAN. It acts a resource manager that ensures that the system has adequate resources to deliver the contracted quality of service to connections that have been established with QoS guarantees.

The 3-bit QoS field indicates the Quality of Service class to which the corresponding frame belongs. For a frame going in the upstream direction, its QoS field determines the level of priority it receives at each NBAN device.



■ **Figure 5.** The NMAC and NPHY framing structure. The NMAC frame encapsulated with the length indicator is mapped into the payload portion of the NPHY framing.

NPHY frames as shown in Fig. 5. The pointer byte in the F/P field in the NPHY framing points to the start of a new NMAC frame within the subsequent payload segment. (If a frame boundary does not occur within an NPHY payload segment, the corresponding pointer byte is set to all 1s.) The LI field attached to the NMAC frames and the pointer in the NPHY F/P field enable NBAN receivers to quickly recover frame boundaries following loss of framing.

NBAN LABEL SWITCHING

Since in an NBAN all downstream frames are switched at every intermediate device they encounter, NBAN devices need to be capable of switching frames at rates up to 1 Gb/s. In order to simplify the switching mechanism in these devices, NBAN uses a 12-bit RID for switching within NBAN.³ Each NBAN device is assigned a RID, which is unique within the NBAN segment served by a single fiber node. A packet destined to a subscriber end-system uses the RID associated with the corresponding BIU.

The RID is carried within a 4-byte NBAN Label that is inserted after the source MAC address field as shown in Fig. 6. The 3-bit control field in the NBAN Label indicates the type of message being carried in the frame, which triggers different kinds of actions at the NBAN devices. For instance, specific control plane messages such as heartbeats, routing table updates, and so on, which require specific actions by the NBAN devices are assigned different control bit settings.

The 3-bit QoS field indicates the Quality of Service class to which the corresponding frame belongs. For a frame going in the upstream direction, its QoS field determines the level of priority it receives at each NBAN device. The use of the QoS field and the overall QoS management scheme are explained in detail in the next section.

As mentioned earlier, the 12-bit RID field carries the routing ID. For upstream frames generated by an NBAN subscriber, the BIU, after receiving the frames over the 100BaseT interface, attaches the NBAN label with the RID field filled with its own RID. The QoS field is determined on the basis of the flow to which the frame belongs. Since the destination for all upstream subscriber frames within the NBAN is the router attached to the ONDS, the RID field

is ignored by all intermediate devices as the frame gets switched upstream. When the ONDS receives the NBAN frame, it strips the NBAN label from it and forwards the Ethernet frame to the router over the Gigabit Ethernet interface. The ONDS learns the association between the MAC addresses of subscriber devices and the RIDs of the BIUs to which they are attached, as well as MAC-to-RID mappings of all NBAN devices, from frames flowing in the upstream direction.

For a downstream frame destined for an NBAN subscriber, the ONDS attaches the NBAN label on receiving the Ethernet frame from the router. For switching downstream frames, the ONDS fills the RID field with the routing identifier of the BIU to which the destination subscriber device is attached. At each NDS or SAS, the frame gets switched in the downstream direction on the basis of the RID. The QoS field in the downstream direction is ignored by all intermediate devices. When the NBAN frame reaches the BIU, the NBAN label is stripped and the Ethernet frame is forwarded to the subscriber over the 100BaseT interface.

QoS MANAGEMENT

Other than its high bandwidth, the aspect of NBAN that makes it an ideal vehicle for next-generation services is its elaborate QoS management scheme that delivers different levels of service to different end-user applications based on their QoS requirements. The NBAN QoS management scheme accomplishes this through a combination of CAC, priorities, flow control, and ingress traffic policing.

FOUR QoS CLASSES

NBAN allows four QoS classes: real-time constant bit rate (RT-CBR) service, real-time variable bit rate (RT-VBR) service, non-real-time VBR (NRT-VBR) service with throughput guarantees, and unspecified bit rate (UBR) service. The RT-CBR and RT-VBR services have delay objectives of 5 ms and 15 ms, respectively, and also a guaranteed throughput. The NRT-VBR service does not have an explicit delay objective; however, it has a guaranteed throughput. The UBR service has no QoS guarantees. As described earlier, the NBAN label attached to a frame as it enters the NBAN has a 3-bit QoS field

³ The RID is more appropriately a switching ID, since it is used to speed up downstream layer 2 switching.

