Bringing Home the Bandwidth: Optimal HFC Access Architectures for New Builds

Chris Bonang
Harmonic, Inc.

Abstract

As consumer demand for services such as high-speed internet access, multiple voice lines, and video on demand continues to grow, HFC networks are increasingly being recognized worldwide as the only single, proven, residential access network that can deliver the enormous bandwidth necessary to supply these services. Any network built today must be "future-proof", capable of scaling to whatever amount of bandwidth that will be necessary to support any future services and applications that may appear in the next 10-15 years. Even in existing HFC systems with extensive legacy equipment, it is possible to apply solutions developed for new builds when planning system upgrades.

In this paper, optimal HFC designs for such "greenfield" builds are reviewed. The expected growth in bandwidth demand from cable subscribers over the next several years is first reviewed. Deep fiber optical node segmentation schemes which will allow highly scalable bandwidth delivery, and at the same time minimize or eliminate RF actives from the coax plant are then discussed. It is then shown how DWDM can provide significant cost savings and deployment convenience by reducing costly hub real estate and minimizing or eliminating expensive SONET transmission systems. Cost-effective implementation of a combination of digital transmission and DWDM in the return path of such deep fiber architectures is examined. Finally, the discussion of digital return is extended to include the possibility of demodulation, and even reduced CMTS functionality, in the node.

BANDWDITH DEMAND

Over the past few years, the exploding popularity of the internet has revealed a bottleneck in the local access network. Uptake of internet services has been delayed by the frustratingly slow speed of 14 kbps dial-up modems. Recently, new technologies such as cable modems and xDSL have improved peak available downstream residential bandwidth to a bearable 500 kbps. Although the associated customer premises equipment (CPE) is capable of much higher bandwidths, the shared access networks themselves typically do not yet support them. Network operators in the process of deploying new systems do not want to face any such service bandwidth limitations, and so are typically choosing hybrid fiber-coax (HFC) architectures, due to their unparalleled bandwidth delivering capabilities – 800 MHz or 5 Gbps per laser transmitter. The key to making this enormous bandwidth available to subscribers is dedicating narrowcasting services to transmitters and optical nodes serving small service areas in a scalable fashion. Before discussing the details of this process, it is first reasonable to characterize the services whose demand is giving incentive to cable operators to build new HFC networks or upgrade their existing ones.

Peak data rates for current typical internet usage is more than sufficient at 1 Mbps. The increasing popularity of applications like streaming video probably add another 1 Mbps to the potential bandwidth demand. The expanding prevalence of digital cameras, both for still pictures and video, will increase both up- and downstream bandwidth demand, as individuals exchange such material over the internet. This bandwidth forecast is probably conservative, since internet backbone traffic is forecasted to increase by x10 per year, and this increase must also then be reflected in the access network. In addition, the improving economics of video-on-demand (VOD) delivery systems will likely soon result in 1 or 2 movies per household at 4 Mbps per movie commonly being purchased. Based on this analysis, the peak rate to a single subscriber could be as much as 10 Mbps in the next 1-2 years. The services mentioned above are particularly downstream intensive, so a traffice asymmetry of 10:1 downstream to upstream is assumed. Peak upstream bandwidth might be on the order of 1 Mbps in this case. New services such as video conferencing, interactive gaming, and future services that have yet to be conceived must also be taken into account. Consumers will likely use as much bandwidth as is made available to them. Therefore, network operators must choose architectures that will scale to support this almost unlimited bandwidth demand.