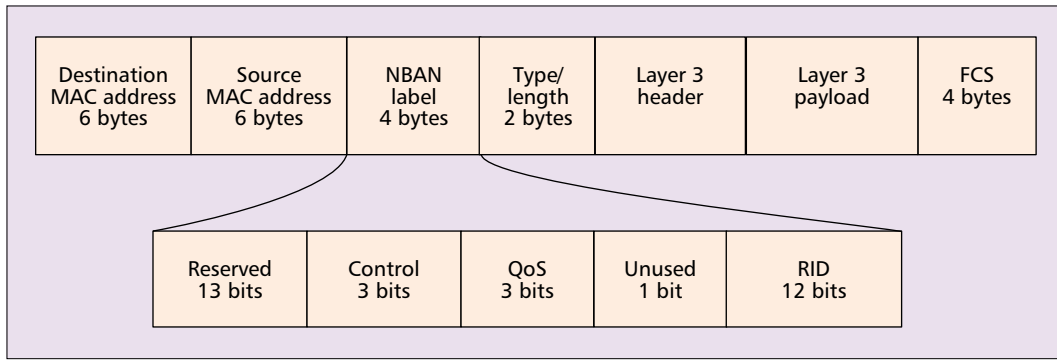


Figure 6. An NMAC frame with the position of the NBAN label. The NMAC frame structure is identical to Ethernet with the NBAN label occupying the same position as the 802.1P/Q VLAN label.

which identifies the QoS class associated with that frame. Once an upstream frame enters the NBAN, its treatment at the intermediate devices it encounters depends entirely on the value contained in the QoS field of its NBAN label.

CALL ADMISSION CONTROL

As we saw earlier, the CAC server is the principal resource manager responsible for managing QoS within the NBAN. It keeps track of bandwidth utilization by different QoS classes along the critical segments of NBAN, and allows the establishment of a connection requesting a QoS level other than UBR only if the resulting bandwidth utilization is within specified limits. When a connection is established or torn down, the CAC server updates its relevant bandwidth utilization figures, and informs the BIU through which the connection accesses the NBAN about the corresponding event.

UPSTREAM FLOW CLASSIFICATION AND POLICING AT THE BIU

Every NBAN device plays a role in the management of QoS. The BIU represents the ingress point for upstream traffic, that is, for traffic generated by an NBAN subscriber. The BIU is informed of every connection requiring a QoS level other than UBR. For every Ethernet frame received from the subscriber, the BIU identifies the flow to which it belongs using appropriate identifiers (e.g., source IP address and port number, or MAC address, or other flow identifier), and attaches to it an NBAN label with the corresponding QoS marker. The BIU polices the incoming upstream Ethernet frames and drops any frames that violate the flow constraints for the corresponding connection. It then places the accepted frames in the egress queues based on their QoS class. A transmission scheduler transmits the frames in the egress queues according to their priorities.

UPSTREAM PRIORITY SCHEDULING AT THE SAS AND NDS

Intermediate NBAN devices, namely SASs and NDSs, play an important role in QoS management although they are not aware of individual connections or flows. For traffic flowing upstream (i.e., toward the router), each NDS

and SAS maintains one queue for each of the top three priority classes and per link queues for UBR traffic. Upstream NBAN frames received by an intermediate device are placed in the appropriate queue based on the QoS marking in their NBAN labels. The NBAN frames in these queues are taken up for transmission further upstream by a transmission scheduler that employs a combination of priority scheduling with a weighted round-robin service. Each QoS class has strict nonpreemptive priority over all classes at a lower priority level. NBAN frames within a given queue are served in a FIFO manner. NBAN frames belonging to the UBR class are taken up for transmission only if the queues associated with the higher-priority classes are empty. Since there are multiple queues corresponding to the UBR class, the scheduler uses a scheme based on a weighted round-robin service discipline to decide the relative order in which frames waiting in different UBR queues get served. The weights used in making these scheduling decisions are chosen to ensure that the subscriber populations attached to different parts of the network are treated equitably in the UBR throughput they receive, independent of the number of hops to the ONDS.