NODE SEGMENTATION

The number of subscribers served by an optical node for a typical HFC system has steadily decreased over the last decade. Formerly, 1000-2000 home nodes were the norm, but today new builds and upgrades are more likely to be in the 100-500 homes per node range. Pushing fiber deeper into the network results in improved bandwidth, performance and reliability, particularly due to the reduction RF actives from the coax plant. Completely eliminating RF amplifiers from the system is an excellent goal, since the system would have a 33-50% less active components resulting in greater reliability and reduced power consumption (~40%). Completely passive coax networks are possible for cases of very high subscriber density – multiple dwelling units (MDU's) with +200 subscribers per mile. However, for more typical densities of 100 subs/mile, it is necessary to segment down to approximately 35 homes per optical node in order to create a completely passive coax network. This assumes a four output node with 51-53 dBmV outputs. Unfortunately, for most cable plant, it is difficult or impossible to effectively use all four outputs from the node, and there is not a tremendous cost difference between nodes utilizing 2, 3 or 4 ports. Although it is generally better to push fiber as deep as possible, 35 homes per node is probably not a cost-effective alternative at this time. A practical trade-off is to employ a four-output node passing 100 subscribers and add one line extender to each node output. This provides plenty of RF level at the home, and potentially anywhere from 20 to 50 Mbps of dedicated downstream bandwidth. With only a single line extender in each path, high performance based primarily on the optical link(s) is easily obtained, and higher order modulation formats like 256-QAM and even 1024-QAM can be supported.

DEEP FIBER ARCHITECTURES

Traditional Double-hop HFC Architecture

A traditional HFC architecture is shown in Figure 1. Broadcast analog video is transmitted by an externally modulated 1550 nm transmitter from the head end to the hub. A redundant transmitter and path are included for signal protection. Narrowcasting services (data, voice, and possibly digital video) are typically transmitted via a SONET link, although other options include ATM, IP or proprietary baseband digital transport are also available. At the hub, the digital narrowcast signals are processed for transmission over the HFC network by the appropriate interface unit. A cable modem termination system (CMTS) converts the downstream data signal to QAM, and QPSK demodulates the upstream signal, as well as supplying media access control. It is assumed that any CMTS/Cable modems deployed in the future will be based on the DOCSIS standard. The host digital terminal (HDT) performs the same tasks for cable telephony. Video servers are shown at the head end, but could also be located in the hub to provide digital video and VOD service. After processing and QAM modulation, the 550-860 MHz narrowcast signals are combined with the 50-550 MHz broadcast analog signal and fed to 1310 nm transmitters, which in turn feed multiple or individual optical nodes. Each optical node could serve anywhere from 100-1000 subscribers; for the deep fiber architecture, we will assume 100 subs/node. Since each hub serves on the order of 50,000 subs, the hubs must be relatively large buildings in order to house all of the equipment and fiber connections. The cost of the building and real estate in metropolitan areas can be as much as $2M, assuming that a suitable site can be located.

Several drawbacks exist for operators deploying traditional double-hop architectures in order to serve today's bandwidth-hungry subscribers. Aside from the difficulty and expense of locating and building a large hub site, there are the functional problems of redundancy and sheer numbers of fibers necessary to feed ~500 nodes from each hub. For services such as telephony, most operators require both equipment and fiber path redundancy down to at least the 500-1000 home level. This is difficult to achieve in the double-hop architecture unless there are redundant receivers in each node fed by redundant fiber rings and transmitters. This arrangement provides redundancy down to the 100 home level, which might be considered overkill, and is definitely expensive. A star architecture from the hub to the nodes would mitigate the redundancy problem, but would be prohibitively expensive in terms of fiber cable, since there would be a unique fiber cable to each node.

An additional disadvantage of the double-hop architecture is the high cost of SONET or proprietary digital transport as bandwidth demand increases beyond 1 Mbps per home. For a hub serving 50,000 subscribers, providing just 2 Mbps to each requires 100 Gbps from the head end to the hub. This corresponds to 10 OC-192 terminals in both the head end and the hub, a great expense in both equipment and space. Also, much of the SONET bandwidth is

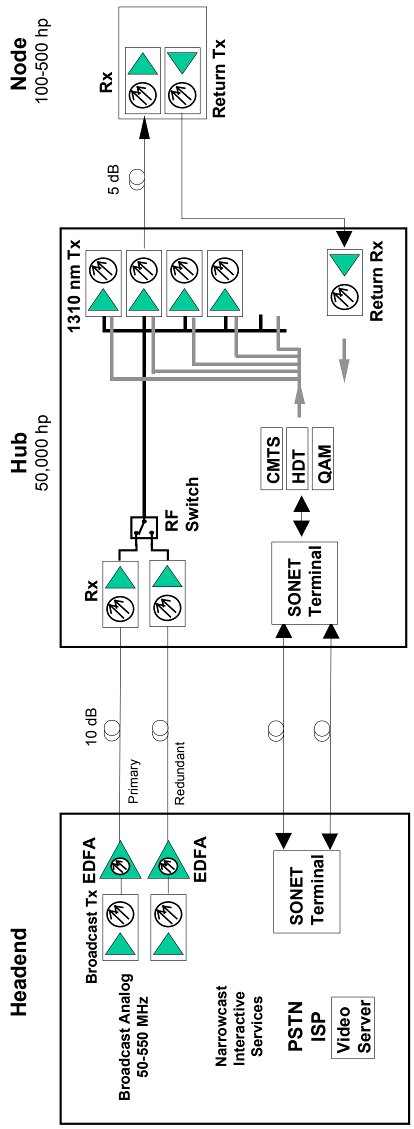

Figure 1: Traditional HFC Architecture

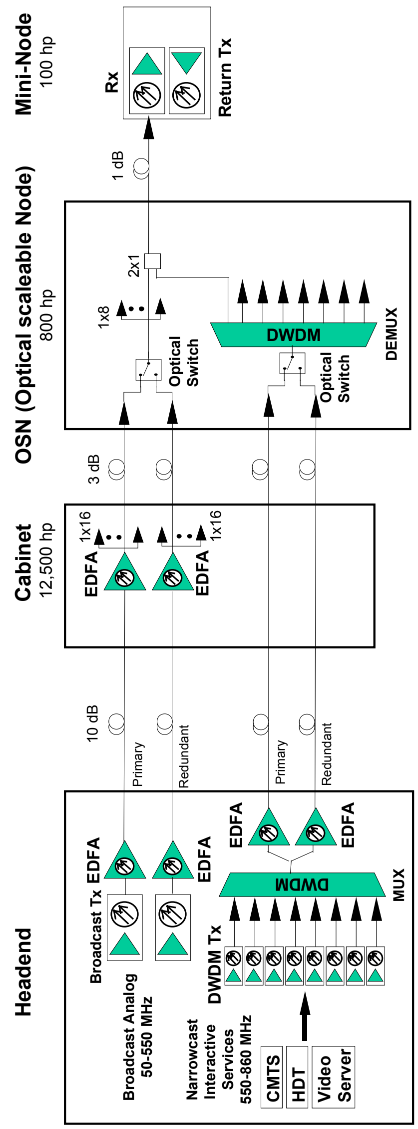

Figure 2: DWDM Deep Fiber HFC Architecture - Forward Path

wasted, due to the 10:1 asymmetry of the downstream and upstream traffic.

DWDM Deep Fiber Architecture

A Dense Wavelength Division Multiplexing (DWDM) deep fiber architecture is shown in Figure 2. As before, the broadcast analog video is redundantly transmitted via an externally modulated 1550 nm transmitter from the head end to a hub. At the hub, the broadcast signal is amplified and split to feed optically scaleable nodes (OSN's). The signal at the OSN, which may be either strand- or cabinet-mounted, is further split to feed 100 home mini-nodes. The narrowcast signals are transmitted via directly modulated DWDM transmitters, whose wavelengths correspond to those of the ITU grid. The wavelengths are multiplexed together at the head end, optically amplified by an erbium doped fiber amplifier (EDFA), pass through the hub site, and are demultiplexed at the OSN. In this ultimate configuration, each wavelength serves a single mini-node, providing 300 MHz of narrowcasting QAM channels (~2 Gbps or 20 Mbps per subscriber). However, the system can be scaled such that initially each wavelength is shared amongst eight mini-nodes. A simple optical splitter is deployed in the OSN instead of a DWDM demultiplexer. As bandwidth demand increases, additional DWDM transmitters are added at the head end, and DWDM demultiplexers replace the optical splitters in the OSN.

In the configuration shown in Figure 2, each OSN serves eight mini-nodes. That limitation presently exists primarily due to the temperature-dependence of DWDM demultiplexers. Wavelength spacing of 200 GHz is necessary for the extreme temperature conditions of a strand-mounted environment. However, next-generation temperature-hardened DWDM couplers will soon be available which allow 100 GHz spacing, and therefore 16 wavelengths can be transmitted to each OSN. Each OSN can then serve 16 mini-nodes, making the system somewhat more cost-effective.