UPSTREAM LINK-BY-LINK FLOW CONTROL

Another critical feature of the NBAN QoS management scheme is link-by-link flow control of upstream traffic. All intermediate devices are aggregation points where the potential rate at which upstream traffic can converge at the device far exceeds the rate at which this traffic can be carried further upstream. In this scenario, there is a potential for significant frame losses even if one were to provide several megabits of expensive buffer memory. The link-by-link upstream flow control scheme implemented in the NBAN solves this problem, practically eliminating frame losses while keeping the buffer memory requirements to a minimum. In this scheme, each device periodically sends flow control signals (per QoS class) to each downstream device directly connected to it. The latter can send upstream traffic belonging to a particular QoS class only if the most recent flow control signal it received for that class permitted it to do so. The flow control bits contained in the F/P field in the NPHY frame (Fig. 5) are used to carry these signals.

The transmission scheduling discipline and

Intermediate NBAN devices, namely SASs and NDSs, play an important role in QoS management although they are not aware of individual connections or flows.

The transmission scheduling discipline and link-by-link flow control scheme implemented at intermediate devices enable the NBAN to meet its QoS objectives of ensuring low delays for high-priority traffic, fairness in the treatment of UBR traffic, and minimizing frame losses within the network.

link-by-link flow control scheme implemented at intermediate devices enable the NBAN to meet its QoS objectives of ensuring low delays for high-priority traffic and providing fairness in the treatment of UBR traffic.

DOWNSTREAM TRAFFIC

Downstream traffic is significantly easier to handle than upstream traffic due to the fact that there are no aggregation points within the NBAN.⁴ The only aggregation point for downstream traffic is at the router, which must provide some buffering and QoS handling for traffic destined to the NBAN. The router must be provisioned to prioritize and shape downstream traffic to the ONDS. Once this traffic enters the NBAN, it flows unimpeded until reaches its destination, with little likelihood of being dropped on the way. Consequently, intermediate NBAN devices ignore the QoS field in their handling of downstream traffic. The SAS, which connects a 1 Gb/s feeder cable to the destination BIU over a 100 Mb/s drop, is the only place where downstream traffic encounters rate adaptation. Limited FIFO buffering is provided for downstream rate adaptation at the SAS. Even with this simple handling of downstream traffic, the delays experienced by all frames have been found to be on the order of 1–2 ms.

CONCLUSIONS

The last mile problem is still the biggest hurdle in the delivery of high-bandwidth high-quality next-generation services to the large pool of residential and business subscribers. The next-generation broadband access network described herein innovatively turns the practically ubiquitous HFC network into a reliable high-bandwidth QoS-enabled access medium, and represents a cost-effective solution to this problem.

Reliability is of paramount importance in the CATV and telecommunication industries, and one of the primary objectives of the NBAN is to exceed industry-wide reliability goals while simultaneously providing new levels of bandwidth and service to the consumer. The NBAN enhances the conventional HFC architecture by converting every device in the plant into an IP-addressable managed device. In addition, the SAS includes an automatic switch bypass feature that circumvents the network element in the event of malfunction. Early experience shows that preventive maintenance and rapid fault isolation enabled by state-of-the-art network management and supervision dramatically improves the availability of the HFC plant beyond that of conventional HFC networks.

NBAN will enable high-bandwidth virtual private networking, remote storage, video on demand, and high-quality voice over IP. Once available, high-quality, reliable, and inexpensive broadband access will permit the deployment of many new services. Open access solutions, easily implemented in NBAN, will provide the ability for multiple third-party providers to deliver and bill for services, content, or applications via a shared HFC access network deployed by the cable operator. This will prove to be an incubator for the creation of new bandwidth-hungry applications.