The broadcast and demultiplexed narrowcast signals are combined at the OSN and transmitted over the same fiber to a single receiver in the mini-node. The signals may be combined via a simple 2x1 optical coupler, or by a 2x1 DWDM multiplexer, depending on the available loss budget. If a DWDM mux is necessary due to link budget considerations, then another option is to leave the broadcast and narrowcast signals separate, and transmit them over individual fibers to separate receivers in the mini-node. The cost of a 2x1 DWDM mux and 10 km of fiber is greater than the cost of an additional receiver. In addition, avoiding the 2x1 combiner more easily enables the equipment associated with a 1x16 split to be packed into the OSN, as opposed to only serving 8 mini-nodes from each OSN. An additional advantage of this approach is that it eliminates potential problems associated with the CSO and CTB from the narrowcast signals interfering with the broadcast signal when using a single receiver.

Due to the fact that the forward path signals are passively transmitted through the hub, it is easy to replace the large hub site with a small, inexpensive cabinet or vault. Such cabinets are much easier for which to find locations and only cost on the order of $20K. In

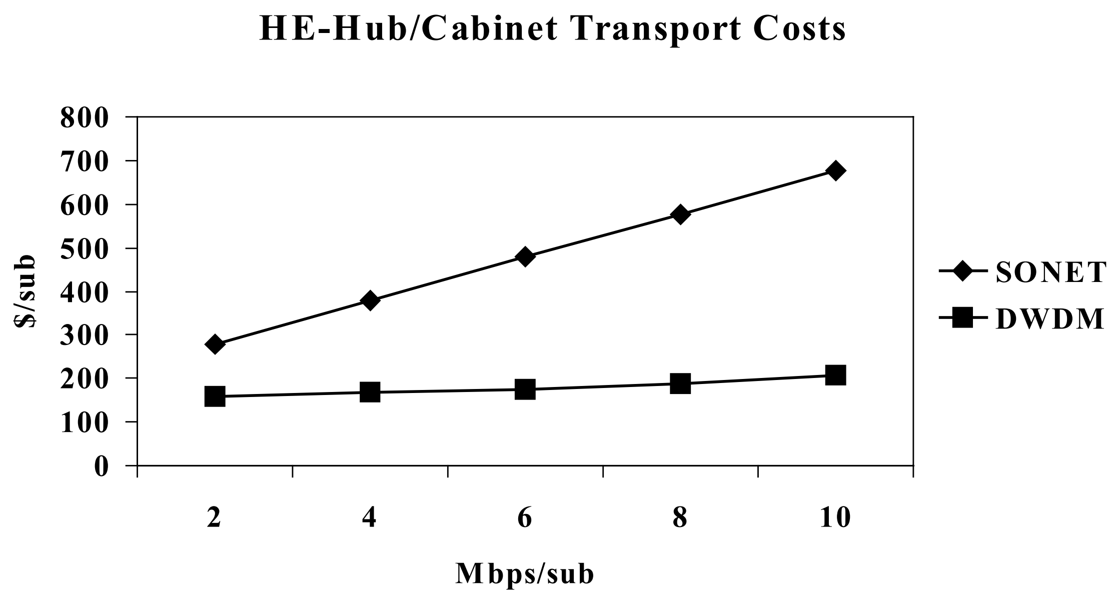

Figure 2, the 50,000 subs formerly served by a single large hub are served by four cabinets. Placing all of the equipment at the head end has the additional advantages of lowering operational costs and permitting less total equipment to be deployed in the initial stages. In the traditional architecture, it is necessary to locate at least one CMTS, HDT and possibly a video server at each hub, regardless of how limited the demand. With the CMTS and HDT pulled back to the head end, this equipment can be shared amongst multiple cabinets/hubs when demand is very low.

The DWDM deep fiber system shown in Figure 2 does not require expensive SONET or other digital transport systems. As shown in Figure 3, when combined with the use of cabinets rather than expensive Hubs, this results in a very large savings as narrowcast bandwidth demand to the home increases into the several Mbps range. The traditional architecture cost increases rapidly with increasing bandwidth demand, but the DWDM architecture only requires additional, and relatively inexpensive, DWDM transmitters. Similar to SONET, the DWDM architecture also provides path and equipment protection via redundant optical amplifiers and optical switches. However, for certain high-priority services, such as telephony, many operators are more comfortable with the extensive protection and monitoring capabilities of SONET. For these operators, a hybrid approach is possible. The low bandwidth, high-priority services like telephony can be transmitted to the hub/cabinet via SONET, while the high bandwidth, nonlifeline services such as VOD can be transmitted via DWDM.

Return Path

Traditionally, multiple return paths are simply combined such that the upstream signal from thousands of subscribers is fed to a single CMTS or HDT. This results in very high noise levels, and limits the number of subscribers who can access the service. As service penetration increases, operators must be able to segment the upstream to serve much smaller numbers of subscribers. Several methods exist for segmenting the return path, thus providing dedicated upstream bandwidth to customers. Figure 4 illustrates the pure DWDM option, which is basically the mirror image of the DWDM downstream. ITU return path transmitters in the mini-node transmit back to the OSN, where the signal is DWDM muxed with the signals from the other mini-nodes served by the OSN. Since the signals are 5-40 MHz analog, it is necessary to amplify with an EDFA before transmitting back to the head end in order to maintain acceptable performance. At the head end, the signals are demultiplexed and fed to individual return path receivers. The DWDM upstream option provides excellent segmentation. But it is not scalable, since the required ITU lasers in every mini-node, and the cabinet EDFA's, combine to make the initial system deployment relatively expensive. The system does provide excellent return path bandwidth of up to approximately 100 Mbps, assuming 16-QAM modulation. This corresponds to 1 Mbps peak rate per subscriber, which may be more than necessary in the early stages of deployment.

A more scalable and less expensive return path option is to combine digital transmission with DWDM. As shown in Figure 5, the 5-40 MHz upstream signal is transmitted by a 1310 nm laser from the mini-node to the OSN. The laser could be either a relatively low-inexpensive uncooled distributed feedback (DFB) laser or a very low-cost Fabry-Perot (FP) laser. The choice between the two depends on how much combining the operator plans to do. FP lasers are more noisy, particularly when no signal is driving them. DFB lasers therefore may be necessary when combining many return path segments, and when high priority services like telephony are offered.

At the OSN, the signal is received and combined with three other upstream signals. The combined 5-40 MHz signals are then digitized by a 10 bit sampling A/D converter. This results in a baseband digital signal of approximately 1 Gbps. This signal is then time division multiplexed (TDM) with the digitized signal from four other combined receivers and transmitted back to the cabinet via a 2.5Gbps ITU transmitter. At the cabinet, the signals are DWDM muxed with other return path wavelengths and transmitted to the head end. Because the signals are in digital format, an EDFA is not necessary. At the head end, the signals are demuxed, converted back to analog format, and fed to the appropriate CMTS, HDT or VOD controller. This system uses fewer ITU lasers than the pure DWDM return option, and no EDFA's, so it is more cost-effective. However, in the initial deployment shown, it only permits segmentation to the 400 home level, which corresponds