⁴ On the downstream, only branching or disaggregation occurs at each NBAN device.

REFERENCES

- [1] DOCSIS, Radio Frequency Interface Specification, SP-RFI v1.1-106-001215, Cable Labs, Dec. 15, 2000, <http://www.cable-modem.com>
- [2] The Emergence of Integrated Broadband Cable Networks, Special Issue, *IEEE Commun. Mag.*, June 2001.
- [3] D. Raskin and D. Stoneback, *Broadband Return Systems for Hybrid Fiber/Coax Cable TV Networks*, Prentice Hall, 1998.
- [4] W. Ciciora, J. Farmer, and D. Large, *Modern Cable Television Technology*, Morgan Kaufman, 1999.
- [5] D. Bertsekas and R. Gallager, *Data Networks*, Prentice Hall, 1991.
- [6] J. Proakis, *Digital Communications*, McGraw Hill, 1995.

BIOGRAPHIES

SUBRA DRAVIDA is vice president of engineering at Narad Networks, responsible for product development. Prior to Narad, he was director of engineering at Cisco, responsible for building voice/data cable and DSL modems. Prior to Cisco, he was VP of engineering at MaxComm Technologies, a voice/data DSL company acquired by Cisco in 1999. Earlier, he worked at the Performance Analysis Department in Bell Laboratories, Holmdel, New Jersey. He holds over 30 patents in networking. He has a Ph.D. and an M.S. in electrical engineering from Rensselaer Polytechnic Institute (RPI) and a B.Tech from the Indian Institute of Technology (IIT), Chennai.

DEV GUPTA is the founder of Narad Networks. Before Narad, he founded two successful companies, MaxComm Technologies and Dagaz Technologies, which were acquired by Cisco. He also served as vice president of architecture and technology for Cisco's access products and held senior management positions at Bell Laboratories. He holds a B.S. in electrical engineering from IIT Kanpur, an M.S.E.E. from the University of Maine at Orono, and a Ph.D. from the University of Massachusetts at Amherst.

SANJIV NANDA (nandas@ieee.org) is director of systems engineering at Narad Networks. His group is responsible for overall system architecture and protocols for Narad's next-generation HFC products and systems. Prior to Narad, he spent over 10 years with the Performance Analysis Department at Bell Laboratories and at WINLAB, Rutgers University, working on many aspects of cellular system design, including system architecture, protocol design, system performance, and link performance. He received a B.Tech degree in electrical engineering from IIT, Kanpur, and M.S. and Ph.D. degrees from RPI.

KIRAN REGE is a senior member of technical staff at Narad Networks. He holds a B.Tech. from IIT, Bombay, an M.E. from the University of Florida, and a Ph.D. from the University of Hawaii, all in electrical engineering. Before joining Narad, he was with Bell Laboratories, Lucent Technologies, where he worked extensively on projects involving performance modeling and analysis of communication systems and networks, switching systems, operations support systems, and manufacturing systems.

JEROME D. STROMBOSKY (M '85) is a system architect at Narad Networks, where he is involved in the design of next-generation cable modem, optical, and wireless-based broadband access systems. In addition to his current position at Narad Networks, he has held various engineering and management positions at MIT Lincoln Laboratory, Lucent Technologies, and Analog Devices Semiconductor, where he was involved in a myriad of circuit design and systems engineering activities. At MIT Lincoln Laboratory, he was principal investigator and project leader in the development of advanced signal processing IC technology and massively parallel/pipelined systems for use in communication and radar signal processing applications. He holds B.S. and M.S. degrees in electrical engineering, and graduated from Worcester Polytechnic Institute with a concentration in communications and signal processing.

MANAS TANDON (tandonm@naradnetworks.com) has a B.Tech in electrical engineering from IIT, Kanpur, where he graduated at the top of the class. Currently, he is manager of systems engineering at Narad Networks, where he has been responsible for developing the system functional specifications and protocols for data and control messaging. Prior to Narad, he worked in the DSL Business Unit of Cisco Systems, where he worked on switching gateways and home networking products.