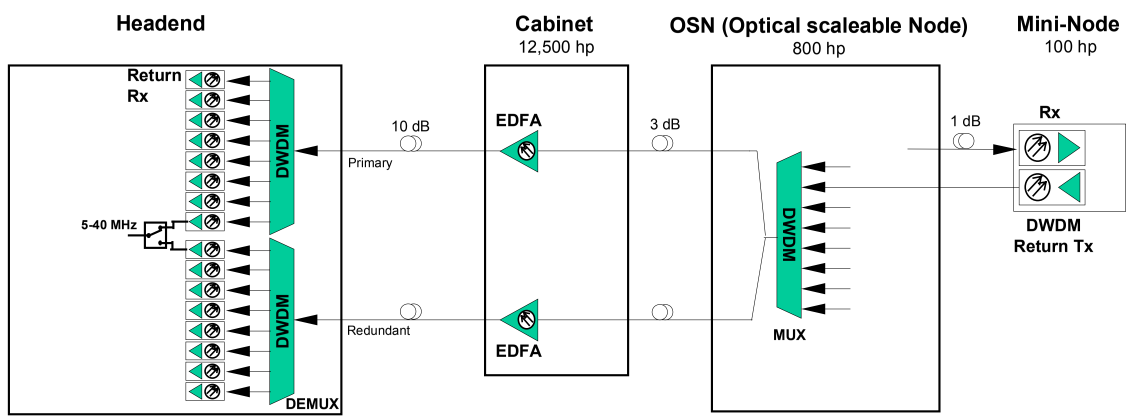

Figure 4: Deep Fiber HFC Architecture - Pure DWDM Return Path

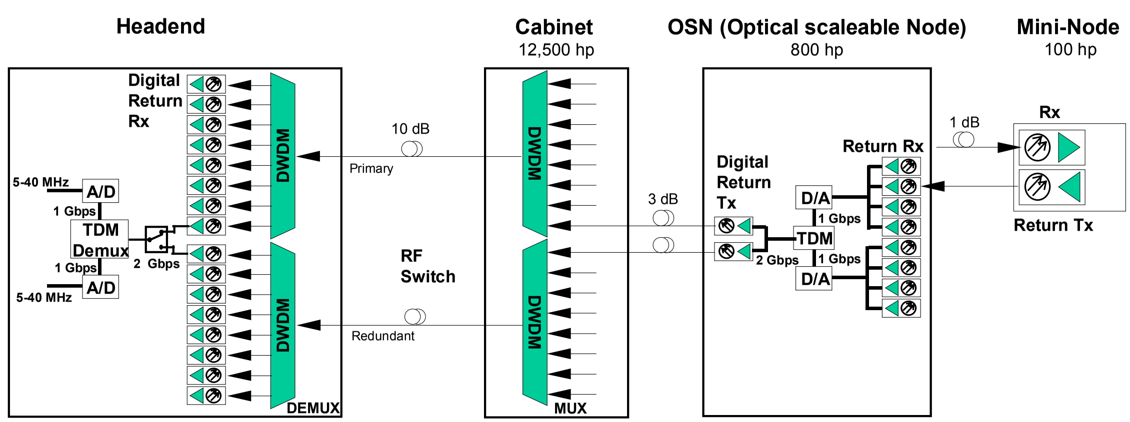

Figure 5: Deep Fiber HFC Architecture - Digital/DWDM Return Path

to a peak upstream rate of 250 Kbps per sub. In order to provide the same segmentation and bandwidth of the pure DWDM option, every return path segment must be digitized separately without combining. Either additional digital 2.5 Gbps ITU transmitters will be necessary, or the eight digitized return segments must all be multiplexed together before transmission. However, this would require a 10 Gbps transmitter, which must be externally modulated, and is probably cost prohibitive compared to a common, directly modulated 2.5Gbps transmitter.

An final return path option is to employ a more efficient form of digitization of the QPSK and 16-QAM signals in the 5-40 MHz return band. Digital sampling of the return path signal is somewhat effective, but very inefficient. The 5-40 MHz waveform is digitized to produce a 1 Gbps signal, despite the fact the maximum useful information carried by the signal, assuming 16-QAM modulation, is only 100 Mbps. A possible solution to the digital upstream efficiency problem under development is to remotely demodulate the DOCSIS QPSK or 16-QAM upstream signal by moving some of the functionality of the CMTS from the head end to the OSN or mini-node. Utilizing such a technique makes the return path more scalable. A further possible step is to locate an entire reduced-functionality CMTS in the mini-node. The device consists of the PHY portion of a regular CMTS (QAM modulator, upconverter, QPSK demodulator) and a rudimentary MAC layer. A two-way ethernet switching fabric transmits the baseband digital signal from the head end to the mini-node.

SUMMARY

Utilizing a Dense Wavelength Division Multiplexing (DWDM) deep fiber architecture overcomes many of the drawbacks associated with traditional HFC architectures. The system is capable of providing enormous amounts of dedicated bandwidth. The architecture is also completely scalable, and cost-effective when compared with traditional dual-hop HFC architectures. The return path uses a combination of digitization and DWDM to provide segmentation and scalability. Next generation technologies may distribute demodulation presently associated with the CMTS out into the node.